



# **Citrix Virtual Apps and Desktops service**

## Contents

<b>Citrix Virtual Apps and Desktops service</b>	<b>3</b>
<b>What's new</b>	<b>13</b>
<b>Known issues</b>	<b>49</b>
<b>Deprecation</b>	<b>51</b>
<b>System requirements</b>	<b>52</b>
<b>Limits</b>	<b>57</b>
<b>Technical security overview</b>	<b>60</b>
<b>Technical security overview for Citrix Managed Azure</b>	<b>67</b>
<b>Delivery methods</b>	<b>80</b>
<b>Get started: Plan and build a deployment</b>	<b>84</b>
<b>Sign up for the service</b>	<b>91</b>
<b>Set up resource locations</b>	<b>94</b>
<b>Microsoft Azure Resource Manager virtualization environments</b>	<b>97</b>
<b>Citrix Hypervisor virtualization environments</b>	<b>120</b>
<b>Microsoft System Center Virtual Machine Manager virtualization environments</b>	<b>122</b>
<b>VMware virtualization environments</b>	<b>125</b>
<b>Amazon Web Services virtualization environments</b>	<b>132</b>
<b>Nutanix virtualization environments</b>	<b>148</b>
<b>Google Cloud Platform virtualization environments</b>	<b>149</b>
<b>Scale and size considerations for Cloud Connectors</b>	<b>169</b>
<b>Create and manage connections</b>	<b>184</b>
<b>Install VDAs</b>	<b>194</b>
<b>Install VDAs using the command line</b>	<b>211</b>

<b>Create machine catalogs</b>	<b>219</b>
<b>Manage machine catalogs</b>	<b>244</b>
<b>Quick Deploy</b>	<b>254</b>
<b>Get started with Quick Deploy</b>	<b>258</b>
<b>Create catalogs using Quick Deploy</b>	<b>261</b>
<b>Manage catalogs in Quick Deploy</b>	<b>271</b>
<b>Azure subscriptions in Quick Deploy</b>	<b>283</b>
<b>Images in Quick Deploy</b>	<b>288</b>
<b>Network connections in Quick Deploy</b>	<b>299</b>
<b>Users and authentication in Quick Deploy</b>	<b>316</b>
<b>Remote PC Access in Quick Deploy</b>	<b>322</b>
<b>Monitor in Quick Deploy</b>	<b>331</b>
<b>Troubleshoot in Quick Deploy</b>	<b>337</b>
<b>Quick Deploy reference</b>	<b>341</b>
<b>Create delivery groups</b>	<b>351</b>
<b>Manage delivery groups</b>	<b>357</b>
<b>Create application groups</b>	<b>376</b>
<b>Manage application groups</b>	<b>383</b>
<b>Remote PC Access</b>	<b>389</b>
<b>Remove components</b>	<b>400</b>
<b>User personalization layer</b>	<b>401</b>
<b>Upgrade</b>	<b>419</b>
<b>Migrate to cloud</b>	<b>422</b>
<b>Print</b>	<b>475</b>

<b>HDX</b>	<b>476</b>
<b>Adaptive transport</b>	<b>486</b>
<b>Rendezvous protocol</b>	<b>495</b>
<b>Citrix ICA virtual channels</b>	<b>499</b>
<b>Double hop in Citrix Virtual Apps and Desktops</b>	<b>509</b>
<b>Devices</b>	<b>511</b>
<b>Generic USB devices</b>	<b>513</b>
<b>Mobile and touch screen devices</b>	<b>513</b>
<b>Serial ports</b>	<b>516</b>
<b>Specialty keyboards</b>	<b>521</b>
<b>TWAIN devices</b>	<b>523</b>
<b>Webcams</b>	<b>523</b>
<b>WIA devices</b>	<b>524</b>
<b>Graphics</b>	<b>524</b>
<b>HDX 3D Pro</b>	<b>526</b>
<b>GPU acceleration for Windows multi-session OS</b>	<b>527</b>
<b>GPU acceleration for Windows single-session OS</b>	<b>529</b>
<b>Thinwire</b>	<b>533</b>
<b>Text-based session watermark</b>	<b>539</b>
<b>Multimedia</b>	<b>541</b>
<b>Audio features</b>	<b>544</b>
<b>Browser content redirection</b>	<b>552</b>
<b>HDX video conferencing and webcam video compression</b>	<b>560</b>
<b>HTML5 multimedia redirection</b>	<b>564</b>



<b>Optimization for Microsoft Teams</b>	<b>567</b>
<b>Monitor, troubleshoot, and support Microsoft Teams</b>	<b>589</b>
<b>Windows Media redirection</b>	<b>596</b>
<b>General content redirection</b>	<b>597</b>
<b>Client folder redirection</b>	<b>597</b>
<b>Host to client redirection</b>	<b>598</b>
<b>Local App Access and URL redirection</b>	<b>602</b>
<b>Generic USB redirection and client drive considerations</b>	<b>610</b>
<b>Policies</b>	<b>620</b>
<b>Manage</b>	<b>621</b>
<b>Autoscale</b>	<b>622</b>
<b>Schedule-based and load-based settings</b>	<b>633</b>
<b>Dynamic session timeouts</b>	<b>652</b>
<b>Restrict Autoscale (cloud burst)</b>	<b>654</b>
<b>Dynamic machine provisioning</b>	<b>660</b>
<b>Force user logoff</b>	<b>666</b>
<b>Cloud Health Check</b>	<b>668</b>
<b>Configuration logging</b>	<b>700</b>
<b>Delegated administration</b>	<b>704</b>
<b>Load balance machines</b>	<b>717</b>
<b>Local Host Cache</b>	<b>717</b>
<b>Manage security keys</b>	<b>727</b>
<b>Scale and size considerations for Local Host Cache</b>	<b>743</b>
<b>Use Search in the Full Configuration management interface</b>	<b>761</b>

<b>Virtual IP and virtual loopback</b>	<b>763</b>
<b>Sessions</b>	<b>766</b>
<b>Tags</b>	<b>773</b>
<b>Zones</b>	<b>783</b>
<b>Licenses</b>	<b>791</b>
<b>User access</b>	<b>792</b>
<b>Monitor</b>	<b>795</b>
<b>Site Analytics</b>	<b>796</b>
<b>Alerts and notifications</b>	<b>805</b>
<b>Filter data to troubleshoot failures</b>	<b>815</b>
<b>Monitor historical trends across a site</b>	<b>817</b>
<b>Monitor Autoscale-managed machines</b>	<b>823</b>
<b>Troubleshoot deployments</b>	<b>826</b>
<b>Troubleshoot applications</b>	<b>826</b>
<b>Application probing</b>	<b>830</b>
<b>Desktop probing</b>	<b>835</b>
<b>Troubleshoot machines</b>	<b>840</b>
<b>Troubleshoot user issues</b>	<b>849</b>
<b>Diagnose session startup issues</b>	<b>850</b>
<b>Diagnose user logon issues</b>	<b>855</b>
<b>Shadow users</b>	<b>862</b>
<b>Send messages to users</b>	<b>863</b>
<b>Resolve application failures</b>	<b>863</b>
<b>Restore desktop connections</b>	<b>865</b>

<b>Restore sessions</b>	<b>865</b>
<b>Run HDX channel system reports</b>	<b>866</b>
<b>Reset a user profile</b>	<b>866</b>
<b>Feature compatibility matrix</b>	<b>869</b>
<b>Delegated administration and monitoring</b>	<b>873</b>
<b>Data granularity and retention</b>	<b>877</b>
<b>Citrix Virtual Apps and Desktops service for Citrix Service Providers</b>	<b>882</b>
<b>Citrix Gateway service</b>	<b>887</b>
<b>SDKs and APIs</b>	<b>887</b>

## Citrix Virtual Apps and Desktops service

July 27, 2021

### Introduction

Citrix Virtual Apps and Desktops provides virtualization solutions that give IT control of virtual machines, applications, and security while providing anywhere access for any device. End users can use applications and desktops independently of the device's operating system and interface.

Using the Citrix Virtual Apps and Desktops service, you can deliver secure virtual apps and desktops to any device, leaving most of the installation, setup, and upgrades to Citrix. You maintain complete control over applications, policies, and users while delivering the best user experience on any device.

The service allows you to manage on-premises data center and public cloud workloads together in a hybrid deployment. You can connect to public clouds Microsoft Azure, Amazon Web Services (AWS), and Google Cloud, plus on-premises hypervisors such as Citrix Hypervisor, Microsoft Hyper-V, Nutanix AHV, and VMware vSphere. The hybrid, multi-cloud approach gives you the flexibility to deploy different applications in different resource locations worldwide.

The service offers several ways to deliver apps and desktops.

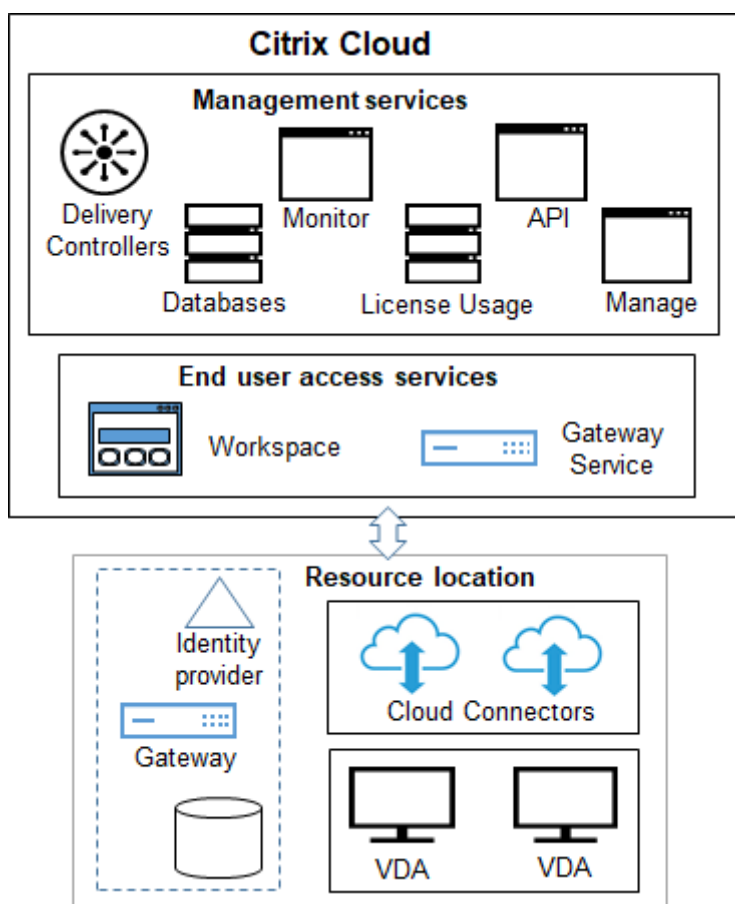
- [Delivery methods](#) describes the primary ways, with use-cases and pros/cons.
- [Delivery models](#) lists more choices, and also offers VDI model comparisons.

Citrix Managed Azure further simplifies the deployment of virtual apps and desktops. With Citrix Managed Azure, Citrix also manages the hosting of Azure workloads.

[Learn more about the advantages of using this service.](#)

### Site overview

The following graphic shows the services and components that Citrix administrators work with in a Citrix Virtual Apps and Desktops service production deployment (also known as a site).



As shown in the graphic, Citrix manages the user access and management services and components in Citrix Cloud. The applications and desktops that you deliver to users reside on machines in one or more resource locations. In a Citrix Virtual Apps and Desktops service deployment, a resource location contains components from the access layer and resource layers. Each resource location is considered a [zone](#).

If you recently migrated from on-premises Citrix Virtual Apps and Desktops, you'll see that the service eliminates most of the component setup work required in an on-premises deployment.

### Components and services managed by Citrix

- **Delivery Controllers:** The Citrix Virtual Apps and Desktops service provides the functionality to load balance applications and desktops, authenticate users, and broker or prioritize connections directly from the cloud, without the need to manage Delivery Controllers, as with on-premises Citrix Virtual Apps and Desktops.
- **Databases:** Site configuration, monitoring, and configuration logging data is stored by the cloud service, eliminating the SQL database requirement of the on-premises Citrix Virtual Apps and Desktops product.

- **Licensing:** Manages licenses and provides [usage statistics].
- **Management interfaces:** See Management interfaces. Many tasks are also available in [service APIs](#).
- **Monitor interface:** The [Monitor](#) interface enables IT support and help desk teams to monitor an environment, troubleshoot issues before they become critical, and perform support tasks for end users. Displays include:
  - Real-time session data from the Broker Service in the Controller, which includes data from the broker agent in the Virtual Deliver Agent (VDA).
  - Historical data from the Monitor Service in the Controller.
  - Data about HDX traffic (also known as ICA traffic).
- **Cloud Connectors:** A Cloud Connector is the communications channel between the components in the Citrix Cloud and components in the resource location. In the resource location, the Cloud Connector acts as a proxy for the Delivery Controller in Citrix Cloud.

Every resource location contains at least one Cloud Connector. Two or more Cloud Connectors are recommended for redundancy.

- When using Full Configuration to provision machines, you first install Cloud Connectors from the Citrix Cloud console. For details, see [Cloud Connectors](#).
- When using Quick Deploy to provision Azure machines, Citrix creates the resource location and Cloud Connectors for you when you create a catalog.

After Cloud Connectors are installed, Citrix manages and updates them. The only tasks handled by the customer are Cloud Connector Windows updates and patching.

## Management interfaces

From the **Manage** tab of the service, you can select the following interfaces.

### Full Configuration

From the **Manage > Full Configuration** interface, you can:

- [Create and manage connections](#) to hosts.
- [Create](#) and [manage](#) catalogs of machines that contain apps and desktops you deliver to your users.
- [Create](#) and [manage](#) delivery groups (and optionally, application groups).
- Create and manage [Citrix policies](#) that affect the use and behavior of HDX technologies and features, plus site-level management. This includes policy settings for sessions, adaptive transport, devices, graphics, multimedia, content redirection, and VDAs.

- Customize [delegated administration](#) to create role-based administrators who have specific scopes of authority.
- Manage the [Autoscale](#) feature to proactively power manage machines that deliver apps and desktops.
- [Load balance machines](#)
- [Run health checks](#) on your VDAs to identify potential issues and fix suggestions.
- [Display configuration log content](#) to see when configuration changes and other administrative activities occurred, and who initiated them.

### Quick Deploy

From the **Manage > Quick Deploy** interface, you can easily deploy and manage Microsoft Azure workloads that use either a Citrix Managed Azure subscription or your own Azure subscription. For more information, see [Quick Deploy](#) and Citrix Managed Azure. From Quick Deploy, you can:

- [Create](#) and [manage](#) catalogs.
- [Create and customize](#) images, either from various Citrix prepared images, or from images you import from your Azure subscription.

For more information, see [Quick Deploy](#).

### Environment Management

From the **Environment Management** interface, you can use intelligent resource management and Profile Management technologies to deliver the best possible performance, desktop logon, and application response times. For more information, see [Workspace Environment Management](#).

### Components and technologies managed by the customer

- **Citrix Gateway:** When users connect from outside the corporate firewall, Citrix Virtual Apps and Desktops can use Citrix Gateway technology to secure these connections with TLS. The Citrix Gateway or VPX virtual appliance is an SSL VPN appliance deployed in the DMZ. It provides a single secure point of access through the corporate firewall.

Citrix installs and manages the Citrix Gateway service in Citrix Cloud. You can also optionally install Citrix Gateway in resource locations.

- **Active Directory:** Active Directory is used for authentication and authorization. It authenticates users and ensures that they are getting access to appropriate resources. A subscriber's identity defines the services to which they have access in Citrix Cloud. This identity comes from Active Directory domain accounts provided from the domains within the resource location.

- **Identity Provider (IdP):** The IdP is the final authority for the user's identity. Supported IdPs include: on-premises Active Directory, Active Directory plus token, Azure Active Directory, Citrix Gateway, and Okta. For more information, see:
  - [Workspace Identity](#)
  - [Identity and access management](#)
- **Virtual Delivery Agents (VDAs):** Each physical or virtual machine that delivers resources (applications and desktops) must have a Citrix VDA installed on it. VDAs establish and manage the connection between the machine on which it's installed and the user device, and apply policies that are configured for the session.

The VDA registers with a Delivery Controller, using a Cloud Connector in the resource location as a proxy.

Several VDA types are available:

- VDAs for Windows multi-session operating systems allow multiple users to connect to the machine at one time. This VDA type is usually installed on Windows servers.
- VDAs for Windows single-session operating systems allow one user to connect to a machine at a time. This VDA type is usually used for VDI.

A core version of this VDA type is available for use with the Remote PC Access feature. It contains a subset of the features in the full single-session VDA.
- Linux VDAs support virtual apps and desktops based on an RHEL, CentOS, SUSE, or Ubuntu distribution.

Throughout this service's documentation, "VDA" often refers to the agent and the machine on which it is installed.

- **Hypervisors and cloud services:** In most production sites, the app and desktop instances (workloads) that you make available (publish) to your users are "hosted" by a [supported hypervisor or cloud service](#). (The Remote PC Access feature is usually used with physical machines. Therefore, it does not use hypervisors or cloud services for machine provisioning.)
  - When using the Full Configuration interface, you create a connection to a supported host hypervisor or cloud service. Then from Full Configuration, you use an image (created through that host) to create a catalog of machines that contain the app and desktop instances. Then you create a delivery group. Citrix provides many tools to simplify and facilitate how these session hosts are built and maintained.
  - When using Quick Deploy to deliver Azure workloads, you only need to create the catalog. Although you can use your own Azure subscription when creating the catalog, using a Citrix Managed Azure subscription eliminates your need to manage the host, too.



The app and desktop instances that you publish can be on-premises, hosted in public clouds, or in a hybrid mixture of both.

- **Citrix StoreFront:** [Citrix StoreFront](#) is the predecessor to the cloud-hosted Citrix Workspace. It is used as the web interface for access to applications and desktops.

You can optionally install StoreFront servers in resource locations. Having local stores can help deliver apps and desktops during network outages. The [Local Host Cache](#) feature requires a customer-managed StoreFront in each resource location.

See [User access](#) for considerations for using StoreFront in a service environment.

## Objects you configure to deliver desktops and applications

You configure the following items to deliver apps and desktops in a production environment.

- **Host connection:** A host connection (mentioned earlier) helps enable communication between components in the control plane (Citrix Cloud) and VDAs in a resource location. Connection specifications include:
  - The address and credentials to access the host
  - The storage method to use, and the machines to use for storage
  - Which network the VMs can use

Remember: When using Quick Deploy, you don't have to create a connection. And if you use Citrix Managed Azure, Citrix manages the hosting, as well.

- **Catalog:** In the Full Configuration and Monitor interfaces, catalogs are called “machine catalogs.”

A catalog is a collection of virtual or physical machines that have the same operating system type (for example, Windows multi-session, Ubuntu single-session).

When creating a catalog, you usually use an image, which is also known as a template. (Remote PC Access catalogs usually contain physical machines, so no image is needed.)

- When using Quick Deploy, Citrix provides several Citrix prepared images you can use to create your own customized images. Or, you can import images from your own Azure subscription.
- When using Full Configuration to create VMs using a supported host type, the image usually must be created and reside on a host machine. When creating the catalog, you provide the path to that image.

Regardless of where the image resides, you can install applications on the image, if you want those apps on all machines created from that image (and don't want to virtualize those apps).

After the image is ready, you create the catalog.

- For VMs, MCS creates the machines and the catalog.
- For Remote PC Access, MCS simply creates the catalog, because the physical machines already exist.

For more information about MCS, see [Image management](#).

- **Delivery group:** A delivery group specifies:
  - One or more machines from a catalog.
  - Users who are allowed to access those machines. Alternatively, you can specify users through the Citrix Cloud Library.
  - The applications and desktops that users can access through Workspace. Alternatively, you can specify applications and users through the Citrix Cloud Library.

When using Quick Deploy, a delivery group is created automatically. (It appears only in the Full Configuration interface.)

- **Application group:** Application groups let you manage collections of applications. You can create application groups for applications shared across different delivery groups or used by a subset of users within delivery groups. Application groups are optional.

## Citrix Managed Azure

Citrix Managed Azure is an option available in several Citrix Virtual Apps and Desktops service editions. Using Citrix Managed Azure simplifies the deployment of virtual apps and desktops from Azure. Citrix manages the infrastructure for hosting Azure workloads.

With Citrix Managed Azure, you get a dedicated Citrix-managed Azure subscription and resource location. In that Azure subscription, you create a catalog of VMs. You can:

- Deploy single-session and multi-session Windows OS machines or Linux OS machines, from various supported versions.
- Choose from a curated list of compute types and storage options in select regions.
- Provision persistent or non-persistent workloads on those machines.
- Choose from several Citrix provided images that have the latest VDA installed. Then, from the Citrix interface, you build your own image from that template, and customize it. You can also import and use images from your own Azure subscriptions.

Even though Citrix manages Azure capacity, if you want to communicate with existing resources on your own Azure subscription, you can use Azure VNet peering to connect resources. You can also use Citrix SD-WAN to connect to your on-premises resources directly.

For information about security and responsibilities when using Citrix Managed Azure, see [Technical security overview for Citrix Managed Azure](#).

## Ordering Citrix Managed Azure

To get a Citrix Managed Azure subscription, you must subscribe to a supported Citrix service offering, and then order Citrix Managed Azure Consumption Funds. You can order the service and consumption funds through Citrix or from Azure Marketplace.

Citrix Managed Azure is supported on the following service offerings:

- Citrix Workspace Premium Plus
- Citrix Virtual Apps and Desktops service, Advanced and Premium editions
- Citrix Virtual Apps and Desktops Standard for Azure edition

For details, see [Sign up for the service](#).

## Citrix Managed Azure benefits summary

Using Citrix Managed Azure offers several benefits:

- Fastest path to hybrid-cloud benefits.
- Offloads IT management of infrastructure. Provides an administration experience that puts IT in control without the management and maintenance challenges.
- Enables you to rapidly scale work solutions.
- Provides a separate Azure subscription that is managed and maintained by Citrix. This isolates activity from your other Azure subscriptions.
- You retain the flexibility to create and manage workloads using your own Azure subscriptions. Your deployment can include workloads that use the Citrix Managed Azure subscription, and workloads that use your own (customer-managed) Azure subscriptions.
- Uses a true consumption-based Infrastructure as a Service (IaaS) model.
- Several technologies are available to create connections to your own on-premises networks (such as Azure VNet peering and SD-WAN). This allows your users to access your network's resources, such as file servers.

Deploying and managing Citrix Managed Azure from this service uses the [Quick Deploy](#) management interface.

For more information, contact your Citrix representative.

## Delivering applications and desktops to users

### Citrix Workspace

Subscribers (users) access their desktops and apps through Citrix Workspace.

After installing and configuring the service, you're provided with a workspace URL link. The workspace URL is posted in two places:

- From the Citrix Cloud console, select **Workspace Configuration** from the menu in the upper left corner. The **Access** tab contains the Workspace URL.
- From the Citrix Virtual Apps and Desktops service **Welcome** page, the workspace URL appears at the bottom of the page.

Test and then share the workspace URL link with your subscribers (users) to give them access to their apps and desktops. Your subscribers can access the workspace URL without any additional configuration.

You configure workspaces from Citrix Cloud.

- Specify which services are integrated with Citrix Workspace.
- Customize the URL that your subscribers use to access their workspace.
- Customize the appearance of subscribers' workspaces, such as logos, color, and preferences.
- Specify how subscribers authenticate to their workspace, such as using Active Directory or Azure Active Directory.
- Specify external connectivity for resource locations used by your subscribers.
- Automate workspace actions with Microapps and optimize workflows.

For more information, see [Citrix Workspace](#).

### **Citrix Workspace app**

From the user side, Citrix Workspace app is installed on user devices and other endpoints, such as virtual desktops. Citrix Workspace app provides users with secure, self-service access to documents, applications, and desktops from any device, including smartphones, tablets, and PCs. Citrix Workspace app provides on-demand access to Windows, web, and Software as a Service (SaaS) applications.

For devices that cannot install Citrix Workspace app software, Citrix Workspace app for HTML5 provides a connection through a HTML5-compatible web browser.

Citrix Workspace app is available for various operating systems. For details, see [Citrix Workspace app](#).

### **Service Level Agreement**

The Citrix Virtual Apps and Desktops service (the Service) is designed using industry best practices to achieve cloud scale and a high degree of service availability.

For complete details about Citrix's commitment for availability of Citrix Cloud services, see the [Service Level Agreement](#).

Performance against this goal can be monitored on an ongoing basis at <https://status.cloud.com>.

### **Limitations**

The calculation of this Service Level Goal will not include loss of availability from the following causes:

- Customer failure to follow configuration requirements for the Service documented in the product documentation on <https://docs.citrix.com>.
- Caused by any component not managed by Citrix including, but not limited to, customer controlled physical and virtual machines, customer installed and maintained operating systems, customer installed and controlled networking equipment or other hardware; customer defined and controlled security settings, group policies and other configuration policies; public cloud provider failures, Internet Service Provider failures or other external to Citrix control.
- Service disruption due to reasons beyond Citrix control, including natural disaster, war or acts of terrorism, government action.

### More information

- [Citrix Virtual Apps and Desktops Service diagrams](#)
- [Citrix Virtual Apps and Desktops Service Reference Architecture and Deployment Methods](#)
- [Technical security overview](#)
- [Network ports](#)
- [Third-party notices](#)
- [System requirements](#)
- Features
  - An introduction to [HDX technologies](#), plus details about [Devices](#), [Graphics](#), and [Multimedia](#).
  - [Remote PC Access](#): Allow users to log on remotely from anywhere to a physical PC in the office. You can configure Remote PC Access from Full Configuration or Quick Deploy.
  - [Publish content](#): Publish an application that is simply a URL or UNC path to a resource.
  - [Server VDI](#): Deliver a desktop from a server operating system for a single user.
- For the Citrix Virtual Apps and Desktops Standard for Azure service, see [its dedicated product documentation](#).
- To learn about feature availability in the Citrix Virtual Apps and Desktops products and editions, see the [Citrix Virtual Apps and Desktops feature matrix](#).
- The Citrix Cloud Learning Series offers education course to get you up and running with Citrix Cloud and its services. You can sequentially view all of the modules, from introductions through planning and building services. You can also choose individual modules or task-specific segments within a module. See [Cloud Learning Series](#).

## Get started

To learn how to set up your deployment, start with [Plan and build a deployment](#). That summary guides you through the major steps in the process, and provides links to more information and detailed procedures.

## What's new

September 9, 2021

A goal of Citrix is to deliver new features and product updates to Citrix Virtual Apps and Desktops service customers when they are available. New releases provide more value, so there's no reason to delay updates. Rolling updates to the service release approximately every three weeks.

This process is transparent to you. Initial updates are applied to Citrix internal sites only, and are then applied to customer environments gradually. Delivering updates incrementally in waves helps ensure product quality and maximize availability.

For details about the Service Level Agreement for cloud scale and service availability, see [Service Level Agreement](#). To monitor service interruptions and scheduled maintenance, see the [Service Health Dashboard](#).

## Virtual Delivery Agents (VDAs)

VDAs for Windows machines generally release at the same time as the on-premises Citrix Virtual Apps and Desktops product.

- For information about new VDA and HDX features, see the [What's new](#) and [Known issues](#) articles for the current on-premises Citrix Virtual Apps and Desktops release.
- For information about VDA platforms and features that are no longer supported, see [Deprecation](#). That article also includes platforms and features that are scheduled to be unsupported in a future release (such as which operating systems support VDA installation).

### Important:

If the Personal vDisk (PvD) component was ever installed on a VDA, that VDA cannot be upgraded to version 1912 LTSR or later. To use the new VDA, you must uninstall the current VDA and then install the new VDA. (This instruction applies even if you installed PvD but never used it.) For details, see [If the VDA has Personal vDisk installed](#).

## September 2021

### New and enhanced features

**Support for non-domain-joined catalogs.** We added an identity type, **Non-domain-joined**, to the **Machine Catalog Setup > Machine Identities** page of the Full Configuration management interface. With that identity type, you can use MCS to create machines that are not joined to any domain. For more information, see [Create machine catalogs](#).

**Support for using a machine profile.** We added an option, **Use a machine profile**, to the **Machine Catalog Setup > Master Image** page of the Full Configuration management interface. The option lets you specify which machine profile you want the image to inherit the configuration from when creating VMs in Azure environments. The image can inherit the following configurations from the selected machine profile:

- Accelerated networking
- Boot diagnostics
- Host disk caching (relating to OS and MCSIO disks)
- Machine size (unless otherwise specified)
- Tags placed on the VM

For more information, see [Create a machine catalog using an Azure Resource Manager image](#).

**Support for Windows Server 2022.** Requires minimum VDA 2106.

## August 2021

### New and enhanced features

**Extend the number of sortable items from 500 to 5,000.** On the **Search** node of the Full Configuration management interface, you can now sort up to 5,000 items by any column heading. When the number of items exceeds 5,000, use filters to reduce the number of items to 5,000 or fewer to enable sorting. For more information, see [Use Search in the Full Configuration management interface](#).

**Support for additional Azure storage types.** You can now select different storage types for virtual machines in Azure environments using MCS. For details, see [Storage types](#).

**Support for selecting the storage type for write-back cache disks.** In the Full Configuration management interface, when creating an MCS catalog, you can now select the storage type for the write-back cache disk. Available storage types include: Premium SSD, Standard SSD, and Standard HDD. For more information, see [Create machine catalogs](#).

**Shut down suspended machines.** In the **Manage > Full Configuration** interface, we added an option, **When no reconnection in (minutes)**, to the **Load-based Settings** page of the Manage Autoscale user interface for single-session OS delivery groups. The option becomes available after you select

**Suspend**, letting you specify when to shut down the suspended machines. Suspended machines remain available to disconnected users when they reconnect but are not available for new users. Shutting the machines down makes them available again to handle all workloads. For more information, see [Autoscale](#).

**Extended support for using CSV files to bulk add machines to a catalog.** In the **Manage > Full Configuration** interface, you can now use a CSV file to bulk add machines already in your data center to a catalog where those machines are power managed. For more information, see [Create machine catalogs](#) and [Manage machine catalogs](#).

## July 2021

### New and enhanced features

**Configuration logging.** The **Logging** user interface has changed in **Manage > Full Configuration**. The following three tabs comprise the interface:

- **Events** (formerly, configuration logging). This tab lets you track configuration changes and administrative activities.
- **Tasks.** This tab lets you view tasks related to machine catalog operations.
- **APIs.** This tab lets you view REST API requests made during a certain time period.

For more information, see [Configuration logging](#).

**Autoscale now provides you with dynamic session timeout options.** You can configure disconnected and idle session timeouts for your peak and off-peak usage times to achieve faster machine draining and cost savings. For more information, see [Dynamic session timeouts](#).

**Support for Google Cloud Platform (GCP) Customer Managed Encryption Keys (CMEK).** You can now use Google's CMEK with MCS catalogs. CMEK provides greater control over keys used to encrypt data within a Google Cloud project. For more information, see [Customer-managed encryption keys \(CMEK\)](#). To configure this feature, see [Using Customer Managed Encryption Keys \(CMEK\)](#). The feature is available on the **Machine Catalog Setup > Disk Settings** page of the **Manage > Full Configuration** interface.

#### Note:

This feature is available as a preview.

**Updates to the Manage tab.** We have updated options in the menu of the **Manage** tab:

- **Full Configuration:** Previously, this option took you to the legacy console. It now takes you to the new, web-based console (Web Studio). The web-based console has full parity with the legacy console and includes several enhancements. We recommend that you start using it now.
- **Legacy Configuration:** This option takes you to the legacy console, which is scheduled for removal in September 2021. After that, **Full Configuration** will be the only interface that offers



access to the full range of configuration and management actions.

**Web Studio now supports choosing a power management connection for a Remote PC Access catalog.** Previously, you were able to use Studio to create a Wake on LAN host connection to your resource location (selecting **Remote PC Wake on LAN** as the connection type). However, PowerShell was your only choice to associate that connection with a Remote PC Access catalog. You can now use Studio to achieve that. For more information, see [Configure Wake on LAN in Studio](#).

## June 2021

### New and enhanced features

**Access Azure Shared Image Gallery images.** When creating a machine catalog, you can now access images from the Azure Shared Image Gallery on the Master Image screen. For details, see [Access images from Azure Shared Image Gallery](#).

**Support shielded virtual machines on Google Cloud Platform (GCP).** You can provision shielded virtual machines on GCP. A shielded virtual machine is hardened by a set of security controls that provide verifiable integrity of your Compute Engine instances, using advanced platform security capabilities like secure boot, a virtual trusted platform module, UEFI firmware and integrity monitoring. For more information, see [Shield VMs](#).

**Enforce either HTTPS or HTTP.** Use registry settings to [enforce HTTPS or HTTP traffic through the XML service](#).

**Always use standard SSD for an identity disk to reduce cost in Azure environments.** Machine catalogs use the standard SSD storage type for identity disks. Azure standard SSDs are a cost-effective storage option optimized for workloads that need consistent performance at lower IOPS levels. For more information about storage types, see [Azure Resource Manager master image](#).

#### Note:

For more information about Azure managed disk pricing, see [Managed Disks pricing](#).

**New feature available in Web Studio.** The following features are now available in the web-based console:

- **Studio now supports authenticating to Azure to create a service principal.** You can now establish a host connection to Azure by authenticating to Azure to create a service principal. This support eliminates the need to manually create a service principal in your Azure subscription before creating a connection in Studio. For more information, see [Microsoft Azure Resource Manager virtualization environments](#).
- **Studio now supports cloning of existing machine catalogs.** This feature enables you to clone an existing machine catalog to use as a template for a new one, eliminating the need to create a similar catalog from scratch. When cloning a catalog, you cannot change settings associated

with operating system and machine management. The cloned catalog inherits those settings from the original. For more information, see [Clone a catalog](#).

- **A new node called Settings now available in the Studio navigation pane.** The **Settings** node lets you configure settings that apply to the entire site (your deployment of a Citrix Virtual Apps and Desktops service product). The following settings are available:
  - **Load balance multi-session catalogs.** Select the load balancing option that meets your needs. This setting applies to all your catalogs. Previously, you accessed this feature by clicking the gear icon in the upper right corner of the console. For more information, see [Load balance machines](#).
- **Enhanced search experience in Studio.** This release enhances your Studio search experience. When you use filters to perform an advanced search, the Add filters window appears in the foreground, leaving the background view unchanged. For more information, see [Use Search in the Full Configuration management interface](#).
- **Ability to suspend and resume Google Cloud VMs in MCS.** You can now suspend and resume Google Cloud VMs in MCS as you would any VM. For details, see [Manage delivery groups](#). To enable this ability, set the `compute.instances.suspend` and `compute.instances.resume` permissions in the Google Cloud service account. The Compute Admin role comes with these permissions.

In Citrix Virtual Apps and Desktops, you can also use the `New-BrokerHostingPowerAction` PowerShell command to suspend and resume the VMs. For details, see [New-Brokerhostingpoweraction](#).

Google Cloud enforces some limitations on the type and configuration of instances that can be suspended. For additional information, refer to [Suspending and resuming an instance](#) on the Google Cloud site.

## May 2021

### New and enhanced features

**Session reconnection after disconnect from machine in maintenance mode.** Previously, when pooled (random) single-session desktop (VDI) users were disconnected from a machine in maintenance mode, session reconnection was not allowed to any machine in the pool. Multi-session and static single-session machines always allowed session reconnection in that circumstance.

Now, using PowerShell, you can control at the delivery group level whether session reconnection is allowed after a disconnect occurs on a machine in maintenance mode. This applies to all VDAs in the group (single-session and multi-session).

For details, see [Control session reconnection when disconnected from machine in maintenance mode](#).

**Application probing and Desktop probing support in all Citrix Virtual Apps and Desktops Service editions.** In addition to the existing **Premium** edition support, Application probing and Desktop probing are now available in **Citrix Virtual Apps Advanced Service** and **Citrix Virtual Apps and Desktops Advanced Service** editions.

**New feature available in Web Studio.** The following feature is now available in the web-based console:

- **Studio now supports selecting Azure Availability Zones.** Previously, PowerShell was your only choice to provision machines into a specific Availability Zone in Azure environments. When using Studio to create a machine catalog, you can now select one or more Availability Zones into which you want to provision machines. If no zones are specified, Machine Creation Services (MCS) lets Azure place the machines within the region. If more than one zone is specified, MCS randomly distributes the machines across them. For more information, see [Provision machines into specified availability zones](#).

**Azure ephemeral disk.** Citrix Virtual Apps and Desktops service supports Azure ephemeral disk. An ephemeral disk allows you to repurpose the cache disk to store the OS disk for an Azure-enabled virtual machine. This functionality is useful for Azure environments that require a higher performant SSD disk over a standard HDD disk.

**Note:**

Persistent catalogs do not support ephemeral OS disks. Also, when using this feature, consider that the extra performant disk incurs an extra cost. It's beneficial to reuse the cache disk to store the OS disk instead of paying for an extra managed disk.

Ephemeral OS disks require that your provisioning scheme use managed disks and a Shared Image Gallery. For more information, see [Azure ephemeral disks](#).

**Improved performance for MCS managed VDAs on Azure.** The Citrix Virtual Apps and Desktops service improves performance for VDAs managed with Machine Creation Services (MCS) on Azure. This enhancement changes the default values for *Absolute Simultaneous actions* for the hosting connection to 500, and *Maximum new actions per minute* for the hosting connection to 2,000. No manual configuration tasks are required to take advantage of this enhancement. For details, see [Azure throttling](#).

**New features available in Cloud Health Check.** Cloud Health Check has been updated to a new version with features including:

- **Automatically discovering VDA machines.** Cloud Health Check can now automatically discover and retrieve VDAs from your Citrix Virtual Apps and Desktops service deployments. For more information, see [Retrieve VDA machines](#).
- **Scheduling health checks.** Cloud Health Check now lets you set up schedules for performing periodic health checks. For more information, see [Cloud Health Check scheduler](#).

- **Cloud Health Check version information.** You can now check which version of Cloud Health Check you are using. To view version information, click the gear icon in the upper right corner of the Cloud Health Check main window.
- **Automatic fix (preview).** Cloud Health Check now supports automatically detecting and fixing certain issues identified on machines where it is running. For more information, see [Automatic fix](#).

**Note:**

Automatic fix is available as a preview.

## April 2021

### New and enhanced features

**Retrieve dynamic instances using AWS API.** The Citrix Virtual Apps and Desktops service now queries AWS to retrieve instance types dynamically. This functionality removes the need to create a custom `InstanceTypes.xml` file for those customers wishing to use machine sizes beyond those defined in Citrix Virtual Apps and Desktops service. This information was previously supplied by the `InstanceTypes.xml` file. To facilitate this dynamic access to the available AWS instance types, users must update the permissions on their service principals to include `ec2:DescribeInstanceTypes` permissions. To support backward compatibility for customers who choose not to update their service principal permissions, the AWS instance types listed in the `InstanceTypes.xml` are used. This process generates a warning message to the MCS CDF log.

**Note:**

Citrix Studio does not display the warning message contained in the CDF log.

For more information about permissions, see [Defining IAM permissions](#) and [About AWS permissions](#).

**New feature available in Web Studio.** The following feature is now available in the web-based console:

- **Studio now displays date and time of your time zone.** Previously, Studio displayed only date and time based on the system clock and time zone. Studio now supports displaying date and time local to your time zone when you hover the mouse pointer over an event item. The time is expressed in UTC.

**MCS I/O support for Azure VMs without temporary storage.** MCS I/O now supports machine catalog creation for VMs that do not have temporary disks or attached storage. With this support:

- The snapshot (managed disk) is retrieved from the source VM *without* temporary storage. The VMs in the machine catalog have no temporary storage.
- The snapshot (managed disk) is retrieved from the source VM *with* temporary storage. The VMs in the machine catalog have temporary storage.

For more information, see [Machine Creation Services \(MCS\) storage optimization](#)

**New feature available in Web Studio.** The following feature is now available in the web-based console:

- **Force log off.** Autoscale now lets you log off sessions existing on machines by force when the established grace period is reached, making the machine eligible for shutdown. Doing that enables Autoscale to power off machines much faster, thus reducing costs. You can send notifications to users before they are logged off. For more information, see [Autoscale](#).

**New update for Automated Configuration.** Automated Configuration has been updated to a new version with features including:

- **Merging multiple sites** – you can merge multiple sites into a single site while avoiding name collisions using prefixes and suffixes. For more information, see [Merging multiple sites into a single site](#).
- **Site activation** – you can select whether your on-premises or cloud deployment controls resources such as reboot schedules and power schemes. For more information, see [Activating sites](#).

Other updates to Automated Configuration include:

- The ability to migrate administrator roles and scopes.
- A `Quiet` parameter for select cmdlets to suppress console logging.
- A `SecurityFileFolder` parameter to allow placing of the `CvadAcSecurity.yml` file in a secure network file share that requires authentication.
- The ability to filter by machine name in machine catalogs and delivery groups.
- Improvements to component selection parameters to use the switch parameter method, eliminating the need to add a `$true` after the component name.
- A new cmdlet (`New-CvadAcZipInfoForSupport`) to zip all your log files to send to Citrix for support.

Download Automated Configuration at [Citrix Downloads](#). For more information on Automated Configuration, see [Migrating to cloud](#).

**Preserve GCP instances across power cycles.** Non-persistent Google Cloud Platform (GCP) instances are no longer deleted when powering off. Instead, the instances are preserved across power cycles. When a non-persistent instance is powered off, the OS disk is detached and deleted. When the instance is powered on, the OS disk is recreated from the base disk, and attached to the existing instance.

**Support for Azure Gen2 images.** You can now provision a Gen2 VM catalog by using either a Gen2 snapshot or a Gen 2 managed disk to improve boot time performance. For more information, see [Create machine catalogs](#). The following operating systems are supported for Azure Gen2 images:

- Windows Server 2019, 2016, 2012 and 2012 R2
- Windows 10

**Note:**

Creating a Gen2 machine catalog using a Gen1 snapshot, or managed disk, is not supported. Similarly, creating a Gen1 machine catalog using a Gen2 snapshot, or managed disk, is also not supported. For more information, see [Support for generation 2 VMs on Azure](#).

**Disabling table storage accounts.** Machine Creation Services (MCS) no longer creates table storage accounts for catalogs that use managed disks when provisioning VDAs on Azure. For more information, see [Azure table storage](#).

**Eliminating locks in storage accounts.** When creating a catalog in Azure using a managed disk, a storage account is no longer created. Storage accounts created for existing catalogs remain unchanged. This change is applicable for managed disks only. For unmanaged disks, there is no change in the existing behavior. Machine Creation Services (MCS) continues creating storage accounts and locks.

**New features available in Web Studio.** The following features are now available in the web-based console:

- **Use a customer-managed encryption key to encrypt data on machines.** Studio now adds a setting called **Customer-managed encryption key** to the **Machine Catalog Setup > Disk Settings** page. The setting lets you choose whether to encrypt data on the machines to be provisioned in the catalog. For more information, see [Customer-managed encryption key](#).
- **Studio now supports restricting Autoscale to tagged machines.** Previously, you had to use PowerShell to restrict Autoscale to certain machines in a delivery group. You can now also use Studio. For more information, see [Restrict Autoscale to certain machines in a delivery group](#).

## March 2021

### New and enhanced features

**Azure dedicated hosts.** Azure dedicated hosts allow you to provision virtual machines on hardware dedicated to a single customer. While using a dedicated host, Azure ensures that your virtual machines would be the only machines running on that host. This provides more control and visibility to customers thereby ensuring they meet their regulatory or internal security requirements. A pre-configured Azure host group, in the region of the hosting unit, is required when using the `HostGroupId` parameter. Also, Azure auto-placement is required. For more information, see [Azure dedicated hosts](#).

**Tip:**

When using Azure dedicated hosts, selecting the **Azure Availability Zone** has no effect. The vir-

tual machine is placed by the Azure auto-placement process.

**Support for Azure server side encryption.** Citrix Virtual Apps and Desktops service supports customer-managed encryption keys for Azure managed disks. With this support you can manage your organizational and compliance requirements by encrypting the managed disks of your machine catalog using your own encryption key. For more information, see [Azure server side encryption](#).

**Provision machines into specified availability zones on Azure.** You can now provision machines into a specific availability zone in Azure environments. With this functionality:

- You can specify one or multiple Availability Zones on Azure. Machines are nominally equally distributed across all provided zones if more than one zone is provided.
- The virtual machine and the corresponding disk are placed in the specified zone (or zones).
- You can browse Availability Zones for a given service offering or region. Valid Availability Zones are displayed via PoSH. View service offering inventory items using `Get-Item`.

For more information, see [Provision machines into specified availability zones on Azure](#).

**New features available in Web Studio.** The following features are now available in the web-based console:

- **Studio now supports associating apps with custom icons.** Previously, you had to use PowerShell to add custom icons for use with published applications. You can now also use Studio to do that. For more information, see [Manage application groups](#).
- **Studio now supports applying tags to machine catalogs.** Previously, you might use Studio to create or delete tags for use with a catalog. However, you had to use PowerShell to apply tags to the catalog. You can now also use Studio to apply or remove a tag to or from a catalog as you do with delivery groups. For more information, see [Apply tags to machine catalogs](#).
- **Studio now supports switching between “horizontal load balancing” and “vertical load balancing” modes.** Previously, PowerShell was your only choice to switch between horizontal and vertical load balancing modes. Studio now gives you more flexibility to control how to load balance multi-session OS machines. For more information, see [Load balance machines](#).
- **Studio now supports including machines in maintenance mode in restart schedules.** Previously, PowerShell was your only choice to configure scheduled restarts for machines in maintenance mode. You can now also use Studio to control whether to include those machines in a restart schedule. For more information, see [Create a restart schedule](#).
- **Studio now supports configuring Wake on LAN for Remote PC Access.** Previously, you had to use PowerShell to configure Wake on LAN for Remote PC Access. You can now also use Studio to configure the feature. For more information, see [Configure Wake on LAN](#).
- **Studio now supports applying AWS instance properties and tagging operational resources.** When creating a catalog to provision machines in AWS by using MCS, you can specify whether to



apply the IAM role and tag properties to those machines. You can also specify whether to apply machine tags to operational resources. You have the following two options:

- **Apply machine template properties to virtual machines**
- **Apply machine tags to operational resources**

For more information, see [Applying AWS instance properties and tagging operational resources](#).

**Azure Shared Image Gallery.** Citrix Virtual Apps and Desktops service supports Azure Shared Image Gallery as a published image repository for MCS provisioned machines in Azure. Administrators have the option of storing an image in the gallery to accelerate the creation and hydration of OS disks. This process improves the boot and application launch times for non-persistent VMs. For details about this feature, see [Azure shared image gallery](#).

**Note:**

Shared Image Gallery functionality only works with managed disks. It is not available for legacy machine catalogs.

**Storage buckets created in same Google Cloud Platform region as the machine catalog.** In previous releases, MCS created temporary storage buckets during provisioning as part of the disk upload process. These buckets spanned multiple regions, which [Google](#) defines as a large geographic area containing two or more geographic places. These temporary buckets resided in the United States geographic location, no matter where the catalog was provisioned. MCS now creates storage buckets in the same region where you provision your catalogs. Storage buckets are no longer temporary; they remain in your Google Cloud Platform project after you complete the provisioning process. Future provisioning operations use the existing storage bucket, if one exists in that region. A new storage bucket is created if one does not exist in the specified region.

## February 2021

### New and enhanced features

**Support for Azure Gen2 images.** You can now provision managed disks using Gen2 VMs in Azure environments to improve boot time performance. The following operating systems are supported:

- Windows Server 2019, 2016, 2012 and 2012 R2
- Windows 10

**Note:**

With this support, only a subset of VMs is supported. For example, some VMs can be both Gen1 and Gen2 types, while other VMs can only be Gen1. For more information, see [Support for generation 2 VMs on Azure](#).

**Machine restart schedules.** Citrix Studio now adds an option called **Restart all machines after draining sessions** to the **Restart duration** menu. The option lets you choose whether to restart all



machines after draining all sessions. When the restart time is reached, machines are put into the drain state and restarted when all sessions are logged off. For more information, see [Create a restart schedule](#).

**New features available in Web Studio.** The following features are now available in the web-based console:

- **Studio now supports using CSV files to bulk add machines to a catalog.** This feature enables you to use a CSV file to:
  - Bulk add machines to a multi-session or single-session OS catalog where machines are not power managed through Studio.
  - Bulk add machines to a Remote PC Access catalog. Previously, you had to choose OUs to bulk add machines to a Remote PC Access catalog. Doing that, however, is not easy in scenarios with OU structure restrictions. The feature gives you more flexibility to bulk add machines. You can add only machines (for use with user auto-assignments) or add machines along with user assignments.

For more information, see [Create machine catalogs](#) and [Manage machine catalogs](#).

- **Extended support for Citrix Managed Azure.** [Citrix Managed Azure](#) is now available in the following Citrix Virtual Apps and Desktops service editions: Standard for Azure, Advanced, Premium, and Workspace Premium Plus.
- **Support for placing master images in Azure Shared Image Gallery.** Studio now provides you an option to place master images in Azure Shared Image Gallery (SIG). SIG is a repository for managing and sharing images. It lets you make your images available throughout your organization. We recommend that you store a master image in SIG when creating large non-persistent machine catalogs because doing that enables faster reset of VDA OS disks. For more information, see [Microsoft Azure Resource Manager virtualization environments](#).
- **Retain system disk for MCS machine catalogs in Azure.** Studio now lets you control whether to retain system disks for VDAs during power cycles. Ordinarily, the system disk is deleted on shutdown and recreated on startup. This ensures that the disk is always in a clean state but results in longer VM restart times. If system writes are redirected to the cache and written back to the cache disk, the system disk remains unchanged. To avoid unnecessary disk recreation, use the **Retain system disk during power cycles** option, available on the **Machine Catalog Setup > Disk Settings** page. Enabling the option reduces VM restart times but increases your storage costs. The option can be useful in scenarios where an environment contains workloads with sensitive restart times. For more information, see [MCS storage optimization](#).
- **Studio now supports creating MCS machine catalogs with persistent write-back cache disk.** Previously, PowerShell was your only choice to create a catalog with persistent write-back cache disk. You can now use Studio to control whether the write-back cache disk persists for the provisioned VMs in Azure when you are creating a catalog. If disabled, the write-back cache

disk is deleted during each power cycle to save storage costs, causing any data redirected to the disk to be lost. To retain the data, enable the **Use persistent write-back cache disk** option, available on the **Machine Catalog Setup > Disk Settings** page. For more information, see [MCS storage optimization](#).

## January 2021

**New features available in Web Studio.** The following features are now available in the web-based console:

- **Studio now supports associating apps with custom icons.** Previously, you had to use PowerShell to add custom icons for use with published applications. You can now also use Studio to do that. For more information, see [Manage application groups](#).
- **Studio now supports applying tags to machine catalogs.** Previously, you might use Studio to create or delete tags for use with a catalog. However, you had to use PowerShell to apply tags to the catalog. You can now also use Studio to apply or remove a tag to or from a catalog as you do with delivery groups. For more information, see [Apply tags to machine catalogs](#).
- **Studio now supports switching between “horizontal load balancing” and “vertical load balancing” modes.** Previously, PowerShell was your only choice to switch between horizontal and vertical load balancing modes. Studio now gives you more flexibility to control how to load balance multi-session OS machines. For more information, see [Load balance machines](#).
- **Studio now supports including machines in maintenance mode in restart schedules.** Previously, PowerShell was your only choice to configure scheduled restarts for machines in maintenance mode. You can now also use Studio to control whether to include those machines in a restart schedule. For more information, see [Create a restart schedule](#).
- **Studio now supports configuring Wake on LAN for Remote PC Access.** Previously, you had to use PowerShell to configure Wake on LAN for Remote PC Access. You can now also use Studio to configure the feature. For more information, see [Configure Wake on LAN](#).
- **Studio now supports applying AWS instance properties and tagging operational resources.** When creating a catalog to provision machines in AWS by using MCS, you can specify whether to apply the IAM role and tag properties to those machines. You can also specify whether to apply machine tags to operational resources. You have the following two options:
  - **Apply machine template properties to virtual machines**
  - **Apply machine tags to operational resources**

For more information, see [Applying AWS instance properties and tagging operational resources](#).

- **AWS dedicated host.** Citrix Studio now adds an option called **Use dedicated host** to the **Machine Catalog Setup > Security** page. This setting is suitable for deployments with licensing

restrictions or security requirements that need your use of a dedicated host. With a dedicated host, you own an entire physical host and are billed on an hourly basis. Owning that host lets you spin up as many EC2 instances as that host permits, without more charges. For more information, see [AWS tenancy](#).

- **Studio now supports running a restart schedule immediately.** Studio now lets you run a restart schedule immediately to restart all applicable machines in the schedule. For more information, see [Immediately run a restart schedule](#).
- **Autoscale.** Autoscale provides the following new features and enhancements:
  - **Studio now supports displaying machines in drain state.** Previously, PowerShell was your only choice to identify machines in drain state. You can now use Studio to identify machines that are in drain state. For more information, see [Display machines in drain state](#).
  - **Studio now supports defining peak times at a granular level of 30 minutes for VDI delivery groups.** Previously, you had to use PowerShell to define the peak times for the days included in a schedule at a granular level of 30 minutes for VDI delivery groups. You can now also use Studio to do that. This support enables you to set the minimum number of machines running in a VDI Delivery Group separately for each half hour of the day.

**Azure Shared Image Gallery.** Citrix Virtual Apps and Desktops service supports Azure Shared Image Gallery as a published image repository for MCS provisioned machines in Azure. Administrators have the option of storing an image in the gallery to accelerate the creation and hydration of OS disks from the master image. This process improves the boot and application launch times for non-persistent VMs.

The gallery contains the following three elements:

- Gallery. Images are stored here. MCS creates one gallery for each machine catalog.
- Gallery Image Definition. This definition includes information (operating system type and state, Azure region) about the master image. MCS creates one image definition for each master image created for the catalog.
- Gallery Image Version. Each image in a Shared Image Gallery can have multiple versions, and each version can have multiple replicas in different regions. Each replica is a full copy of the master image. Citrix Virtual Apps and Desktops service always creates one Standard\_LRS image version (version 1.0.0) for each image with the appropriate number of replicas in the catalog's region. This configuration is based on the number of machines in the catalog, the configured replica ratio, and the configured replica maximum.

**Note:**

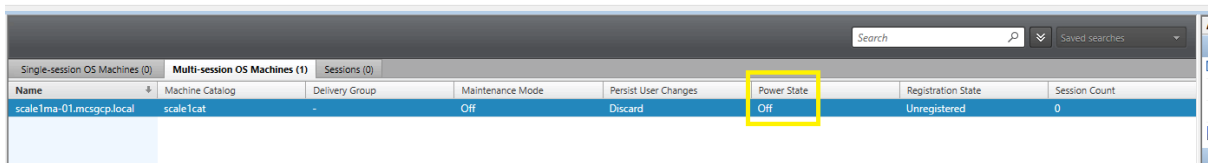
Shared Image Gallery functionality only works with managed disks. It is not available for legacy machine catalogs.

For details about this feature, see [Configure shared image gallery](#).

**Storage buckets created in same Google Cloud Platform region as the machine catalog.** In previous releases, MCS created temporary storage buckets during provisioning as part of the disk upload process. These buckets spanned multiple regions, which Google defines as a large geographic area containing two or more geographic places. These temporary buckets resided in the United States geographic location, no matter where the catalog was provisioned. MCS now creates storage buckets in the same region where you provision your catalogs. Storage buckets are no longer temporary; they remain in your Google Cloud Platform project after you complete the provisioning process. Future provisioning operations use the existing storage bucket. If one exists in that region, or a new storage bucket is created if one does not exist in the specified region.

**PowerShell option that sets default to re-use pooled VDAs during an outage.** A new PowerShell command option (`-DefaultReuseMachinesWithoutShutdownInOutage`) extends the ability to reuse pooled desktop VDAs that haven't been shut down during an outage, by default. See [Application and desktop support](#).

**Google Cloud Platform on-demand provisioning.** Citrix Virtual Apps and Desktops service updates how the Google Cloud Platform (GCP) provisions machine catalogs. When creating a machine catalog, the corresponding machine instance is not created in GCP and the power state is set to **OFF**. Machines are not provisioned at catalog creation time but rather the first time the machines are powered on. For example, after you create a catalog, the VM power state is set to **Off**:



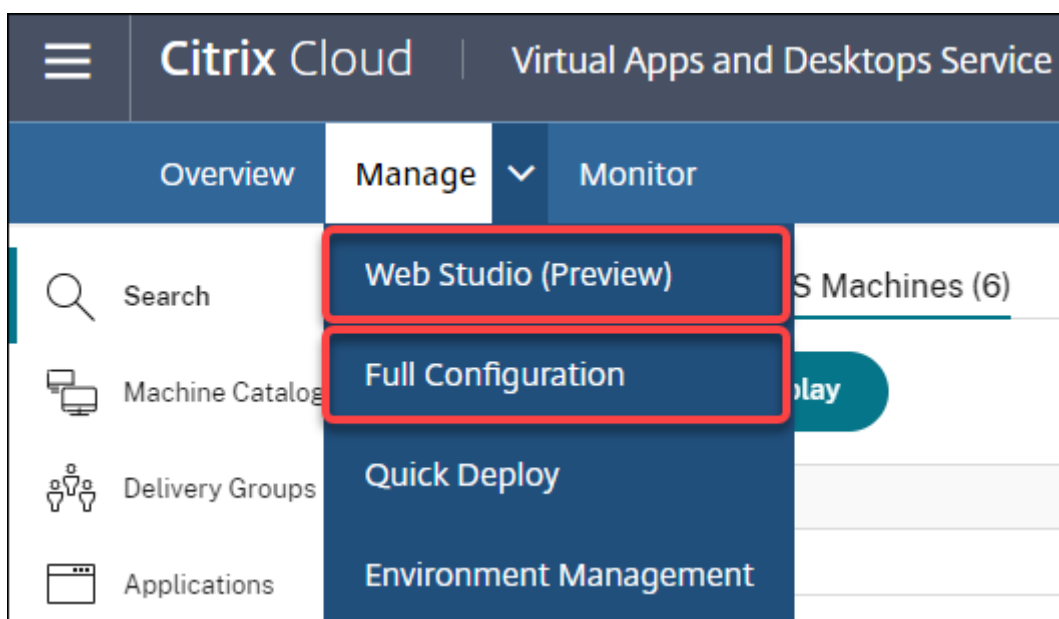
Name	Machine Catalog	Delivery Group	Maintenance Mode	Persist User Changes	Power State	Registration State	Session Count
scale1ma-01.mcs-gcp.local	scale1cat	-	Off	Discard	Off	Unregistered	0

## December 2020

### New and enhanced features

**Web Studio is available as a preview.** A new, web-based console is now available. We are in the process of migrating the full set of Studio functionalities from the legacy console to the new, web-based console. The web-based console generally responds faster than the legacy console. By default, you automatically log on to the web-based console. You can easily switch between the web-based console and the legacy console from within the **Manage** tab to perform your configuration or site management tasks. Click the down arrow next to **Manage** and select an option:

- **Web Studio (Preview).** Takes you to the new, web-based console.
- **Full Configuration.** Takes you to the legacy console.



The following features are available only in the web-based console:

- **Standard SSD disk type support for Azure.** Studio now adds support for standard SSD disk type. Azure standard SSDs are a cost-effective storage option optimized for workloads that need consistent performance at lower IOPS levels. For more information, see [Create a machine catalog using an Azure Resource Manager master image](#).
- **Studio now supports configuring the power-off delay for static VDI delivery groups.** Previously, you might configure the power-off delay for static VDI delivery groups only through the PowerShell SDK. Studio now lets you configure the power-off delay in the Autoscale user interface for static VDI delivery groups. For more information, see [Autoscale](#).

## October 2020

### New and enhanced features

**Dismiss multiple hypervisor alerts.** Citrix Monitor now supports automatic dismissal of hypervisor alerts older than a day. For more information, see [Hypervisor Alerts Monitoring](#).

**Remove external IP address.** An external IP address on a temporary virtual machine that is used to prepare a provisioned image in the Google Cloud Platform (GCP) is no longer required. This external IP address enables the temporary virtual machine to access the Google public API to complete the provisioning process.

Enable Private Google Access to permit the VM to access the Google public API directly from the subnet. For more information, see [Enable Google Private access](#).

**New model addresses how machine identities are managed.** Machine identities used in machine catalogs have been managed and maintained using Active Directory. All machines created by MCS will

now be joining Active Directory. The new Citrix Virtual Apps and Desktops service model addresses how machine identities are managed. This model allows the creation of machine catalogs using *workgroup*, or, non-domain joined machines.

**Tip:**

This functionality supports a new identity service, *FMA trust*, added to Citrix Cloud for non-domain joined machines.

MCS communicates with the new FMA trust service for identity management. Identity information is stored in the identity disk as a pair of GUID and private key pairs, instead of the domain SID and machine account password paradigm used by Active Directory. VDAs using non-domain joined machines use this GUID and private key combination for broker registration. For more information, see [Configure support for non-domain joined catalogs](#).

**Use direct upload for Azure managed disks.** This release allows you to use direct upload when creating managed disks in an Azure environment. This functionality reduces costs associated with extra storage accounts. You no longer have to stage the VHD into a storage account before converting it to a managed disk. Also, direct upload eliminates the need to attach an empty managed disk to a virtual machine. Directly uploading to an Azure managed disk simplifies the workflow by enabling you to copy an on-premises VHD directly for use as a managed disk. Supported managed disks include Standard HDD, Standard SSD, and Premium SSD.

For more information on this feature, see the Microsoft Azure [blog](#).

For more information about Azure managed disks, see the [documentation page](#).

**Single Resource Group in Azure.** You can now create and use a single Azure resource group for updating and creating catalogs in Citrix Virtual Apps and Desktops. This enhancement applies to both the full scope and narrow scope service principals.

The previous limit of 240 VMs per 800 managed disks per Azure Resource Group has been removed. There is no longer a limit on the number of virtual machines, managed disks, snapshots, and images per Azure Resource Group.

For more information, see [Microsoft Azure Resource Manager virtualization environments](#).

## September 2020

### New and enhanced features

**Quick Deploy.** The new [Quick Deploy](#) feature replaces the earlier Azure Quick Deploy. The new feature offers a quick way to get started with the Citrix Virtual Apps and Desktops service using Microsoft Azure. You can use Quick Deploy to deliver desktops and apps, and configure Remote PC Access.

**Session Administrator (built-in role).** Citrix Studio now adds a new built-in role called **Session Administrator**. The role lets an administrator view delivery groups and manage their associated ses-

sions and machines on the **Filters** page of the **Monitor** tab. With this feature, you can configure access permissions of existing administrators or administrators you invite in a way that aligns with their role in your organization. For more information about the built-in role, see [Built-in roles and scopes](#). For information about how to assign the built-in role to an administrator, see [Delegated administration and monitoring](#).

For a more granular level of control over access to the **Filters** page related to sessions and machines, create a custom role and select one of the following for the Director object: **View Filters page - Machines only**, **View Filters page - Sessions only**. For information about creating a custom role, see [Create and manage roles](#).

**Support for a new machine type.** This release adds support for the NV v4 and the DA v4 series of AMD machines, when configuring Premium Disks for a machine catalog. For more information, see [Create delivery groups](#).

## August 2020

### New and enhanced features

**Limited access to the Remote PowerShell SDK during an outage.** Previously, you might not use PowerShell commands during an outage. Now, Local Host Cache allows limited access to the Remote PowerShell SDK during an outage. See [What is unavailable during an outage](#).

**Support for two new Citrix Virtual Apps and Desktops service editions.** Citrix Monitor now supports two new Citrix Virtual Apps and Desktops service editions, namely, **Citrix Virtual Apps Advanced service** and **Citrix Virtual Apps and Desktops Advanced service**. For more information, see the Citrix Monitor [Feature compatibility matrix](#).

**Support for shared Virtual Private Cloud (VPC) in Google Cloud Platform.** The Citrix Virtual Apps and Desktops service supports Shared VPC on Google Cloud Platform as a host resource. You can use Machine Creation Services (MCS) to provision machines in a Shared VPC and manage them using Citrix Studio. For information about Shared VPC, see [Shared Virtual Private Cloud](#).

**Zone selection support for Google Cloud Platform.** The Citrix Virtual Apps and Desktops service supports zone selection on Google Cloud Platform. This feature allows administrators to specify one or multiple zones within a region for catalog creation.

For sole-tenant type VMs, zone selection provides administrators with the ability to place sole tenant nodes across zones of their choice. For non-sole tenant VMs, zone selection provides the ability to place VMs deterministically across zones of their choice thereby providing flexibility in designing the deployment. For configuration information, see [Enable zone selection](#).

Also:



- Sole tenancy provides exclusive access to a sole tenant node, which is a physical compute engine server dedicated to hosting only your project's VMs. These nodes allow you to group your VMs together on the same hardware or separate your VMs from other project's VMs.
- Sole tenant nodes help you meet dedicated hardware requirements for Bring Your Own License (BYOL) scenarios. They also enable you to comply with network access control policy, security, and privacy requirements such as HIPAA.

**Note:**

Sole tenancy is the only route to using Windows 10 VDI deployments on Google Cloud. Server VDI also supports this method. A detailed description for sole tenancy can be found on the [Google documentation site](#).

**Improved boot performance for Azure system disks.** This release supports improved boot performance for Citrix Cloud implementations using Azure when MCSIO is enabled. With this support, you can retain the system disk. This provides the following advantages:

- VMs and applications now boot and launch with performance similar to how the golden image is served.
- Reduction in API quota consumption, deleting and creating the system disk, and state transition delay caused when you delete a VM.

For example, use the PowerShell `PersistOsDisk` custom property in the `New-ProvScheme` command to configure this feature.

```
1 -CustomProperties '<CustomProperties xmlns="http://schemas.citrix.com
2   /2014/xd/machinecreation" xmlns:xsi="http://www.w3.org/2001/
3   XMLSchema-instance">
4   <Property xsi:type="StringProperty" Name="UseManagedDisks" Value="true"
5     />
6   <Property xsi:type="StringProperty" Name="StorageAccountType" Value="
7     Premium_LRS" />
8   <Property xsi:type="StringProperty" Name="ResourceGroups" Value="
9     benvaldev5RG3" />
10  <Property xsi:type="StringProperty" Name="PersistOsDisk" Value="true"
11    />
12 </CustomProperties>'
13 <!--NeedCopy-->
```

For more configuration information, see [Improve boot performance](#).



## July 2020

### New and enhanced features

**Support for granular, role-based access to the Filters page.** Citrix Studio now provides more granular control over access to the **Monitor > Filters** page when you create a custom role. Specifically, you can assign permissions to view any combination of **Machines, Sessions, Connections, and Application Instances** to a custom role. The following are four more options for the **Director** object in the **Create Role** window:

- View Filters page - Application Instances only
- View Filters page - Connections only
- View Filters page - Machines only
- View Filters page - Sessions only

For information about creating roles, see [Create and manage roles](#).

**Power-off delay support for assigned VDI machines (PowerShell only).** In earlier releases, the power-off delay applied only to unassigned machines. Starting with this release, the power-off delay applies to both assigned and unassigned machines. For more information, see [How Autoscale power manages machines](#).

**Support for Windows Client licenses.** The Citrix Virtual Apps and Desktops service now supports using Windows Client licenses to provision VMs in Azure. To run Windows 10 VMs in Azure, verify that your volume licensing agreement with Microsoft qualifies for this usage. For more information, see [Create a machine catalog using an Azure Resource Manager master image](#).

## May 2020

### New and enhanced features

**Machine restart schedules.** You can now indicate whether a restart schedule affects machines that are in maintenance mode. This feature is available only in PowerShell. For details, see [Scheduled restarts for machines in maintenance mode](#).

**Resource availability.** You can now ensure resource availability during an outage without having to publish resources in every zone (resource location). For details, see [Resource availability](#).

## April 2020

### New and enhanced features

**Enhanced scheduling granularity for VDI Delivery Groups (PowerShell only).** Autoscale now supports defining the peak times for the days included in a schedule at a granular level of 30 minutes. You can set the minimum number of machines running in a VDI Delivery Group separately for each

half hour of the day. Also, Autoscale can now scale up or down the number of powered-on machines in VDI Delivery Groups on a half-hourly basis instead of on an hourly basis. For more information, see [Broker PowerShell SDK commands](#).

**MTU Discovery.** The Citrix protocol Enlightened Data Transport (EDT) now has MTU Discovery capabilities. MTU Discovery allows EDT to automatically determine and set the payload size for the session. This feature enables the ICA session to adjust to networks with non-standard Maximum Transmission Unit (MTU) or Maximum Segment Size (MSS) requirements. The ability to adjust avoids packet fragmentation that might result in degraded performance or failure to establish an ICA session. This update requires a minimum of Citrix Workspace app 1911 for Windows. If using Citrix Gateway, the minimum Citrix ADC firmware version required is 13.0.52.24 or 12.1.56.22. For more information, see [EDT MTU Discovery](#).

## March 2020

### New and enhanced features

**PVS target device metrics.** Citrix Monitor now provides a PVS target device metrics panel on the Machine Details page. Use the panel to view the status of Provisioning target devices for single-session and multi-session OS machines. Several metrics for Network, Boot, and Cache are available on this panel. These metrics help you monitor and troubleshoot PVS target devices to ensure that they are up and running. For more information, see [PVS target device metrics](#).

**AWS instance property capturing.** MCS now reads properties from the instance from which the AMI was taken and applies the IAM role and tags of the machine to the machines provisioned for a given catalog. When using this optional feature, the catalog creation process finds the selected AMI source instance, reading a limited set of properties. These properties are then stored in an AWS Launch Template, which is used to provision machines for that catalog. Any machine in the catalog inherits the captured instance properties. For more information, see [AWS instance property caching](#).

**AWS operational resource tagging.** This release introduces an option to tag resources created by Citrix components during provisioning. Each tag represents a label consisting of a customer-defined key and an optional value that improve your ability to manage, search for, and filter resources. For more information, see [AWS operational resource tagging](#).

**Secure transfer in Azure storage.** Machine Creation Services (MCS) provides an enhancement for storage accounts created by MCS-provisioned catalogs in Azure Resource Manager environments. This enhancement automatically enables the secure transfer required property. This option enhances storage account security by only allowing requests to the account from secure connections. For more information, see [Require secure transfer to ensure secure connections](#) on the Microsoft site.

Enable the **Secure transfer required** property when creating a storage account in Azure:

**Create storage account**
✕

---

Basics
**Advanced**
Tags
Review + create

---

**SECURITY**

Secure transfer required ⓘ  Disabled  Enabled

**VIRTUAL NETWORKS**

Allow access from  All networks  Selected network  
 ⓘ All networks will be able to access this storage account. [Learn more](#)

**DATA LAKE STORAGE GEN2 (PREVIEW)**

Hierarchical namespace ⓘ  Disabled  Enabled

Review + create

Previous

Next: Tags >

**Support for Azure SSD managed disks.** Machine Creation Services (MCS) supports standard SSD managed disks for Azure virtual machines. This disk type provides consistent performance, and delivers better availability compared to HDD disks. For more information, see [Standard SSD Disks for Azure Virtual machine workloads](#).

Use the PowerShell `StorageAccountType` custom property in the `New-ProvScheme` command or `Set-ProvScheme` command to configure this feature:

```

1 <Property xsi:type="StringProperty" Name="UseManagedDisks" Value="true" /><Property xsi:type="StringProperty" Name="StorageAccountType" Value="StandardSSD_LRS" /><Property xsi:type="StringProperty" Value="Windows_Server" />
2 <!--NeedCopy-->

```

**Note:**

This feature is only available when using managed disks, that is, the custom property `UseManagedDisks` is set to **true**. For unmanaged disks only Standard HDD and Premium SSD are supported.

## January 2020

### New and enhanced features

**Language bar in Citrix Studio.** Starting with this release, Citrix Studio provides a language bar to facilitate correct keyboard mapping.

- If the language of Citrix Cloud or the display language of your browser is set to **English** or **Japanese**, the language bar does not appear.
- If the language of Citrix Cloud or the display language of your browser is set to **German**, **Spanish**, or **French**, the language bar appears after you log on to Citrix Studio. There are two language options on the language bar list. Select an option that matches the topmost language of your browser.

**Tip:**

- Settings that you configure for the language bar might not take effect. In this case, log out and log back on.
- You might fail to input certain symbols and localized characters by using the language bar. To resolve the issue, you need to configure the language of Citrix Cloud, the display language of your browser, and the local keyboard layout. For more information, see Knowledge Center article [CTX310743](#).

**Restart schedule maximum delay timer (PowerShell only).** If a scheduled restart of machines in a Delivery Group doesn't begin because of a site database outage, you can specify how long to wait beyond the scheduled start time. If the database connection is restored during that interval, the restarts begin. If the connection isn't restored during that interval, the restarts don't begin. For details, see [Scheduled restarts delayed due to database outage](#).

**Vertical load balancing (PowerShell only).** Previously, the service used horizontal load balancing for all RDS launches, which assigns incoming load to the least-loaded RDS machine. That remains the default. Now, you can use PowerShell to enable vertical load balancing as a site-wide setting.

When vertical load balancing is enabled, the broker assigns incoming load to the most-loaded machine that has not reached a high watermark. This saturates existing machines before moving on to new machines. As users disconnect and free up existing machines, new load is assigned to those machines.

By default, horizontal load balancing is enabled. To view, enable, or disable vertical load balancing, the `Get-BrokerSite`, and `Set-BrokerSite` cmdlets now support the `UseVerticalScalingForRdsLaunches` setting. For more information, see [Load manage machines in Delivery Groups](#).

## December 2019

### New and enhanced features

**Service for Citrix Service Providers (CSP).** CSPs can now onboard tenant customers to the Virtual Apps and Desktops service, configure customer administrator access to the service, and provide shared or dedicated workspaces to customers' users using federated domains. For more information, see [Citrix Virtual Apps and Desktops service for Citrix Service Providers](#).

**Support for determining why a machine is in maintenance mode (PowerShell only).** Using PowerShell, you can now determine why a machine is in maintenance mode. To do so, use the parameter `-MaintenanceModeReason`. The feature is useful for administrators to troubleshoot issues with machines in maintenance mode. For details, see <https://developer-docs.citrix.com/projects/citrix-virtual-apps-desktops-sdk/en/latest/Broker/Get-BrokerMachine/>.

**Autoscale.** Autoscale now provides the capability to create machines and delete them dynamically. You can use the capability by using a PowerShell script. The script helps you dynamically scale up or down the number of machines in the Delivery Group based on the current load conditions. For more information, see [Dynamically provision machines with Autoscale](#).

## November 2019

### New and enhanced features

**GroomStartHour.** Monitor now supports **GroomStartHour** - a new configuration that helps administrators to determine the time of the day that grooming should start running. For more information, see the [Citrix Virtual Apps and Desktops SDK](#) documentation.

**OData Pagination.** Monitor now supports **OData pagination**. All OData v4 endpoints return a maximum of 100 records per page with a link to the next 100 records in the response. For more information, see [Accessing Monitor Service data using the OData v4 endpoint in Citrix Cloud](#).

## October 2019

### New and enhanced features

**App-V.** App-V functionality is now available in Citrix Cloud. You can add App-V packages to the Delivery Controller in your Citrix Cloud configuration, in either single- or dual admin mode. The *Virtual Apps and Desktops Service App-V package discovery module*, available in [Citrix Downloads](#), allows you to import App-V packages and register Microsoft App-V servers. The apps they contain are then available to your users. This PowerShell module allows you to register Microsoft App-V Management and Publishing Servers using DNS URLs, avoiding the need for servers behind load balancing mechanisms to be registered using their actual machine URL. For more information, see [Citrix Virtual Apps and Desktops service discovery module for App-V packages and servers](#).

**Google Cloud Platform.** Citrix Virtual Apps and Desktops service now adds support for using Machine Creation Services (MCS) to provision machines on the Google Cloud Platform (GCP). For more information, see [Google Cloud Platform virtualization environments](#).

## September 2019

### New and enhanced features

**VDA support for Azure Virtual Desktop.** For supported operating systems and VDA versions see [VDAs in an Azure Virtual Desktop environment](#).

**Enhanced power policy.** In earlier releases, a VDI machine transitioning to a time period where an action (disconnect action=“**Suspend**” or “**Shutdown**”) was required remained powered on. This scenario occurred if the machine disconnected during a time period (peak or off-peak times) where no action (disconnect action=“**Nothing**”) was required.

Starting with this release, Autoscale suspends or powers off the machine when the specified disconnection time elapses, depending on the disconnect action configured for the destination time period. For more information, see [Power manage VDI machines transitioning to a different time period with disconnected sessions](#).

**Machine catalogs: Tags.** You can now use PowerShell to apply tags to machine catalogs. For more information, see [Apply tags to machine catalogs](#).

**Session startup duration.** Monitor now displays the session startup duration divided into Workspace App Session Startup and VDA Session Startup time periods. This data helps you to understand and troubleshoot high session startup duration. Further, the time duration for each phase involved in the session startup helps in troubleshooting issues associated with individual phases. For example, if the Drive Mapping time is high, you can check if all the valid drives are mapped properly in the GPO or script. This feature is available on VDAs 1903 or later. For more information, see [Diagnose session startup issues](#).

## August 2019

### New and enhanced features

**Session Auto Reconnect.** The Sessions page on the Trends tab now includes information about the number of auto reconnects. Auto reconnects are attempted when the Session Reliability or Auto Client Reconnect policies are in effect. The auto reconnect information helps you view and troubleshoot network connections having interruptions, and also analyze networks having a seamless experience.

The drilldown provides additional information like Session Reliability or Auto Client Reconnect, time stamps, Endpoint IP, and Endpoint Name of the machine where the Workspace app is installed. This

feature is available for Citrix Workspace app for Windows, Citrix Workspace app for Mac, Citrix Receiver for Windows, and Citrix Receiver for Mac. This feature requires VDAs 1906 or later. For more information, see:

- [Sessions](#)
- [Auto client reconnect policy settings](#)
- [Session reliability policy settings](#)
- [Session Auto Reconnect](#)

## July 2019

### New and enhanced features

**Configuration Logging.** You can now use the Remote PowerShell SDK to periodically delete Configuration Logging database content. For details, see [Schedule periodic data deletion](#).

**Autoscale.** Autoscale now provides the flexibility to power manage only a subset of machines in a Delivery Group. This feature can be useful in cloud bursting use cases, where you want to use on-premises resources to handle workloads before cloud-based resources address other demands (that is, burst workloads). For more information, see [Restrict Autoscale to certain machines in a Delivery Group](#).

**Local App Access and URL redirection.** Citrix Studio now lets you add the Add Local App Access Application option to the Studio user interface for your Site by using the PowerShell SDK. For more information, see [Provide access only to published applications](#).

**Operating system name changes.** Operating system names on the **Create Machine Catalog > Machine Catalog Setup > Operating System** and the **Monitor** pages have changed:

- Multi-session OS (formerly Server OS): The multi-session OS machine catalog provides hosted shared desktops for a large-scale deployment of standardized Windows multi-session or Linux OS machines.
- Single-session OS (formerly Desktop OS): The single-session OS machine catalog provides VDI desktops ideal for various users.

**Citrix Profile Management duration in Profile Load.** Monitor now includes profile processing duration in the Profile Load bar of the Logon duration chart. This is the duration Citrix Profile Management takes to process user profiles. This information helps administrators to troubleshoot high profile load durations with greater accuracy. This enhancement is available on VDAs 1903 and later. For more information, see [Profile Load](#).

**Desktop probing.** Desktop probing is a feature of the Citrix Virtual Apps and Desktops service. It automates health checks of virtual desktops published on a site, which improves user experience. To initiate desktop probing, install and configure the Citrix Probe Agent on one or more endpoints.

Desktop probing is available for Premium licensed Sites. This feature requires Citrix Probe Agent 1903 or later. For more information, see [Desktop Probing](#).

**Note:**

Citrix Probe Agent now supports TLS 1.2.

## June 2019

### New and enhanced features

**Restrict by tags.** Tags are strings that identify items such as machines, applications, desktops, Application Groups, and policies. After creating a tag and adding it to an item, you can tailor certain operations to apply to only items that have a specified tag. For more information, see [Application Groups](#) and [Tags](#).

**Email Notifications.** Citrix Virtual Apps and Desktops service sends email notifications related to alerting and probing directly. This eliminates the need to configure the SMTP email server. The **Notification Preferences** box is enabled by default and Citrix Cloud sends alert notifications to the email addresses provided in the **Notification Preferences** section. Ensure that the email address [donotreplynotifications@citrix.com](mailto:donotreplynotifications@citrix.com) is white-listed in your email setup.

## May 2019

### New and enhanced features

**Autoscale.** Autoscale is a feature of Citrix Virtual Apps and Desktops service that provides a consistent, high-performance solution to proactively power manage your machines. It aims to balance costs and user experience. Autoscale incorporates the deprecated Smart Scale technology into the Studio power management solution. For more information, see [Autoscale](#). You can monitor the metrics of Autoscale-managed machines from the Trends pages on the **Monitor** tab. For more information, see [Monitor Autoscale-managed machines](#).

## February 2019

### New and enhanced features

**Hypervisor alerts monitoring.** Alerts from Citrix Hypervisor and VMware vSphere are now displayed on the **Monitor > Alerts** tab to help monitor the following states/parameters of hypervisor health:

- CPU usage
- Memory usage
- Network usage
- Hypervisor connection unavailable



- Disk usage (vSphere only)
- Host connection or power state (vSphere only)

For more information, see the Hypervisor alerts monitoring section in [Alerts and Notifications](#).

**Communications over earlier TLS versions.** To improve the security of the service, Citrix will block any communication over Transport Layer Security (TLS) 1.0 and 1.1 as of March 15, 2019, allowing only TLS 1.2 communications. For more information, see [TLS versions](#). For comprehensive guidance, see [CTX247067](#).

**Application Groups.** Application Groups let you manage collections of applications. You can create Application Groups for applications shared across different Delivery Groups or used by a subset of users within Delivery Groups. For more information, see [Create Application Groups](#).

**Logon Performance - Profile Drilldown.** The **Logon Duration** panel on the **User Details** page within **Monitor** now includes information on the **Profile load phase** drilldown of the logon process. Profile drilldown provides useful information about user profiles for the current session that can help administrators troubleshoot high profile load issues. A tool tip with the following user profiles information is displayed:

- Number of files
- Profile size
- Number of large files

A detailed drill-down provides information about the individual folders, their size, and the number of files. This feature is available on VDAs 1811 and later. For more information, see [Diagnose user logon issues](#).

**Microsoft RDS license health.** Monitor the Microsoft RDS (Remote Desktop Services) license status on the **Machine Details** panel on the Machine Details and the User Details page for Server OS machines. An appropriate message is displayed for the license status. You can hover over the info icon to see further details. For more information, see the Microsoft RDS license health section in [Troubleshooting Machines](#).

**Application Probing.** This feature automates the assessment of the health of Virtual Apps published in a Site.

To initiate application probing:

- On one or more endpoint machines, install the Citrix Application Probe Agent
- Configure the Citrix Application Probe Agent with the credentials of Citrix Workspace and Citrix Virtual Apps and Desktops service.
- Configure the applications to be probed, the endpoint machines to run the probe on, and the scheduled probe time in **Monitor > Configuration** of the Citrix Virtual Apps and Desktops service.

The agent tests the launch of selected applications via Citrix Workspace and reports back the probe results on the **Monitor** tab of Citrix Virtual Apps and Desktops service in:

- the Applications page – the last 24-hours' data and the **Trends > Application Probe Results** page
- the historical probe data along with the stage when the probe failure occurred - Workspace Reachability, WorkspaceAuthentication, WorkspaceEnumeration, ICA download, or Application launch

The failure report is emailed to configured email addresses. You can schedule application probes to run during off-peak hours across multiple geographical locations. In that way, you can use the results to proactively troubleshoot issues related to provisioned applications, hosting machines, or connections before the users experience them. For more information, see [Application Probing](#)

## January 2019

### New and enhanced features

**Delegated Administration with custom scope.** Monitoring now supports custom scope for built-in delegated administrator roles. For more information on the available built-in roles for monitoring and how to assign them, see [Delegated administrator roles](#).

## December 2018

### New and enhanced features

The date after which Citrix will block communication over Transport Layer Security (TLS) 1.0 and 1.1 has changed from December 31, 2018 to January 31, 2019. For details, see [Deprecation of TLS versions](#).

## November 2018

### New and enhanced features

**Machine historical data available using OData API:** Historical data containing machine analytics is now available through the OData API. This data is collected on an hourly basis and rolled up for the day.

- Number of powered on machines (for power managed machines)
- Number of registered machines
- Number of machines in maintenance mode
- Total number of machines

The data is aggregated for the time period during which the Monitoring Service is running. For more information on the usage of the OData API and examples, see [Citrix Monitor Service 7 1808](#). The database schema is available at [Monitor Service Schema](#).

**Logon Performance - Interactive session drilldown:** The **Logon Duration** panel on the **User and Session Details** view includes information on the **Interactive Session** phase of the logon process. The time taken for each of the three subphases (**Pre-userinit**, **Userinit**, and **Shell**) is displayed on the **Interactive Session** bar as a tooltip. This provides more granular troubleshooting and remediation of this phase of the logon. The cumulative time delay between the subphases and a link to the documentation is also provided. This feature is available on Delivery Controller version 7 1808 and later. The **Interactive Session** drilldown bar shows the time duration for the current session only. For more information, see [Diagnose user logon issues](#).

**Logon Performance - GPO drilldown:** The **Logon Duration** panel on the **User and Session** details view contains the GPO (Group Policy Objects) duration. This is the total time taken to apply the GPOs on the virtual machine during the logon process. Now, you can see the drilldown of each policy applied as per CSEs (Clients-Side Extension) as a tool tip on the GPO bar. For each policy application, the drilldown displays the status and the time taken. This additional information eases troubleshooting and remediation of issues involving high GPO duration. The time durations in the drilldown represent the CSE processing time only and do not add up to the total GPO time. This feature is available on Delivery Controller version 7 1808 and later. For more information, see [Diagnose user logon issues](#).

## Fixes

Custom report queries saved during monitoring are not available after a Cloud upgrade. [DNA-23420]

## October 2018

### New and enhanced features

**Applications: Limit per machine.** You can now limit the number of application instances per machine. This limit applies to all machines in the Site. This limit is an addition to the existing application limit for all users in the Delivery Group and the limit per user. This capability is available only through PowerShell, not in Studio. For details, see [Configure application limits](#).

**Windows Server 2019.** You can now install VDAs for multi-session OS (formerly VDAs for Server OS) on Windows Server 2019 machines, as noted in [System requirements](#).

## September 2018

### New and enhanced features

**Delegated Administration.** With Delegated Administration, you can configure the access permissions that all of your administrators need, in accordance with their role in your organization. For details, see [Delegated Administration](#). Monitoring supports allocation of built-in roles. Built-in roles are available with full scope. For more information about built-in roles for monitoring and how to assign them, see [Delegated administrator roles](#).

**Configuration Logging.** Configuration Logging allows administrators to keep track of configuration changes and administrative activities. For details, see [Configuration Logging](#).

Several PowerShell cmdlets in the Remote PowerShell SDK that were previously disabled are now enabled, for use with Configuration Logging:

- Log:GetLowLevelOperation
- Log:GetHighLevelOperation
- Log:GetSummary
- Log:GetDataStore
- Log:ExportReport

**Local Host Cache.** Local Host Cache is now fully available. Local Host Cache enables connection brokering operations to continue when a Cloud Connector in a resource location cannot communicate with Citrix Cloud. For details, see [Local Host Cache](#).

**Citrix Provisioning.** To provision VDAs, you can now use Citrix Provisioning or the existing Machine Creation Services. For Citrix Provisioning information specific to the cloud environment, see [Citrix Provisioning managed by Citrix Cloud](#).

### Fixes

In earlier versions, when using Azure on-demand provisioning, all VMs were deleted when powered-off. Now, only pooled VMs are deleted. Persistent (dedicated) VMs are not deleted when powered-off.

## August 2018

### • New product names

If you've been a Citrix customer or partner for a while, you'll notice new names in our products and in this product documentation. If you're new to this Citrix product, you might see different names for a product or component.

The new product and component names stem from the expanding Citrix portfolio and cloud strategy. Articles in this product documentation use the following names.

- **Citrix Virtual Apps and Desktops:** Citrix Virtual Apps and Desktops offers a virtual app and desktop solution, provided as a cloud service and as an on-premises product, giving employees the freedom to work from anywhere on any device while cutting IT costs. Deliver Windows, Linux, web, and SaaS applications or full virtual desktops from any cloud: public, on-premises, or hybrid. Virtual Apps and Desktops was formerly XenApp and XenDesktop.
- **Citrix Workspace app:** The Citrix Workspace app incorporates existing Citrix Receiver technology and other Citrix Workspace client technologies. It has been enhanced to deliver more capabilities to provide end users with a unified, contextual experience where they can interact with all the work apps, files, and devices they must do their best work. For more information, see this blog post.
- **Citrix SD-WAN:** NetScaler SD-WAN, a crucial technology for our customers and partners transforming their branch networks and WANs with cloud technology, is now Citrix SD-WAN.
- **Citrix Secure Web Gateway:** As the Citrix Networking portfolio expands, we're proud to offer our robust Citrix Secure Web Gateway Service, previously known as NetScaler Secure Web Gateway.
- **Citrix Gateway:** Our robust NetScaler Unified Gateway, which allows secure, contextual access to the apps and data you must do your best work, is now Citrix Gateway.
- **Citrix Content Collaboration and Citrix Files for Windows:** The advanced access, collaboration, workflows, rights management, and integration features of ShareFile are now available in the Citrix Content Collaboration component set in our secure, contextual, integrated Citrix Workspace. Citrix Files for Windows allows you to access your Content Collaboration files directly through a mapped drive, providing a native Windows Explorer experience.
- **Citrix Hypervisor:** The technology from XenServer for virtualization infrastructure, based on the XenProject hypervisor, is now Citrix Hypervisor.

Here's a quick recap:

Is	Was
Citrix Virtual Apps and Desktops	XenApp and XenDesktop
Citrix Workspace app	Incorporates Citrix Receiver and extensive enhancements
Citrix SD-WAN	NetScaler SD-WAN
Citrix Secure Web Gateway	NetScaler Secure Web Gateway
Citrix Gateway	NetScaler Unified Gateway
Citrix Content Collaboration	ShareFile

Is	Was
Citrix Files for Windows	ShareFile Desktop App, ShareFile Sync, ShareFile Drive Mapper
Citrix Hypervisor	XenServer
Citrix Provisioning	Citrix Provisioning Services

Implementing this transition in our products and their documentation is an ongoing process.

- In-product content might still contain former names. For example, you might see instances of earlier names in console text, messages, and directory/file names.
- It is possible that some items (such as commands and MSIs) might continue to retain their former names to prevent breaking existing customer scripts.
- Related product documentation and other resources (such as videos and blog posts) that are linked from this product's documentation might still contain former names.
- For Citrix Hypervisor: The new name is used on the Citrix website and in informational product materials from September 2018. You will also see the new name on the administrator consoles of some Citrix products, such as Citrix Virtual Apps and Desktops. The XenServer product release and technical documentation materials continue to use XenServer 7.x until early 2019.

Your patience during this transition is appreciated.

For more detail about our new names, see <https://www.citrix.com/about/citrix-product-guide/>.

#### • **Product and component version number changes**

Citrix installs and manages most of the Citrix Virtual Apps and Desktops components, so you won't be concerned with those version numbers. However, you might see version numbers when installing Cloud Connectors, and when installing or upgrading VDAs in resource locations.

Citrix Virtual Apps and Desktops product and component version numbers are displayed in the format: **YYMM.c.m.b**

- YYMM = Year and month when the product or component released. For example, a September 2018 release appears as 1809.
- c = Citrix Cloud release number for the month.
- m = Maintenance version (if applicable).
- b = Build number. This field is shown only on the About page of the component, and in the OS's feature for removing or changing programs.

For example, **Citrix Virtual Apps and Desktops 1809.1.0** indicates that the component released in September 2018. It is associated with Citrix Cloud release 1 in that month, and

is not a maintenance version. Some displays show only the version's year and month: for example, **Citrix Virtual Apps and Desktops 1809**.

In earlier releases (7.18 and earlier), version numbers were displayed in the format: *7.version*, where version incremented by one for each release. For example, the VDA release following XenApp and XenDesktop 7.17 was 7.18. Earlier releases (7.18 and earlier) will not be updated with the new numbering format.

- **Deprecation of TLS versions.** To improve the security of the Citrix Virtual Apps and Desktops service, Citrix will block any communication over Transport Layer Security (TLS) 1.0 and 1.1, effective December 31, 2018. For details, see [Deprecation of TLS versions](#).
- **Google Cloud Platform virtualization environment.** The Citrix Virtual Apps and Desktops service supports the ability to manually power cycle Virtual Apps and Desktops VMs on the Google Cloud Platform (GCP). For more information, see [Google Cloud Platform virtualization environments](#).

## July 2018

- **Export of Filters data.** You can now export real-time monitoring data on the **Monitor > Filters** tab to CSV format files. The export feature is available from the Machines, Sessions, Connections, and Application Instances Filters pages. You can select a predefined custom filter or select suitable filter criteria, choose required columns on the table, and export the data. Data of up to 100,000 records can be exported. The exported CSV files give a comprehensive view of the real-time data, and helps ease analysis of large data sets.

## June 2018

- **Azure Resource Manager connections.** In the Studio connection creation wizard, the Azure environment selection on the **Connection** page includes all Azure Clouds that are valid for your Azure subscription. General availability for Azure US Government Cloud and Azure Germany Cloud replaces the preview versions of those two environments in earlier releases.

## May 2018

- **Azure Quick Deploy.** When your resource location uses Azure Resource Manager machines to deliver applications and desktops, you can now choose a deployment method:
  - Full Configuration: This existing method uses the Studio management console, which guides you through creating a machine catalog and then creating a Delivery Group.
  - Azure Quick Deploy: This new option offers a simpler interface that offers faster deployment of apps and desktops.

- **Citrix Health Assistant link.** The Machine Details page of an unregistered machine on the Monitoring console now contains a **Health Assistant** button. Currently, the button links to [Troubleshoot machines](#) and to the Knowledge Center article, [Citrix Health Assistant - Troubleshoot VDA Registration and Session Launch](#) where you can download the tool. Citrix Health Assistant is a tool to troubleshoot configuration issues in unregistered VDAs. The tool automates several health checks to identify possible root causes for common VDA registration, session launch, and time zone redirection configuration issues.
- **Interactive Session drilldown.** In the monitoring console, the **User Details view > Logon Duration** panel now includes information on the **Interactive Session** stage of the logon process. To provide more granular troubleshooting and remediation of this phase of the logon, **Interactive Session** now has three subphases: **Pre-userinit**, **Userinit**, and **Shell**. In this release, hovering over **Interactive Session** displays a tooltip showing the subphases and a link to the documentation. For a description of the subphases and how to improve the performance of each phase, see [Diagnose user logon issues](#).

## March 2018

- **Application instance prediction (Preview feature).** This is the first monitoring feature based on predictive analytics. Predicting resource usage patterns is important for administrators to organize resources and the required number of licenses on each resource. The Application instance prediction feature indicates the number of hosted application instances that are likely to be launched per Site or Delivery Group over time. Machine learning algorithms based on data models created with existing historical data are used to do the prediction. Tolerance level indicates the prediction quality.

For more information see [Application instance prediction](#) in Director. Submit your feedback regarding the usefulness and usability of this feature in the [Citrix Cloud discussion forum](#).

- **Delivery Groups APIs - Preview**

The Delivery Groups APIs Preview provides a set of REST APIs that you can use to automate the management of Delivery Groups. The complete set of available APIs can be viewed and tried out in the Citrix Cloud API documentation at <https://developer.cloud.com/>.

- **Web Studio authentication**

The service management console on Citrix Cloud now uses a bearer token to authenticate customers. The bearer token is required to authenticate access to the Delivery Groups REST API.

- **Access Monitor Service data using OData Version 4 API (Preview feature)**

You can create your customized monitoring and reporting dashboards based on the Monitor Service data by using the OData V.4 endpoint. OData V.4 is based on ASP .Net Web API and supports aggregation queries. Use your Citrix Cloud user name and bearer token to access the data with



the V4 endpoint. For more information and examples, see [Access Monitor Service data using the OData v4 endpoint in Citrix Cloud](#).

Share your feedback on the utility of this feature in the [Citrix Cloud discussion forum](#).

## Fixes

- You can rename, move, and delete application folders. [#STUD-2376]

## January 2018

- **RDS license check.** Creation of machine catalogs containing Windows Server OS machines now includes an automatic RDS license check. Any RDS license issues found are displayed, so that you can take the appropriate steps to prevent a gap in service. For details, see [Create machine catalogs](#).
- **Access to machine console from Monitor.** The Machine Details panel from Monitor now provides access to consoles of machines hosted on the XenServer hypervisor version 7.3. You can now troubleshoot issues in VDAs directly from Monitor. For more information, see [Machine Console access](#) in Troubleshoot machines.

## December 2017

### New and enhanced features

- **Citrix Workspace.** Citrix Workspace is now available for **new** XenApp and XenDesktop Service customers. For more information, see [Workspace Configuration](#).
- **Applications Analytics.** You can now analyze and monitor the performance of applications efficiently with the new Application Analytics page available from **Monitor > Applications** tab. The page provides a consolidated view of the health and usage of all applications published on your Site. It shows metrics such as the number of instances per application, and faults and errors associated with the published applications. This feature requires VDAs Version 7.15 or later.

For more information, see [Application Analytics](#) section in Monitor.

## November 2017

### New and enhanced features

- **Local Host Cache.** Local Host Cache enables connection brokering operations to continue when a Cloud Connector in a resource location cannot communicate with Citrix Cloud. For details, see [Local Host Cache](#).

- **Azure Managed Disks.** Azure Managed Disks are now used by default for MCS-provisioned VMs in Azure Resource Manager environments. Optionally, you can use conventional storage accounts. For details, see [Microsoft Azure Resource Manager virtualization environments](#).
- **Help desk administrator.** When managing service administrators for a Citrix Cloud customer account, you now have a new choice: Help Desk Administrator. A help desk administrator can access the Monitor functions on the service. For details, see [Manage](#).

## Fixes

- You can now use the service management console wizard to create a Remote PC Access machine catalog. In earlier releases, you had to use a PowerShell cmdlet to create a catalog (as documented in [CTX220737](#)). Then, you had to return to the management console to create a Delivery Group. Now, you create the catalog and the Delivery Group sequentially on the management console.
- MCS-created catalogs can use existing Active Directory machine accounts. [#DNA-24566]
- When monitoring a deployment, scrolling in a sorted **Trends > Sessions** table displays accurate results. [DNA-51257]

## More information

- [Known issues](#).
- For information about third-party software that is included in the service, see [Third party notifications](#).

## Known issues

September 2, 2021

The Citrix Virtual Apps and Desktops service has the following known issues:

- In AWS environments, launching and terminating volume worker instances fails to remove their associated network interfaces. To resolve this issue, manually delete network interfaces with conditions matching the following state: `Available && Description: "XD NIC"&& tag : "XdConfig : XdProvisioned=true"`. [PMCS-20775]
- The current hypervisor communication implementation, *Remote HCL*, might throw exceptions by the target hypervisor platform. As a result, the connection between the cloud controller and the cloud connector fails and is then recreated. If any other Remote HCL operations are in progress and using the same connection, those connections can also fail. This causes machine power and registration states to fall out of sync. Other problems can arise as a result because

the issue affects all types of Remote HCL operations, not just power states. Azure and GCP hypervisors hosting connections are not affected. These connections do not use Remote HCL. [CCVADHELP-483]

- VMware machines fail to restart and cannot be forcefully restarted. This issue applies to all versions of VMware, including VMC on AWS. The problem occurs in machine catalogs that have persistent (dedicated) VMs, or, VMs that are power managed. To resolve this issue, use the [New-Brokerhostingpoweraction](#) cmdlet to restart or force restart your machines. [PMCS-15797]
- The drop-down arrow icon for the Average IOPS, Session Control, and Power Control buttons might not appear on the **User Details** and the **Machine Details** pages. However, the functionality works as expected. To view all the items on the menu, click anywhere on the button. [DIR-11875]
- If you use Azure AD Domain Services: Workspace (or StoreFront) logon UPNs must contain the domain name that was specified when enabling Azure AD Domain Services. Logons cannot use UPNs for a custom domain you create, even if that custom domain is designated as primary.
- When deploying to Azure and creating an MCS catalog version 7.9 or later with write-back cache enabled and the VDA installed on the master image is 1811 or earlier, an error occurs. Also, you cannot create anything related to Personal vDisk for Microsoft Azure. As a workaround, select a different catalog version to deploy to Azure, or disable write-back cache. To disable write-back cache when you create a catalog, clear the **Memory allocated to cache** and **Disk cache size** check boxes on the **Machines** page.
- The **Console** link on **Monitor > Machine Details** does not launch the Machine Console in the Microsoft Edge 44 and Firefox 68 ESR browsers. [DIR-8160]
- Changing the name of an AWS Virtual Private Cloud (VPC) in the AWS console breaks the existing hosting unit in Citrix Cloud. When the hosting unit is broken, you cannot create catalogs or add machines to existing catalogs. [PMCS-7701]
- When you try to use the 'Restart' option in Workspace App web or desktop, the 'Restarting' dialogue never closes and never reports success. The hypervisor shows the machine has shut down but has not started. As a workaround, after some time the user can close the 'Restarting' dialogue and launch the desktop and the desktop will start. [BRK-5564]
- When you deploy machines in an MCS catalog, the provisioning task can fail and the following error message appears: "Terminating Error: Desktop Studio closed." The error details might show that no AD accounts were created. The catalog might complete successfully later without intervention. The issue is seen in large, complex deployments. [PMCS-8869]
- In a Mozilla Firefox browser, when you use the **Full Configuration** management interface to configure policy settings or assign the policy, the wizard might not behave as expected. If you type Japanese characters through an Input Method Editor (IME) and then press the **Enter** key to

finish, you exit the current window with your changes saved. The issue occurs because pressing the **Enter** key also triggers the shortcut key used to exit the window. [GP-912]

- Cloud Library cannot be used to assign resources in deployments that include on-premises StoreFront. [CCVADHELP-625]

For issues related to current VDAs, see [Known issues](#).

## Deprecation

July 21, 2021

This article gives you advanced notice of Citrix Virtual Apps and Desktops service features that are being phased out, so that you can make timely business decisions. Citrix monitors customer use and feedback to determine when features are withdrawn. Announcements can change in subsequent releases and might not include every deprecated feature or functionality. For details about product lifecycle support, see the [Product Lifecycle Support Policy](#) article.

### Note:

Citrix Virtual Apps and Desktops deprecations and removals are described in their own [Deprecation](#) article.

## Deprecations and removals

The following list shows the Citrix Virtual Apps and Desktops service features that are deprecated or removed.

*Deprecated* items are not removed immediately. Citrix continues to support them but they will be removed in a future release.

*Removed* items are removed, or no longer supported, in the Citrix Virtual Apps and Desktops service. Dates in **bold** face indicate the latest updates.

Item	Deprecation announced in release	Removed in release	Alternative
Legacy console (MMC-based console)	July 2021	Target: September 30, 2021	Use <b>Mange &gt; Full Configuration</b> to access the full range of configuration and management actions.

Item	Deprecation announced in release	Removed in release	Alternative
Azure Quick Deploy	September 2020		Use <a href="#">Quick Deploy</a> .
Ability to import Citrix Provisioning target devices to create catalogs in cloud Studio.	August 2020	February 2021	Use the Citrix Provisioning Export Devices Wizard to push Citrix Provisioning VMs into Delivery Controllers/MCS for catalog creation.

## System requirements

September 1, 2021

### Introduction

System requirements for components that are not covered here (such as Citrix Workspace app and Citrix Provisioning) are described in their respective documentation.

Specific recommendations for sizing VMs that deliver desktops and applications cannot be provided because of the complex and dynamic nature of hardware offerings. Every deployment has unique needs. Generally, sizing a VM is based on the hardware and not the user workloads (except for RAM; you need more RAM for applications that consume more.) The [Citrix VDI Handbook and Best Practices](#) contains the latest guidance on VDA sizing.

**Remember:** In a Citrix Virtual Apps and Desktops service deployment, you don't need to install or manage the core components (Delivery Controllers, the site database, or management and monitoring consoles). For Virtual Delivery Agent (VDA) installation guidance, see:

- [Install VDAs](#)
- [Install VDAs using the command line.](#)

### Cloud Connectors

For details, see [Cloud Connector Technical Details](#).

## VDA in an Azure Virtual Desktop environment

Supported operating systems:

- Windows 10 multi-session
- Windows 10 single-session
- Windows 7 single-session
- [Windows Server 2022](#) (requires minimum VDA 2106)
- Windows Server 2019
- Windows Server 2016
- Windows Server 2012 R2

### Notes:

All currently supported VDA versions are supported for use with the Citrix Virtual Apps and Desktops service. See the [Citrix Product Matrix](#) for details about supported versions.

Windows 7 is supported only with VDA 7.15 CU5 (or later CUs).

Windows Server 2012 R2 is supported only with VDA 7.15 CU5 (and later CUs) or VDA 1912 (and later CUs).

Windows Server requires [Microsoft RDS licensing](#).

For information about Azure Virtual Desktop, see the Microsoft [documentation](#).

## VDA for single-session OS

The following information applies to the latest VDA release.

Supported operating systems:

- Windows 10
  - For edition support, see [CTX224843](#). That article also contains links to Citrix known issues with the supported Windows versions.
  - Desktop composition redirection and legacy graphics mode are not supported on Windows 10.

Requirements:

- Microsoft .NET Framework 4.8 is installed automatically, if it (or a later version) is not already installed.
- Microsoft Visual C++ 2017 Runtime, 32-bit and 64-bit.

Remote PC Access uses this VDA, which you install on physical office PCs. This VDA supports Secure Boot for Citrix virtual desktops Remote PC Access on Windows 10.

Several multimedia acceleration features (such as HDX MediaStream Windows Media Redirection) require that Microsoft Media Foundation be installed on the machine on which you install the VDA. If the

machine does not have Media Foundation installed, the multimedia acceleration features will not be installed and will not work. Do not remove Media Foundation from the machine after installing the Citrix software. Otherwise, users will not be able to log on to the machine. On most supported Windows desktop OS editions, Media Foundation support is already installed and cannot be removed. However, N editions do not include certain media-related technologies. You can obtain that software from Microsoft or a third party.

More information:

- For Linux VDA information, see the [Linux Virtual Delivery Agent](#) product documentation.
- To use the server VDI feature, you can use the command line interface to install a single-session VDA on a supported Windows Server machine. See [Server VDI](#) for guidance.
- For information about installing a VDA on a Windows 7 machine, see [Earlier operating systems](#).
- See also VDAs in an Azure Virtual Desktop environment.

## VDA for multi-session OS

The following information applies to the latest VDA release.

Supported operating systems:

- [Windows Server 2022](#) (requires minimum VDA 2106)
- Windows Server 2019, Standard and Datacenter Editions
- Windows Server 2016, Standard and Datacenter Editions

The installer automatically deploys the following requirements:

- Microsoft .NET Framework 4.8 is installed automatically, if it (or a later version) is not already installed.
- Microsoft Visual C++ 2017 Runtime, 32-bit and 64-bit.

The installer automatically installs and enables Remote Desktop Services role services, if they are not already installed and enabled. This triggers a restart.

Several multimedia acceleration features (such as HDX MediaStream Windows Media Redirection) require that the Microsoft Media Foundation be installed on the machine on which you install the VDA. If the machine does not have Media Foundation installed, the multimedia acceleration features will not be installed and will not work. Do not remove Media Foundation from the machine after installing the Citrix software. Otherwise, users will not be able to log on to the machine. On most Windows Server versions, the Media Foundation feature is installed through the Server Manager. However, N editions do not include certain media-related technologies. You can obtain that software from Microsoft or a third party.

If Media Foundation is not present on the VDA, these multimedia features do not work:

- Flash Redirection
- Windows Media Redirection
- HTML5 Video Redirection
- HDX RealTime Webcam Redirection

More information:

- For Linux VDA information, see the [Linux Virtual Delivery Agent](#) articles.
- For information about installing a VDA on a Windows Server 2008 R2 machine, see [Earlier operating systems](#).
- See also VDAs in an Azure Virtual Desktop environment.

## Hosts / virtualization resources

The following host/virtualization resources (listed alphabetically) are supported. Where applicable, the *major.minor* versions are supported, including updates to those versions. [CTX131239](#) contains the most current hypervisor version information, plus links to known issues.

- **Amazon Web Services (AWS)**

- You can provision applications and desktops on supported Windows server operating systems.
- The Amazon Relational Database Service (RDS) is not supported.

For more information, see [Amazon Web Services virtualization environments](#).

- **Citrix Hypervisor (formerly XenServer)**

[CTX131239](#) contains current version information, plus links to known issues.

For more information, see [Citrix Hypervisor virtualization environments](#).

- **Google Cloud Platform**

For more information, see [Google Cloud Platform virtualization environments](#) and [Deployment Guide](#).

- **Microsoft Azure Resource Manager**

For more information, see [Microsoft Azure Resource Manager virtualization environments](#).

- **Microsoft System Center Virtual Machine Manager**

Includes any version of Hyper-V that can register with the supported System Center Virtual Machine Manager versions.

[CTX131239](#) contains current version information, plus links to known issues.

For more information, see [Microsoft System Center Virtual Machine Manager virtualization environments](#).



- **Nutanix Acropolis**

[CTX131239](#) contains current version information, plus links to known issues.

For more information, see [Nutanix virtualization environments](#).

- **Oracle Cloud Infrastructure (OCI) Classic**

For more information, see [Deploying Citrix Cloud XenApp and XenDesktop Service with Oracle Cloud Infrastructure Classic](#).

- **VMware Cloud on AWS**

[CTX131239](#) contains current version information, plus links to known issues.

For more information, see [VMware virtualization environments](#).

- **VMware vSphere(vCenter + ESXi)**

No support is provided for vSphere vCenter Linked Mode operation.

[CTX131239](#) contains current version information, plus links to known issues.

For more information, see [VMware virtualization environments](#).

## **Active Directory functional levels**

The following functional levels for the Active Directory forest and domain are supported:

- Windows Server 2016
- Windows Server 2012 R2
- Windows Server 2012
- Windows Server 2008 R2

For more information about Active Directory, see [Active Directory](#).

## **HDX technologies**

For specific HDX feature support and requirements, see [HDX](#).

## **Universal Print Server**

The Universal Print Server comprises client and server components. The UpsClient component is included in the VDA installation. You install the UpsServer component on each print server where shared printers reside that you want to provision with the Citrix Universal Print Driver in user sessions.

The UpsServer component is supported on:

- Windows Server 2019

- Windows Server 2016

Requirements:

- Microsoft Visual C++ 2017 Runtime, 32-bit and 64-bit
- Microsoft .NET Framework 4.8 (minimum)

For multi-session VDAs, user authentication during printing operations requires the Universal Print Server to be joined to the same domain as the VDA.

Standalone client and server component packages are also available for download.

For more information, see [Provision printers](#).

## Service connectivity

See [System and Connectivity Requirements](#) for Internet connection information. That information includes requirements that are common to most Citrix Cloud services, plus [requirements specific to Citrix Virtual Apps and Desktops](#).

## Other

- The Microsoft Group Policy Management Console (GPMC) is required if you store Citrix policy information in Active Directory rather than the site configuration database. The machine on which you install `CitrixGroupPolicyManagement_x64.msi` must have Visual Studio 2015 runtime installed. For more information, see the Microsoft documentation.
- For product components and features that you can install on Windows Servers, Server Core and Nano Server installations are not supported, unless noted.
- For details about resource limits in a deployment, see [Limits](#).
- For supported StoreFront versions, see the [StoreFront system requirements](#).
- For globalization information, see [CTX119253](#).
- For information about ports that the service uses, see [Communications Ports Used by Citrix Technologies](#).
- For information about requirements when using the Quick Deploy management interface, see [Requirements](#).

## Limits

July 28, 2021

The values in this article indicate the limits of a single Citrix Virtual Apps and Desktops service instance. When the number of concurrent users exceeds 100,000, Citrix can scale and combine multiple Citrix Virtual Apps and Desktops service instances, to deliver a unified experience at any scale.

The information in this article is dynamic. Check back frequently for updates. If you have current requirements that the published limits do not address, contact your Citrix representative for assistance.

## Configuration limits

If policies exceed the limit, Citrix recommends using the [Workspace Environment Management service](#) or [Active Directory Group Policy Objects \(GPOs\)](#).

---

Resource	Limit
Active Directory domains	50
Application folders	1,000
Application Groups	250
Applications	2,000
Catalogs	1,000
Delivery Groups	1,000
Host connections	200
Resource locations	50
Management console (Studio) policies	100
Tags	500
VDAs	60,000

---

## Resource location limits

The following table lists the limits for each resource location.

If your requirements exceed these limits, Citrix recommends using additional resource locations.

---

Resource	Limit
Single-session VDAs	5,000
Multi-session VDAs	500
Active Directory domains	1

---

Resource	Limit
Host connections	20

### Machine Creation Services (MCS) limits

The MCS limits in the following table are the recommended maximums for a single public provider subscription.

You are likely to reach quota limits from your public cloud vendor at lower levels. In such cases, contact the vendor to raise your subscription quota. For larger-scale deployments, Citrix recommends a hub-and-spoke model, where VDAs are distributed across multiple subscriptions and host connections.

For more information, see the following reference architectures:

- [Citrix Virtual Apps and Desktops Service on Azure](#)
- [Reference Architecture for Citrix Virtual Apps and Desktops on AWS](#)

Resource	Limit
VDAs per Amazon Web Services account per region	1,500
VDAs per Google Cloud Platform project	1,000
VDAs per Microsoft Azure subscription	2,500

### Usage limits

For information about administrator roles and the differences between them, see:

- [Manage \(Studio\) administrators](#)
- [Monitor \(Director\) administrators](#)

Resource	Limit
Concurrent Monitor (Director) full administrators	20
Concurrent Monitor (Director) help desk administrators	130
Concurrent Monitor (Director) session administrators	50

Resource	Limit
Concurrent Manage (Studio) cloud administrators	100
Concurrent Manage (Studio) help desk administrators	60
Concurrent end users	100,000
Resources published to a single user	250
Session launches per minute	3,000


## Technical security overview

August 31, 2021

### Security overview

This document applies to Citrix Virtual Apps and Desktops services hosted in Citrix Cloud. This includes Citrix Virtual Apps Essentials and Citrix Virtual Desktops Essentials.

Citrix Cloud manages the operation of the control plane for Citrix Virtual Apps and Desktops environments. This includes the Delivery Controllers, management consoles, SQL database, license server, and optionally StoreFront and Citrix Gateway (formerly NetScaler Gateway). The Virtual Delivery Agents (VDAs) hosting the apps and desktops remain under the customer's control in the data center of their choice, either cloud or on-premises. These components are connected to the cloud service using an agent called the Citrix Cloud Connector. If customers elect to use Citrix Workspace, they can also choose to use the Citrix Gateway Service instead of running Citrix Gateway within their data center. The following diagram illustrates the service and its security boundaries.



### Citrix cloud-based compliance

As of January 2021, the use of Citrix Managed Azure Capacity with various Citrix Virtual Apps and Desktops service editions and Workspace Premium Plus has not been evaluated for Citrix SOC 2 (Type 1 or 2), ISO 27001, HIPAA, or other cloud compliance requirements. Visit the [Citrix Trust Center](#) for more information regarding Citrix Cloud Certifications, and check back frequently for updates.

### Data flow

The VDAs are not hosted by the service, so the customer’s application data and images required for provisioning are always hosted in the customer setup. The control plane has access to metadata, such as user names, machine names, and application shortcuts, restricting access to the customer’s Intellectual Property from the control plane.

Data flowing between the cloud and customer premises uses secure TLS connections over port 443.

### Data isolation

The Citrix Virtual Apps and Desktops service stores only the metadata needed for the brokering and monitoring of the customer’s applications and desktops. Sensitive information, including images, user profiles, and other application data remains on the customer premises or in their subscription with a public cloud vendor.

### Service editions

The capabilities of the Citrix Virtual Apps and Desktops service vary by edition. For example, Citrix Virtual Apps Essentials supports only Citrix Gateway service and Citrix Workspace. Consult that product

documentation to learn more about supported features.

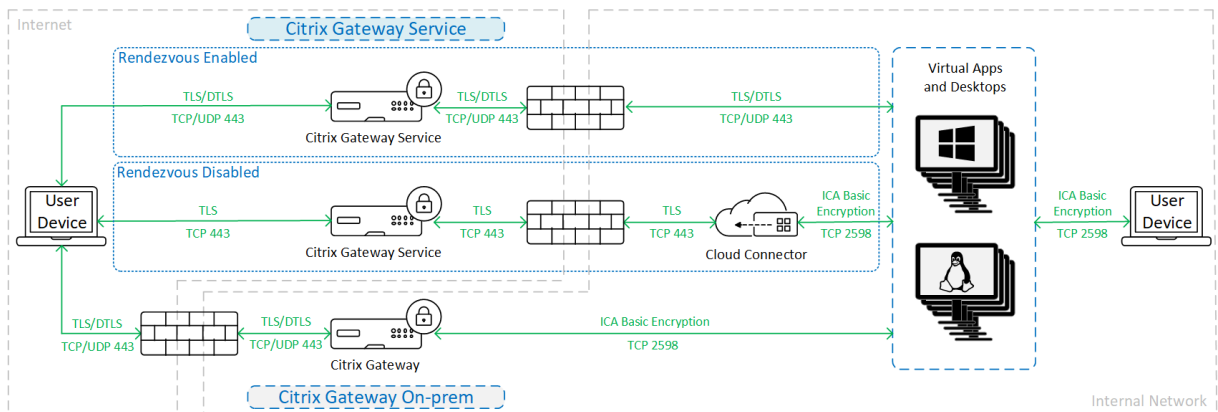
### ICA Security

The Citrix Virtual Apps and Desktops service provides several options for securing ICA traffic in transit. The following are the options available:

- **Basic encryption:** The default setting.
- **SecureICA:** Allows encrypting session data using RC5 (128-bit) encryption.
- **VDA TLS/DTLS:** Allows using network-level encryption using TLS/DTLS.
- **Rendezvous protocol:** Available only when using the Citrix Gateway Service. When using the Rendezvous protocol, ICA sessions are encrypted end-to-end using TLS/DTLS.

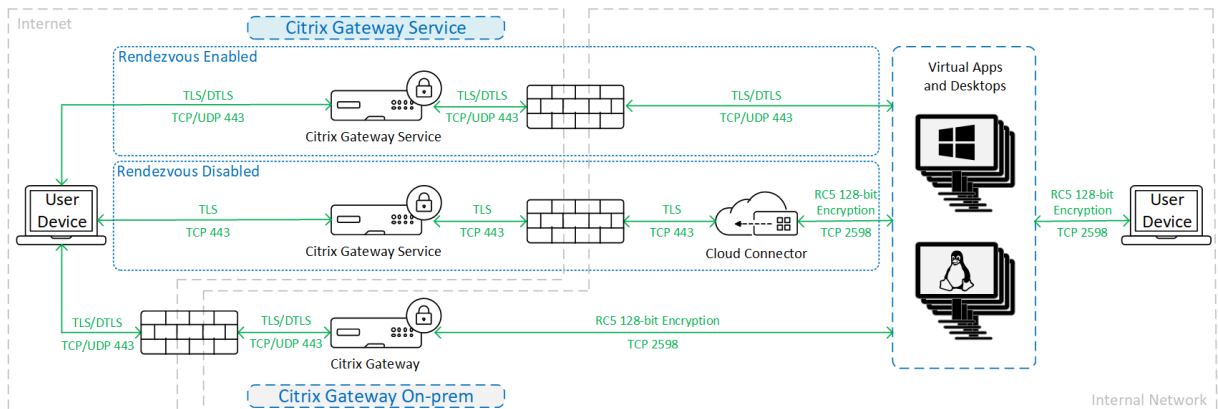
#### Basic encryption

When using basic encryption, traffic is encrypted as shown in the following graphic.



#### SecureICA

When using SecureICA, traffic is encrypted as shown in the following graphic.

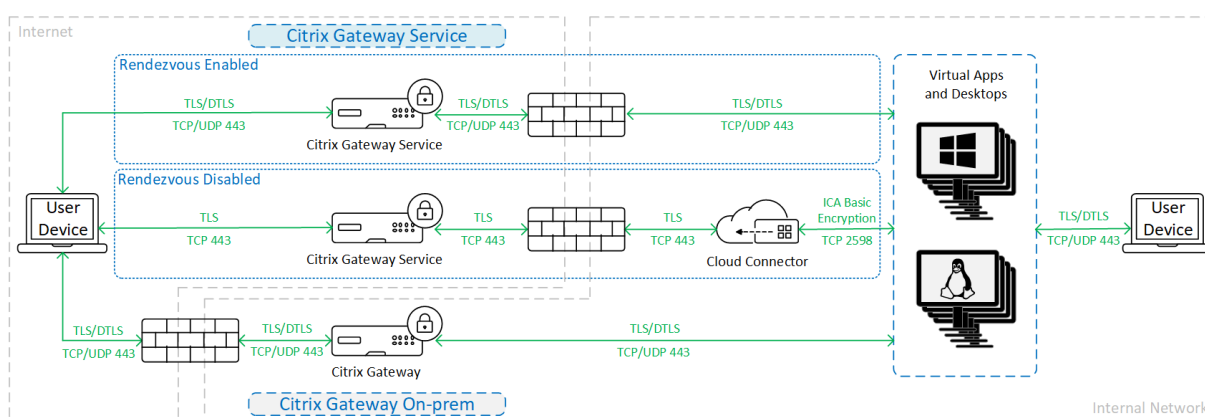


**Note:**

SecureICA is not supported when using Workspace app for HTML5 or Chrome OS.

**VDA TLS/DTLS**

When using VDA TLS/DTLS encryption, traffic is encrypted as shown in the following graphic.



**Note:**

When using the Gateway Service without Rendezvous, the traffic between the VDA and the Cloud Connector is not TLS encrypted, because the Cloud Connector does not support connecting to the VDA with network-level encryption.

**More resources**

For more information about the ICA security options and how to configure them, see:

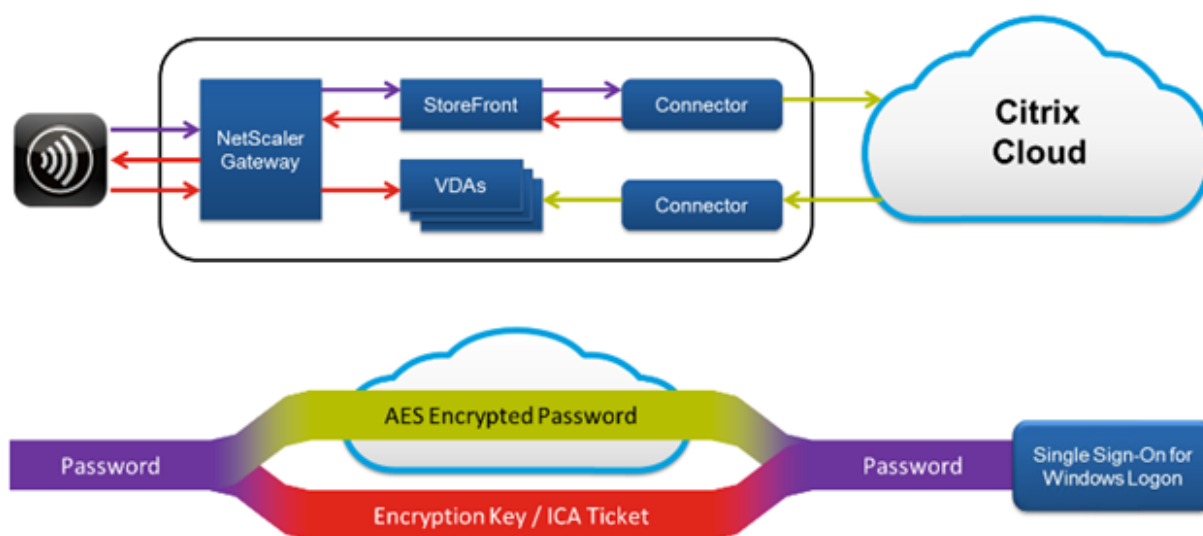
- SecureICA: [Security policy settings](#)
- VDA TLS/DTLS: [Transport Layer Security](#)
- Rendezvous protocol: [Rendezvous protocol](#)

**Credential handling**

The service handles four types of credentials:

- User Credentials: When using a customer-managed StoreFront, the Cloud Connector encrypts user credentials using AES-256 encryption and a random one-time key generated for each launch. The key is never passed into the cloud, and returned only to Citrix Workspace app. The Citrix Workspace app then passes this key to the VDA to decrypt the user password during session launch for a single sign-on experience. The flow is shown in the following figure.





- Administrator Credentials: Administrators authenticate against Citrix Cloud. This generates a one-time signed JSON Web Token (JWT) which gives the administrator access to the Citrix Virtual Apps and Desktops service.
- Hypervisor Passwords: On-premises hypervisors that require a password for authentication have a password generated by the administrator and directly stored encrypted in the SQL database in the cloud. Peer keys are managed by Citrix to ensure that hypervisor credentials are only available to authenticated processes.
- Active Directory (AD) Credentials: Machine Creation Services uses the Cloud Connector for creating machine accounts in a customer's AD. Because the machine account of the Cloud Connector has only read access to AD, the administrator is prompted for credentials for each machine creation or deletion operation. These credentials are stored only in memory, and are held only for a single provisioning event.

## Deployment considerations

Citrix recommends that users consult the published best practices documentation for deploying Citrix Gateway applications and VDAs within their environments.

## Citrix Cloud Connector network access requirements

The Citrix Cloud Connectors require only port 443 outbound traffic to the internet, and can be hosted behind an HTTP proxy.

- The communication used in Citrix Cloud for HTTPS is TLS. (See Deprecation of TLS versions.)
- Within the internal network, the Cloud Connector needs access to the following for the Citrix Virtual Apps and Desktops service:

- VDAs: Port 80, both inbound and outbound. plus 1494 and 2598 inbound if using Citrix Gateway service
- StoreFront servers: Port 80 inbound.
- Citrix Gateways, if configured as a STA: Port 80 inbound.
- Active Directory domain controllers
- Hypervisors: Outbound only. See [Communications Ports Used by Citrix Technologies](#) for specific ports.

Traffic between the VDAs and Cloud Connectors is encrypted using Kerberos message-level security.

### **Customer-managed StoreFront**

A customer-managed StoreFront offers greater security configuration options and flexibility for deployment architecture, including the ability to maintain user credentials on-premises. The StoreFront can be hosted behind the Citrix Gateway to provide secure remote access, enforce multifactor authentication, and add other security features.

### **Citrix Gateway service**

Using the Citrix Gateway service avoids the need to deploy Citrix Gateway within customer data centers.

For details, see [Citrix Gateway service](#).

All TLS connections between the Cloud Connector and Citrix Cloud are initiated from the Cloud Connector to the Citrix Cloud. No in-bound firewall port mapping is required.

### **XML trust**

The XML trust setting applies to deployments that use:

- An on-premises StoreFront.
- A subscriber (user) authentication technology that does not require passwords. Examples of such technologies are domain pass-through, smart cards, SAML, and Veridium solutions.

Enabling the XML trust setting allows users to successfully authenticate and then start applications. The Cloud Connector trusts the credentials sent from StoreFront. Enable this setting only when you have secured communications between your Citrix Cloud Connectors and StoreFront (using firewalls, IPsec, or other security recommendations).

This setting is disabled by default.

Use the Citrix Virtual Apps and Desktops Remote PowerShell SDK to manage the XML trust setting.

- To check the XML trust setting's current value, run `Get-BrokerSite` and inspect the value of `TrustRequestsSentToTheXMLServicePort`.
- To enable XML trust, run `Set-BrokerSite -TrustRequestsSentToTheXMLServicePort $true`
- To disable XML trust, run `Set-BrokerSite -TrustRequestsSentToTheXMLServicePort $false`

### Enforce HTTPS or HTTP traffic

To enforce either HTTPS or HTTP traffic through the XML Service, configure one of the following registry value sets on each of your Cloud Connectors.

After you configure the settings, restart the Broker Service and the Remote Broker Provider Service on each Cloud Connector.

In `HKLM\Software\Citrix\DesktopServer\`:

- To enforce HTTPS (ignore HTTP) traffic: Set `XmlServicesEnableSsl` to 1, and `XmlServicesEnableNonSsl` to 0.
- To enforce HTTP (ignore HTTPS) traffic: Set `XmlServicesEnableNonSsl` to 1, and `XmlServicesEnableSsl` to 0.

### Deprecation of TLS versions

To improve the security of the Citrix Virtual Apps and Desktops service, Citrix began blocking any communication over Transport Layer Security (TLS) 1.0 and 1.1 as of March 15, 2019.

All connections to Citrix Cloud services from Citrix Cloud Connectors require TLS 1.2.

To ensure successful connection to Citrix Workspace from user endpoint devices, the installed Citrix Receiver version must be equal to or newer than the version listed in the following table.

Receiver	Version
Windows	4.2.1000
Mac	12.0
Linux	13.2
Android	3.7
iOS	7.0
Chrome/HTML5	Latest (browser must support TLS 1.2)

To upgrade to the latest Citrix Receiver version, go to <https://www.citrix.com/products/receiver/>.

Alternatively, upgrade to the [Citrix Workspace app](#), which uses TLS 1.2. To download the Citrix Workspace app, go to <https://www.citrix.com/downloads/workspace-app/>.

If you must continue using TLS 1.0 or 1.1 (for example, with a thin client based on an earlier Receiver for Linux version), install a StoreFront in your resource location and have all the Citrix Receivers point to it.

## More information

The following resources contain security information:

- [Technical security overview for Citrix Managed Azure](#).
- [Citrix security site](#).
- [Security and Compliance Information](#): The security and compliance center contains security bulletins that can help you stay informed. The center also has documentation about standards and certifications that are important in maintaining a secure and compliant IT environment.
- [Secure Deployment Guide for the Citrix Cloud Platform](#): This guide provides an overview of security best practices when using Citrix Cloud and describes the information Citrix Cloud collects and manages. This guide also contains links to comprehensive information about the Citrix Cloud Connector.
- [System and Connectivity Requirements](#).
- [Security considerations and best practices](#).
- [Smart cards](#).
- [Transport Layer Security \(TLS\)](#).

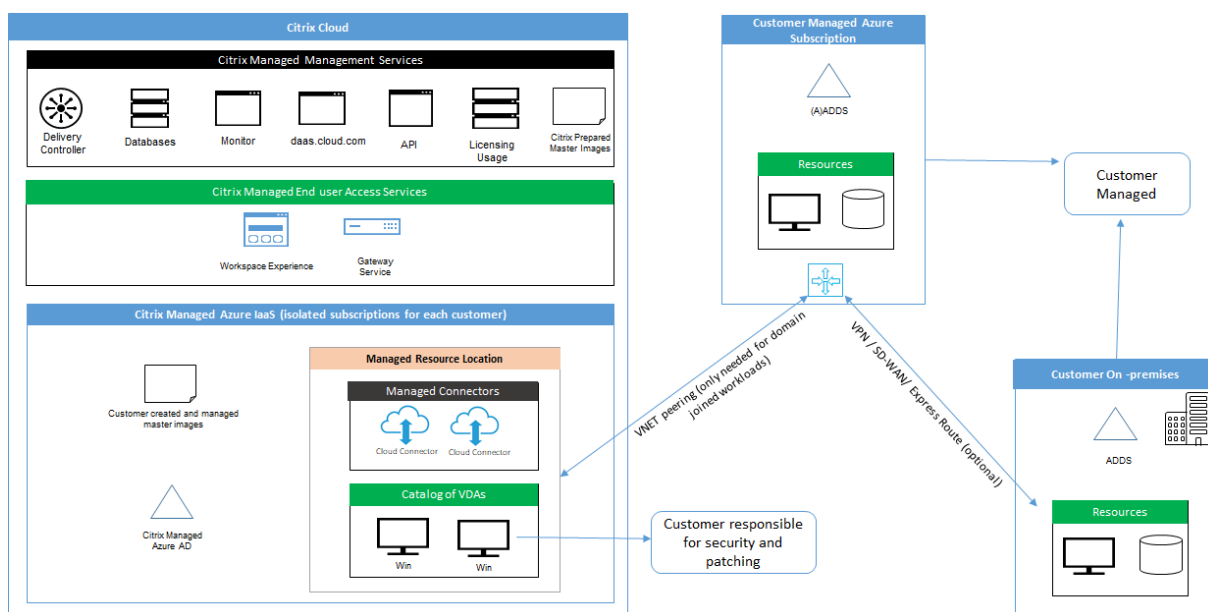
### Note:

This document is intended to provide the reader with an introduction to and overview of the security functionality of Citrix Cloud; and to define the division of responsibility between Citrix and customers with regard to securing the Citrix Cloud deployment. It is not intended to serve as a configuration and administration guidance manual for Citrix Cloud or any of its components or services.

## Technical security overview for Citrix Managed Azure

July 21, 2021

The following diagram shows the components in a Citrix Virtual Apps and Desktops service deployment that uses Citrix Managed Azure. This example uses a VNet peering connection.



With Citrix Managed Azure, the customer’s Virtual Delivery Agents (VDAs) that deliver desktops and apps, plus Citrix Cloud Connectors, are deployed into an Azure subscription and tenant that Citrix manages.

### Citrix cloud-based compliance

As of January 2021, the use of Citrix Managed Azure Capacity with various Citrix Virtual Apps and Desktops service editions and Workspace Premium Plus has not been evaluated for Citrix SOC 2 (Type 1 or 2), ISO 27001, HIPAA, or other cloud compliance requirements. Visit the [Citrix Trust Center](#) for more information regarding Citrix Cloud Certifications, and check back frequently for updates.

### Citrix responsibility

#### Citrix Cloud Connectors for non-domain-joined catalogs

When using a Citrix Managed Azure subscription, the Citrix Virtual Apps and Desktops service deploys at least two Cloud Connectors in each resource location. Some catalogs may share a resource location if they are in the same region as other catalogs for the same customer.

Citrix is responsible for the following security operations on non-domain-joined catalog Cloud Connectors:

- Applying operating system updates and security patches
- Installing and maintaining antivirus software
- Applying Cloud Connector software updates

Customers do not have access to the Cloud Connectors. Therefore, Citrix is wholly responsible for the performance of the non-domain-joined catalog Cloud Connectors.

### **Azure subscription and Azure Active Directory**

Citrix is responsible for the security of the Azure subscription and Azure Active Directory (AAD) that are created for the customer. Citrix ensures tenant isolation, so each customer has their own Azure subscription and AAD, and cross-talk between different tenants is prevented. Citrix also restricts access to the AAD to the Citrix Virtual Apps and Desktops service and Citrix operations personnel only. Access by Citrix to each customer's Azure subscription is audited.

Customers employing non-domain-joined catalogs can use the Citrix-managed AAD as a means of authentication for Citrix Workspace. For these customers, Citrix creates limited privilege user accounts in the Citrix-managed AAD. However, neither customers' users nor administrators can execute any actions on the Citrix-managed AAD. If these customers elect to use their own AAD instead, they are wholly responsible for its security.

### **Virtual networks and infrastructure**

Within the customer's Citrix Managed Azure subscription, Citrix creates virtual networks for isolating resource locations. Within those networks, Citrix creates virtual machines for the VDAs, Cloud Connectors, and image builder machines, in addition to storage accounts, Key Vaults, and other Azure resources. Citrix, in partnership with Microsoft, is responsible for the security of the virtual networks, including virtual network firewalls.

Citrix ensures the default Azure firewall policy (network security groups) is configured to limit access to network interfaces in VNet peering and SD-WAN connections. Generally, this controls incoming traffic to VDAs and Cloud Connectors. For details, see:

- Firewall policy for Azure VNet peering connections
- Firewall policy for SD-WAN connections

Customers cannot change this default firewall policy, but may deploy additional firewall rules on Citrix-created VDA machines; for example, to partially restrict outgoing traffic. Customers that install virtual private network clients, or other software capable of bypassing firewall rules, on Citrix-created VDA machines are responsible for any security risks that might result.

When using the image builder in the Citrix Virtual Apps and Desktops service to create and customize a new machine image, ports 3389-3390 are opened temporarily in the Citrix-managed VNet, so that the customer can RDP to the machine containing the new machine image, to customize it.

### **Citrix responsibility when using Azure VNet peering connections**

For VDAs in the Citrix Virtual Apps and Desktops service to contact on-premises domain controllers, file shares, or other intranet resources, the service provides a VNet peering workflow as a connectivity option. The customer's Citrix-managed virtual network is peered with a customer-managed Azure virtual network. The customer-managed virtual network may enable connectivity with the customer's on-premises resources using the cloud-to-on-premises connectivity solution of the customer's choice, such as Azure ExpressRoute or IPsec tunnels.

Citrix responsibility for VNet peering is limited to supporting the workflow and related Azure resource configuration for establishing peering relationship between Citrix and customer-managed VNets.

### **Firewall policy for Azure VNet peering connections**

Citrix opens or closes the following ports for inbound and outbound traffic that uses a VNet peering connection.

### **Citrix-managed VNet with non-domain-joined machines**

- Inbound rules
  - Allow ports 80, 443, 1494, and 2598 inbound from VDAs to Cloud Connectors, and from Cloud Connectors to VDAs.
  - Allow ports 49152-65535 inbound to the VDAs from an IP range used by the Monitor shadowing feature. See [Communications Ports Used by Citrix Technologies](#).
  - Deny all other inbound. This includes intra-VNet traffic from VDA to VDA, and VDA to Cloud Connector.
- Outbound rules
  - Allow all traffic outbound.

### **Citrix-managed VNet with domain-joined machines**

- Inbound rules:
  - Allow ports 80, 443, 1494, and 2598 inbound from the VDAs to Cloud Connectors, and from Cloud Connectors to VDAs.
  - Allow ports 49152-65535 inbound to the VDAs from an IP range used by the Monitor shadowing feature. See [Communications Ports Used by Citrix Technologies](#).
  - Deny all other inbound. This includes intra-VNet traffic from VDA to VDA, and VDA to Cloud Connector.
- Outbound rules
  - Allow all traffic outbound.

### **Customer-managed VNet with domain-joined machines**

- It is up to the customer to configure their VNet correctly. This includes opening the following ports for domain joining.
- Inbound rules:
  - Allow inbound on 443, 1494, 2598 from their client IPs for internal launches.
  - Allow inbound on 53, 88, 123, 135-139, 389, 445, 636 from Citrix VNet (IP range specified by customer).
  - Allow inbound on ports opened with a proxy configuration.
  - Other rules created by customer.
- Outbound rules:
  - Allow outbound on 443, 1494, 2598 to the Citrix VNet (IP range specified by customer) for internal launches.
  - Other rules created by customer.

### **Citrix responsibility when using SD-WAN connectivity**

Citrix supports a fully automated way of deploying virtual Citrix SD-WAN instances to enable connectivity between the Citrix Virtual Apps and Desktops service and on-premises resources. Citrix SD-WAN connectivity has several advantages compared to VNet peering, including:

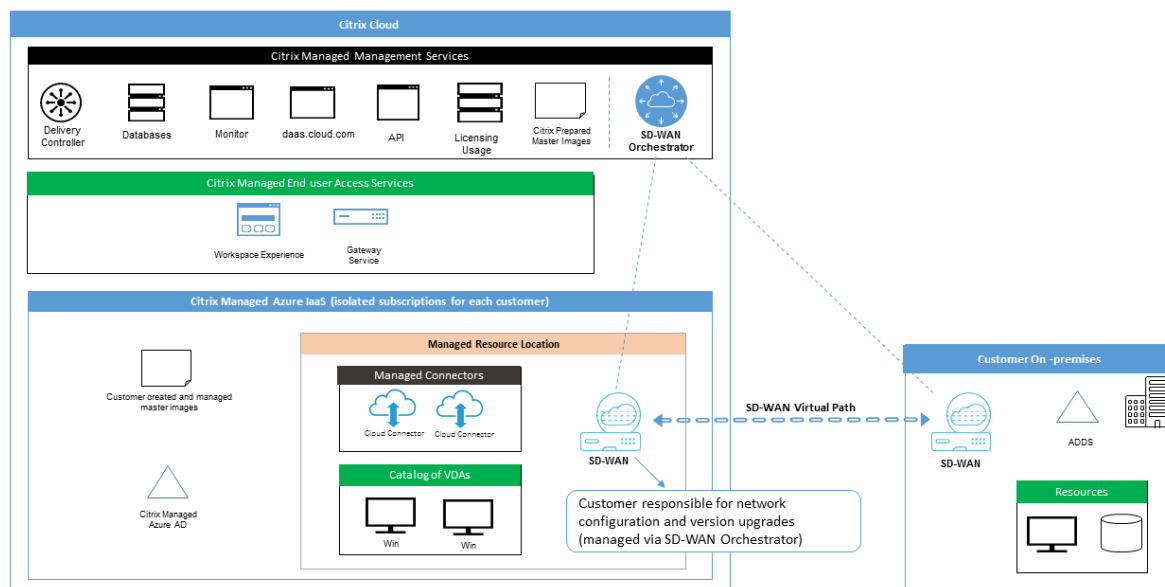
High reliability and security of VDA-to-datacenter and VDA-to-branch (ICA) connections.

- Best end-user experience for office workers, with advanced QoS capabilities and VoIP optimizations.
- Built-in ability to inspect, prioritize, and report on Citrix HDX network traffic and other application usage.

Citrix requires customers who want to take advantage of SD-WAN connectivity for the Citrix Virtual Apps and Desktops service to use SD-WAN Orchestrator for managing their Citrix SD-WAN networks.

The following diagram shows the added components in a Citrix Virtual Apps and Desktops service deployment using a Citrix Managed Azure subscription and SD-WAN connectivity.





The Citrix SD-WAN deployment for Citrix Virtual Apps and Desktops service is similar to the standard Azure deployment configuration for Citrix SD-WAN. For more information, see [Deploy Citrix SD-WAN Standard Edition Instance on Azure](#). In a high availability configuration, an active/standby pair of SD-WAN instances with Azure load balancers is deployed as a gateway between the subnet containing VDAs and Cloud Connectors, and the Internet. In a non-HA configuration, only a single SD-WAN instance is deployed as a gateway. Network interfaces of the virtual SD-WAN appliances are assigned addresses from a separate small address range split into two subnets.

When configuring SD-WAN connectivity, Citrix makes a few changes to the networking configuration of managed desktops described above. In particular, all outgoing traffic from the VNet, including traffic to Internet destinations, is routed through the cloud SD-WAN instance. The SD-WAN instance is also configured to be the DNS server for the Citrix-managed VNet.

Management access to the virtual SD-WAN instances requires an admin login and password. Each instance of SD-WAN is assigned a unique, random secure password that can be used by SD-WAN administrators for remote login and troubleshooting through the SD-WAN Orchestrator UI, the virtual appliance management UI and CLI.

Just like other tenant-specific resources, virtual SD-WAN instances deployed in a specific customer VNet are fully isolated from all other VNets.

When the customer enables Citrix SD-WAN connectivity, Citrix automates the initial deployment of virtual SD-WAN instances used with the Citrix Virtual Apps and Desktops service, maintains underlying Azure resources (virtual machines, load balancers, and so on), provides secure and efficient out-of-the-box defaults for the initial configuration of virtual SD-WAN instances, and enables ongoing maintenance and troubleshooting through SD-WAN Orchestrator. Citrix also takes reasonable measures to

perform automatic validation of SD-WAN network configuration, check for known security risks, and display corresponding alerts through SD-WAN Orchestrator.

### **Firewall policy for SD-WAN connections**

Citrix uses Azure firewall policies (network security groups) and public IP address assignment to limit access to network interfaces of virtual SD-WAN appliances:

- Only WAN and management interfaces are assigned public IP addresses and allow outbound connectivity to the Internet.
- LAN interfaces, acting as gateways for the Citrix-managed VNet, are only allowed to exchange network traffic with virtual machines on the same VNet.
- WAN interfaces limit inbound traffic to UDP port 4980 (used by Citrix SD-WAN for virtual path connectivity), and deny outbound traffic to the VNet.
- Management ports allow inbound traffic to ports 443 (HTTPS) and 22 (SSH).
- HA interfaces are only allowed to exchange control traffic with each other.

### **Access to infrastructure**

Citrix may access the customer's Citrix-managed infrastructure (Cloud Connectors) to perform certain administrative tasks such as collecting logs (including Windows Event Viewer) and restarting services without notifying the customer. Citrix is responsible for executing these tasks safely and securely, and with minimal impact to the customer. Citrix is also responsible for ensuring any log files are retrieved, transported, and handled safely and securely. Customer VDAs cannot be accessed this way.

### **Backups for non-domain-joined catalogs**

Citrix is not responsible for performing backups of non-domain-joined catalogs.

### **Backups for machine images**

Citrix is responsible for backing up any machine images uploaded to Citrix Virtual Apps and Desktops service, including images created with the image builder. Citrix uses locally redundant storage for these images.

### **Bastions for non-domain-joined catalogs**

Citrix operations personnel have the ability to create a bastion, if necessary, to access the customer's Citrix-managed Azure subscription for diagnosing and repairing customer issues, potentially before the customer is aware of a problem. Citrix does not require the customer's consent to create a bastion. When Citrix creates the bastion, Citrix creates a strong randomly generated password for the

bastion and restricts RDP access to Citrix NAT IP addresses. When the bastion is no longer needed, Citrix disposes of it and the password is no longer valid. The bastion (and its accompanying RDP access rules) are disposed of when the operation completes. Citrix can access only the customer's non-domain-joined Cloud Connectors with the bastion. Citrix does not have the password to log in to non-domain-joined VDAs or domain-joined Cloud Connectors and VDAs.

### **Firewall policy when using troubleshooting tools**

When a customer requests creation of a bastion machine for troubleshooting, the following security group modifications are made to the Citrix-managed VNet:


- Temporarily allow 3389 inbound from the customer-specified IP range to the bastion.
- Temporarily allow 3389 inbound from the bastion IP address to any address in the VNet (VDAs and Cloud Connectors).
- Continue to block RDP access between the Cloud Connectors, VDAs, and other VDAs.

When a customer enables RDP access for troubleshooting, the following security group modifications are made to the Citrix-managed VNet:

- Temporarily allow 3389 inbound from the customer-specified IP range to any address in the VNet (VDAs and Cloud Connectors).
- Continue to block RDP access between the Cloud Connectors, VDAs, and other VDAs.

### **Customer-managed subscriptions**

For customer-managed subscriptions, Citrix adheres to the above responsibilities during deployment of the Azure resources. After deployment, everything above falls to the customer's responsibility, because the customer is the owner of the Azure subscription.



## Customer responsibility

### VDAs and machine images

The customer is responsible for all aspects of the software installed on VDA machines, including:

- Operating system updates and security patches
- Antivirus and antimalware
- VDA software updates and security patches
- Additional software firewall rules (especially outbound traffic)
- Follow Citrix [security considerations and best practices](#)

Citrix provides a prepared image that is intended as a starting point. Customers can use this image for proof-of-concept or demonstration purposes or as a base for building their own machine image. Citrix does not guarantee the security of this prepared image. Citrix will make an attempt to keep the operating system and VDA software on the prepared image up to date, and will enable Windows Defender on these images.

### **Customer responsibility when using VNet peering**

The customer must open all ports specified in Customer-managed VNet with domain-joined machines.

When VNet peering is configured, the customer is responsible for the security of their own virtual network and its connectivity to their on-premises resources. The customer is also responsible for security of the incoming traffic from the Citrix-managed peered virtual network. Citrix does not take any action to block traffic from the Citrix-managed virtual network to the customer's on-premises resources.

Customers have the following options for restricting incoming traffic:

- Give the Citrix-managed virtual network an IP block which is not in use elsewhere in the customer's on-premises network or the customer-managed connected virtual network. This is required for VNet peering.
- Add Azure network security groups and firewalls in the customer's virtual network and on-premises network to block or restrict traffic from the Citrix-managed IP block.
- Deploy measures such as intrusion prevention systems, software firewalls, and behavioral analytics engines in the customer's virtual network and on-premises network, targeting the Citrix-managed IP block.

### **Customer responsibility when using SD-WAN connectivity**

When SD-WAN connectivity is configured, customers have full flexibility to configure virtual SD-WAN instances used with Citrix Virtual Apps and Desktops service according to their networking requirements, with the exception of a few elements required to ensure correct operation of SD-WAN in the Citrix-managed VNet. Customer responsibilities include:

- Design and configuration of routing and firewall rules, including rules for DNS and Internet traffic breakout.
- Maintenance of the SD-WAN network configuration.
- Monitoring of the operational status of the network.
- Timely deployment of Citrix SD-WAN software updates or security fixes. Since all instances of Citrix SD-WAN on a customer network must run the same version of SD-WAN software, deployments of updated software versions to Citrix Virtual Apps and Desktops service SD-WAN instances need to be managed by customers according to their network maintenance schedules and constraints.

Incorrect configuration of SD-WAN routing and firewall rules, or mismanagement of SD-WAN management passwords, may result in security risks to both virtual resources in Citrix Virtual Apps and Desktops service, and on-premises resources reachable through Citrix SD-WAN virtual paths. Another possible security risk stems from not updating Citrix SD-WAN software to the latest available patch release. While SD-WAN Orchestrator and other Citrix Cloud services provide the means to address such

risks, customers are ultimately responsible for ensuring that virtual SD-WAN instances are configured appropriately.

### **Proxy**

The customer may choose whether to use a proxy for outbound traffic from the VDA. If a proxy is used, the customer is responsible for:

- Configuring the proxy settings on the VDA machine image or, if the VDA is joined to a domain, using Active Directory Group Policy.
- Maintenance and security of the proxy.

Proxies are not allowed for use with Citrix Cloud Connectors or other Citrix-managed infrastructure.

### **Catalog resiliency**

Citrix provides three types of catalogs with differing levels of resiliency:

- **Static:** Each user is assigned to a single VDA. This catalog type provides no high availability. If a user's VDA goes down, they will have to be placed on a new one to recover. Azure provides a 99.5% SLA for single-instance VMs. The customer can still back up the user profile, but any customizations made to the VDA (such as installing programs or configuring Windows) will be lost.
- **Random:** Each user is assigned randomly to a server VDA at launch time. This catalog type provides high availability via redundancy. If a VDA goes down, no information is lost because the user's profile resides elsewhere.
- **Windows 10 multisession:** This catalog type operates in the same manner as the random type but uses Windows 10 workstation VDAs instead of server VDAs.

### **Backups for domain-joined catalogs**

If the customer uses domain-joined catalogs with a VNet peering, the customer is responsible for backing up their user profiles. Citrix recommends that customers configure on-premises file shares and set policies on their Active Directory or VDAs to pull user profiles from these file shares. The customer is responsible for the backup and availability of these file shares.

### **Disaster recovery**

In the event of Azure data loss, Citrix will recover as many resources in the Citrix-managed Azure subscription as possible. Citrix will attempt to recover the Cloud Connectors and VDAs. If Citrix is unsuccessful recovering these items, customers are responsible for creating a new catalog. Citrix assumes

that machine images are backed up and that customers have backed up their user profiles, allowing the catalog to be rebuilt.

In the event of the loss of an entire Azure region, the customer is responsible for rebuilding their customer-managed virtual network in a new region and creating a new VNet peering or a new SD-WAN instance within Citrix Virtual Apps and Desktops service.

## **Citrix and customer shared responsibilities**

### **Citrix Cloud Connector for domain-joined catalogs**

The Citrix Virtual Apps and Desktops service deploys at least two Cloud Connectors in each resource location. Some catalogs may share a resource location if they are in the same region, VNet peering, and domain as other catalogs for the same customer. Citrix configures the customer's domain-joined Cloud Connectors for the following default security settings on the image:

- Operating system updates and security patches
- Antivirus software
- Cloud Connector software updates

Customers do not normally have access to the Cloud Connectors. However, they may acquire access by using catalog troubleshooting steps and logging in with domain credentials. The customer is responsible for any changes they make when logging in through the bastion.

Customers also have control over the domain-joined Cloud Connectors through Active Directory Group Policy. The customer is responsible for ensuring that the group policies that apply to the Cloud Connector are safe and sensible. For example, if the customer chooses to disable operating system updates using Group Policy, the customer is responsible for performing operating system updates on the Cloud Connectors. The customer can also choose to use Group Policy to enforce stricter security than the Cloud Connector defaults, such as by installing a different antivirus software. In general, Citrix recommends that customers put Cloud Connectors into their own Active Directory organizational unit with no policies, as this will ensure that the defaults Citrix uses can be applied without issue.

## **Troubleshooting**

In the event the customer experiences problems with the catalog in the Citrix Virtual Apps and Desktops service, there are two options for troubleshooting: using bastions and enabling RDP access. Both options introduce security risk to the customer. The customer must understand and consent to undertaking this risk prior to using these options.

Citrix is responsible for opening and closing the necessary ports to carry out troubleshooting operations, and restricting which machines can be accessed during these operations.

With either bastions or RDP access, the active user performing the operation is responsible for the security of the machines that are being accessed. If the customer accesses the VDA or Cloud Connector through RDP and accidentally contracts a virus, the customer is responsible. If Citrix Support personnel access these machines, it is the responsibility of those personnel to perform operations safely. Responsibility for any vulnerabilities exposed by any person accessing the bastion or other machines in the deployment (for example, customer responsibility to add IP ranges to allow list, Citrix responsibility to implement IP ranges correctly) is covered elsewhere in this document.

In both scenarios, Citrix is responsible for correctly creating firewall exceptions to allow RDP traffic. Citrix is also responsible for revoking these exceptions after the customer disposes of the bastion or ends RDP access through the Citrix Virtual Apps and Desktops service.

### **Bastions**

Citrix may create bastions in the customer's Citrix-managed virtual network within the customer's Citrix-managed subscription to diagnose and repair issues, either proactively (without customer notification) or in response to a customer-raised issue. The bastion is a machine that the customer can access through RDP and then use to access the VDAs and (for domain-joined catalogs) Cloud Connectors through RDP to gather logs, restart services, or perform other administrative tasks. By default, creating a bastion opens an external firewall rule to allow RDP traffic from a customer-specified range of IP addresses to the bastion machine. It also opens an internal firewall rule to allow access to the Cloud Connectors and VDAs through RDP. Opening these rules poses a large security risk.

The customer is responsible for providing a strong password used for the local Windows account. The customer is also responsible for providing an external IP address range that allows RDP access to the bastion. If the customer elects not to provide an IP range (allowing anyone to attempt RDP access), the customer is responsible for any access attempted by malicious IP addresses.

The customer is also responsible for deleting the bastion after troubleshooting is complete. The bastion host exposes additional attack surface, so Citrix automatically shuts down the machine eight (8) hours after it is powered on. However, Citrix never automatically deletes a bastion. If the customer chooses to use the bastion for an extended period of time, they are responsible for patching and updating it. Citrix recommends that a bastion be used only for several days before deleting it. If the customer wants an up-to-date bastion, they can delete their current one and then create a new bastion, which will provision a fresh machine with the latest security patches.

### **RDP access**

For domain-joined catalogs, if the customer's VNet peering is functional, the customer can enable RDP access from their peered VNet to their Citrix-managed VNet. If the customer uses this option, the customer is responsible for accessing the VDAs and Cloud Connectors over the VNet peering. Source IP address ranges can be specified so RDP access can be restricted further, even within the customer's



internal network. The customer will need to use domain credentials to log in to these machines. If the customer is working with Citrix Support to resolve an issue, the customer may need to share these credentials with support personnel. After the issue is resolved, the customer is responsible for disabling RDP access. Keeping RDP access open from the customer's peered or on-premises network poses a security risk.

### Domain credentials

If the customer elects to use a domain-joined catalog, the customer is responsible for providing to the Citrix Virtual Apps and Desktops service a domain account (username and password) with permissions to join machines to the domain. When supplying domain credentials, the customer is responsible for adhering to the following security principles:

- **Auditable:** The account should be created specifically for Citrix Virtual Apps and Desktops service usage so that it is easy to audit what the account is used for.
- **Scoped:** The account requires only permissions to join machines to a domain. It should not be a full domain administrator.
- **Secure:** A strong password should be placed on the account.

Citrix is responsible for the secure storage of this domain account in an Azure Key Vault in the customer's Citrix-managed Azure subscription. The account is retrieved only if an operation requires the domain account password.

### More information

For related information, see:

- [Secure Deployment Guide for the Citrix Cloud Platform](#): Security information for the Citrix Cloud platform.
- [Technical security overview](#): Security information for the Citrix Virtual Apps and Desktops service
- [Third party notifications](#)

### Delivery methods

February 24, 2021

A single delivery method likely cannot meet all of your requirements.

You can consider several application delivery methods. Choosing the appropriate method helps improve scalability, management, and user experience.

- **Installed app:** The application is part of the base desktop image. The install process involves dll, exe, and other files copied to the image drive in addition to registry modifications. For details, see [Create machine catalogs](#).
- **Streamed app (Microsoft App-V):** The application is profiled and delivered to the desktops across the network on-demand. Application files and registry settings are placed in a container on the virtual desktop, isolated from the base operating system and each other. This action helps address compatibility issues. For details, see [App-V](#).
- **Layered app (Citrix App Layering):** Each layer contains a single application, agent, or operating system. By integrating one OS layer, one platform layer (for example, VDA) and many application layers, an administrator can easily create new, deployable images. Layering simplifies ongoing maintenance, as an OS, agent and application exists in a single layer. When you update the layer, all deployed images containing that layer are updated. See [Citrix App Layering](#).
- **Hosted Windows app:** An application installed on a multi-user Citrix Virtual Apps host and deployed as an application and not a desktop. A user accesses the hosted Windows app seamlessly from the VDI desktop or endpoint device, hiding the fact that the app is running remotely. For details, see [Create delivery groups](#).
- **Local app:** An application deployed on the endpoint device. The application interface appears within the user's hosted VDI session even though it runs on the endpoint. For details, see [Local App Access and URL redirection](#).
- **Remote PC Access:** Remote PC Access enables employees to remotely access their physical office PCs. When users access their office PCs, they can access all the applications, data, and resources they need to do their work. Remote PC Access eliminates the need to introduce and provide other tools to accommodate teleworking. For details, see [Remote PC Access](#).

For desktops, you can consider Citrix Virtual Apps published desktops or VDI desktops.

## Citrix Virtual Apps published apps and desktops

Use multi-session OS machines to deliver Citrix Virtual Apps published apps and published desktops.

### Use case:

- You want inexpensive server-based delivery to minimize the cost of delivering applications to many users, while providing a secure, high-definition user experience.
- Your users perform well-defined tasks and do not require personalization or offline access to applications. Users can include task workers such as call center operators and retail workers, or users that share workstations.
- Application types: any application.

### Benefits and considerations:

- Manageable and scalable solution within your data center.
- Most cost effective application delivery solution.

- Hosted applications are managed centrally and users cannot modify the application, providing a user experience that is consistent, safe, and reliable.
- Users must be online to access their applications.

**User experience:**

- User requests one or more applications from StoreFront, their Start menu, or a URL you provide to them.
- Applications are delivered virtually and display seamlessly in high definition on user devices.
- Depending on profile settings, user changes are saved when the user's application session ends. Otherwise, the changes are deleted.

**Process, host, and deliver applications:**

- Application processing takes place on hosting machines, rather than on the user devices. The hosting machine can be a physical or a virtual machine.
- Applications and desktops reside on a multi-session OS machine.
- Machines become available through machine catalogs.
- Machines from machine catalogs are organized into delivery groups that deliver the same set of applications to groups of users.
- Multi-session OS machines support delivery groups that host either desktops or applications, or both.

**Session management and assignment:**

- Multi-session OS machines run multiple sessions from a single machine to deliver multiple applications and desktops to multiple, simultaneously connected users. Each user requires a single session from which they can run all their hosted applications.

For example, a user logs on and requests an application. One session on that machine becomes unavailable to other users. A second user logs on and requests an application which that machine hosts. A second session on the same machine is now unavailable. If both users request more applications, no additional sessions are required because a user can run multiple applications using the same session. If two more users log on and request desktops, and two sessions are available on that machine, that one machine now uses four sessions to host four different users.

- Within the delivery group to which a user is assigned, a machine on the least loaded server is selected. A machine with session availability is randomly assigned to deliver applications to a user when that user logs on.

**VM hosted apps**

Use single-session OS machines to deliver VM hosted applications

**Use case:**

- You want a client-based application delivery solution that is secure, provides centralized management, and supports many users per host server. You want to provide those users with applications that display seamlessly in high-definition.
- Your users are internal, external contractors, third-party collaborators, and other provisional team members. Your users do not require offline access to hosted applications.
- Application types: Applications that might not work well with other applications or might interact with the operation system, such as Microsoft .NET framework. These types of applications are ideal for hosting on virtual machines.

**Benefits and considerations:**

- Applications and desktops on the image are securely managed, hosted, and run on machines within your data center, providing a more cost effective application delivery solution.
- Upon logon, users can be randomly assigned to a machine within a delivery group that is configured to host the same application. You can also statically assign a single machine to deliver an application to a single user each time that user logs on. Statically assigned machines allow users to install and manage their own applications on the virtual machine.
- Running multiple sessions is not supported on single-session OS machines. Therefore, each user consumes a single machine within a delivery group when they log on, and users must be online to access their applications.
- This method can increase the amount of server resources for processing applications and increase the amount of storage for users' personal vDisks.

**User experience:**

- The same seamless application experience as hosting shared applications on multi-session OS machines.

**Process, host, and deliver applications:**

- The same as multi-session OS machines except they are virtual single-session OS machines.

**Session management and assignment:**

- Single-session OS machines run a single desktop session from a single machine. When accessing applications only, one user can use multiple applications (and is not limited to a single application). The operating system sees each application as a new session.
- Within a delivery group, logged-on users can access either a statically assigned machine (each time the user logs on to the same machine), or a randomly assigned machine that is selected based on session availability.

**VDI desktops**

Use single-session OS machines to deliver Citrix Virtual Desktops VDI desktops.

VDI desktops are hosted on virtual machines and provide each user with a desktop operating system.

VDI desktops require more resources than Citrix Virtual Apps published desktops, but do not require that applications installed on them support server-based operating systems. Also, depending on the type of VDI desktop you choose, these desktops can be assigned to individual users. This allows users a high level of personalization.

When you create a machine catalog for VDI desktops, you create one of these types of desktops:

- **Random non-persistent desktop, also known as pooled VDI desktop:** Each time a user logs on to one of these desktops, that user connects to a desktop selected from a pool of desktops. That pool is based on a single image. All changes to the desktop are lost when the machine restarts.
- **Static non-persistent desktop:** During the first logon, a user is assigned a desktop from a pool of desktops. (Each machine in the pool is based on a single image.) After the first use, each time a user logs on to use one of these desktops, that user connects to the same desktop that was assigned on first use. All changes to the desktop are lost when the machine restarts.
- **Static persistent desktop:** Unlike other types of VDI desktops, users can fully personalize these desktops. During the first logon, a user is assigned a desktop from a pool of desktops. Subsequent logons from that user connect to the same desktop that was assigned on first use. Changes to the desktop are retained when the machine restarts.

## Get started: Plan and build a deployment

August 24, 2021

If you're not familiar with the components, terminology, and objects used with this service, see [Citrix Virtual Apps and Desktops service](#).

For a customer journey perspective, go to the [Citrix Success Center](#). The Success Center provides guidance for the five key stages of your Citrix journey: plan, build, rollout, manage, and optimize.

- The Success Center information is an essential partner to this product documentation.
- Success Center articles and guides offer a broad solution-based perspective. They also contain links to service-specific details in this product documentation.

If you're migrating from an on-premises Citrix Virtual Apps and Desktops deployment, see [Migrate to cloud](#).

### Important:

To ensure that you get important information about Citrix Cloud and the Citrix services you subscribe to, make sure you can receive all email notifications.

In the upper right corner of the Citrix Cloud console, expand the menu to the right of the customer

name and OrgID fields. Select **Account Settings**. On the **My Profile** tab, select all entries in the **Email Notifications** section.

## How to use this article

To set up your Citrix Virtual Apps and Desktops service deployment, complete the tasks summarized below. Links are provided to each task's details.

Review the entire process before starting the deployment, so you know what to expect. This article also links to other helpful information sources.

### Note:

If you plan to use the Quick Deploy interface to provision Microsoft Azure machines, follow the setup guidance in [Get started with Quick Deploy](#).

## Plan and prepare

Use the Success Center [Plan](#) guidance to help establish goals, define use cases and business objectives, identify potential risks, and create a project plan.

In the Citrix Tech Zone documentation, see a [step-by-step proof of concept guide for this service](#).

## Sign up


[Sign up](#) for a Citrix account and request a Citrix Virtual Apps and Desktops service demo.

## Set up a resource location

[Create a resource location and add Cloud Connectors](#).

More information:

- [What are resource locations and Cloud Connectors?](#)
- Video about installing Cloud Connectors:



If you're using the [Quick Deploy](#) interface to provision Azure VMs, Citrix creates the resource location and Cloud Connectors for you.

On the service's **Overview** page, this step is titled **Connect to infrastructure**.

### **Create a connection to the resource location**

After you add a resource location and Cloud Connectors, [create a connection](#) from the service's control plane to the resource location.

This step isn't necessary in either of the following cases:

- You're building a simple proof of concept deployment
- You're using the [Quick Deploy](#) interface to provision Azure VMs.

More information:

- [What are hosts?](#)
- [What are host connections?](#)


## Install VDAs

Each machine that delivers applications and desktops to users must have a Citrix Virtual Delivery Agent (VDA) installed on it.

- For a simple proof of concept deployment, download and install a VDA on one machine.
- If you're using an image to provision VMs, install a VDA on the image.
- For a [Remote PC Access](#) deployment, install the core version of the VDA for single-session OS on each physical office PC.

How-to and more information:

- [What are VDAs?](#)
- [Installation preparation and instruction](#)
- [Command-line VDA installation](#)
- Video about downloading and installing a VDA:



On the service's **Overview** page, this step is titled **Register resources**.

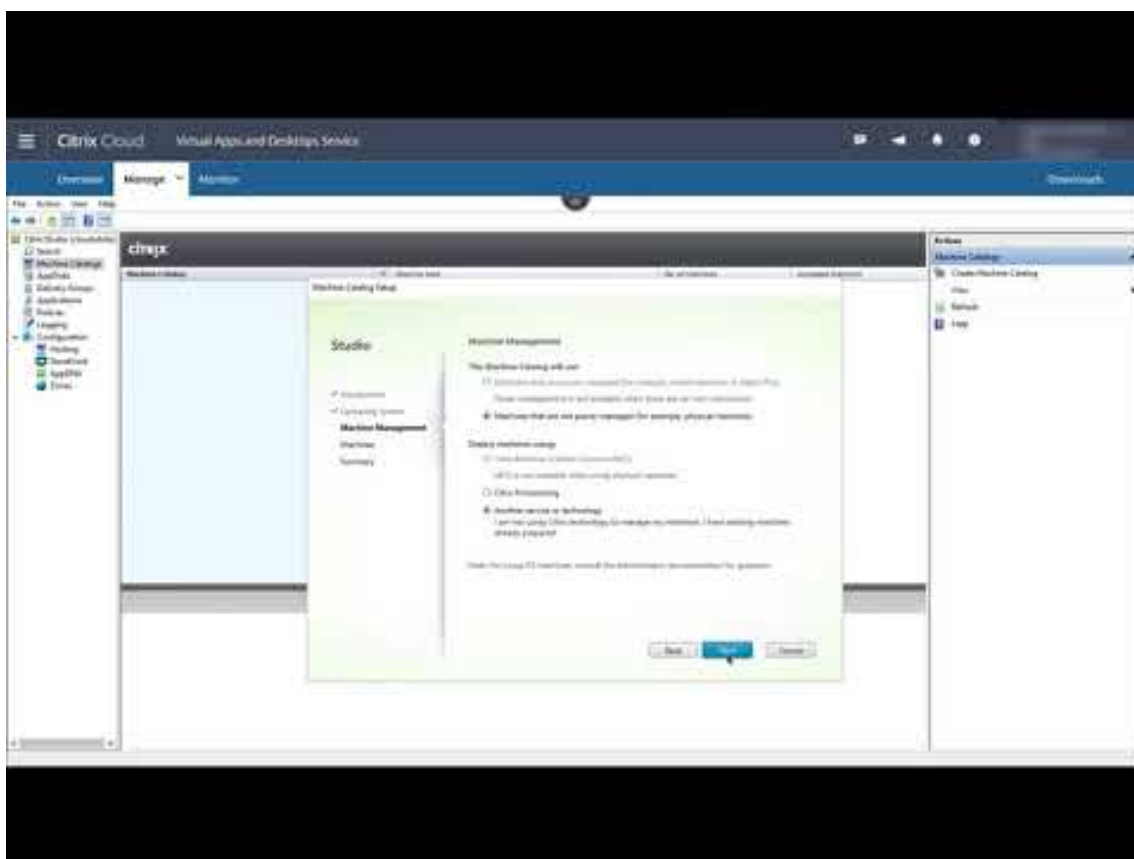
## Create a catalog

After you create a connection to your resource location (if needed), you create a catalog. If you're using the service's full configuration **Manage** interface, the workflow guides you automatically to this step.

How-to and more information:

- [What are catalogs?](#)
- [Create a catalog](#)
- Use the [Quick Deploy](#) interface to deploy a catalog containing Azure VMs.
- Video about creating a catalog using the full configuration management interface:





On the service's **Overview** page, this step is titled **Create a collection of resources**.


### **Create a delivery group**

After you create your first catalog, the **Manage** workflow guides you to create a delivery group.

This step isn't necessary if you're using the [Quick Deploy](#) interface to provision Azure VMs.

How-to and more information:

- [What are delivery groups?](#)
- [Create a delivery group](#)
- Video about how to create a delivery group:



On the service's **Overview** page, this step is titled **Assign users**.

## Deploy other components and technologies

After you complete the tasks above that set up the Citrix Virtual Apps and Desktops service deployment, follow the guidance in the [Build](#) area of the Citrix Success Center. You'll find information about provisioning and configuring other components and technologies in the Citrix solution, such as:

- [Citrix policies](#)
- [StoreFront](#)
- [App Layering](#)
- [Workspace Environment Management \(WEM\) Service](#)
- [Citrix Gateway service](#)
- [Zones](#)
- [Federated Authentication Service \(FAS\)](#)

Complete other tasks that apply to your configuration. For example, if you plan to deliver Windows Server workloads, [configure a Microsoft RDS License Server](#).

## Launch applications and desktops

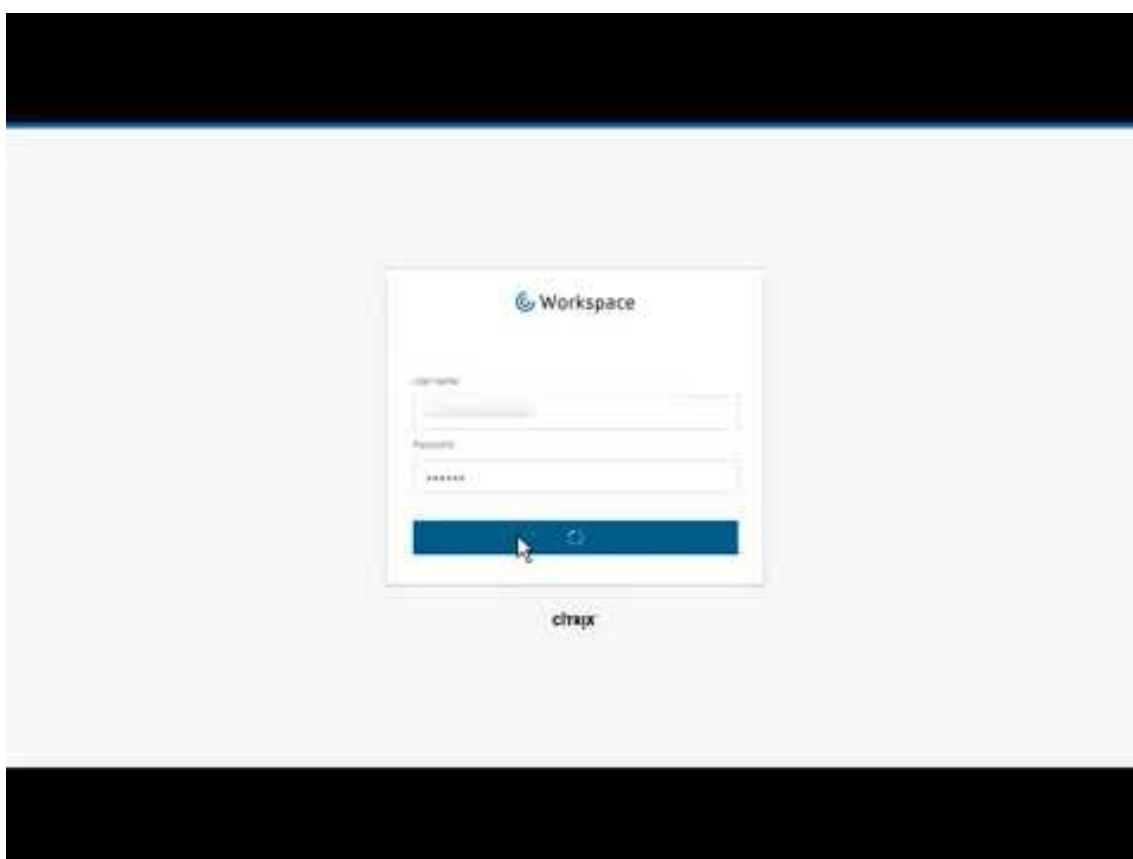
After you configure your deployment, publishing occurs automatically. Applications and desktops that you configured are available to users in their Citrix Workspace. A user simply navigates to their Workspace URL and selects an application or desktop, which launches immediately.

[Send the Workspace URL to your users.](#) You can find the workspace URL in two locations:

- From the Citrix Cloud console, select **Workspace Configuration** from the menu in the upper left corner. The **Access** tab contains the Workspace URL.
- From the Citrix Virtual Apps and Desktops service **Overview** page, the workspace URL appears near the bottom of the page.

More information:

- Video about users launching applications and desktops from their Workspace:



## More information

The Citrix Cloud Learning series provides education courses that are organized by your path:

- If you're new to the Citrix service, see [New to Citrix Virtual Apps and Desktops Learning Path](#).

- If you're migrating from an on-premises Citrix Virtual Apps and Desktops deployment, see [Migrating Citrix Virtual Apps and Desktops to Citrix Cloud Learning Path](#).

## Sign up for the service

July 21, 2021

### Introduction

You can subscribe to the Citrix Virtual Apps and Desktops service through Citrix or through the Azure Marketplace.

If you plan to use [Citrix Managed Azure](#), you can also order the Citrix Azure Consumption Fund through Citrix or through the Azure Marketplace.

- When you order through Citrix, you can order the Citrix Virtual Apps and Desktops service and the Citrix Azure Consumption Fund at the same time.
- When you order through Azure Marketplace, you first order the Citrix Virtual Apps and Desktops service. Then, you can place another order for the Citrix Azure Consumption Fund.

If you order only the Citrix Virtual Apps and Desktops service now, you can order the Citrix Azure Consumption Fund later, either through Azure Marketplace or your Citrix account representative.

### Demos and trials

You can evaluate the service by request through Citrix. From a trial, you can convert to a paid service subscription.

During a trial, you can optionally use a Citrix Managed Azure subscription for catalogs, images, and network connections. If you have Citrix-managed resources at the time you convert to a paid subscription, you must either purchase consumption or delete those Citrix-managed resources. If you do not purchase consumption, those resources are deleted automatically, which might affect users.

### If you currently subscribe to a Citrix Virtual Apps and Desktops service

Generally, a Citrix Cloud account allows you to subscribe to only one of the Citrix Virtual Apps and Desktops services (or one edition) at a time per Citrix OrgID. For example, you can subscribe to Citrix Virtual Apps and Desktops Premium edition OR Citrix Virtual Apps and Desktops Standard for Azure, but not both.

If you currently subscribe to a Citrix Virtual Apps and Desktops service, and want to subscribe to this service, you have two choices:

- Subscribe to this service using a different Citrix Cloud account (OrgID).
- Decommission the service you already have, and then order this service. For decommission instruction, see [CTX239027](#).

## Order through Citrix

You can order this service (and the Citrix Azure Consumption Fund) through Citrix Cloud or through your Citrix account representative.

Through Citrix Cloud:

- Follow the guidance in [Sign up for Citrix Cloud](#) to get a Citrix Cloud account and Organization ID.
- You can request a Citrix Virtual Apps and Desktops service demo. In the Citrix Virtual Apps and Desktops service tile, click **Request Demo**. Provide the requested information.

A Citrix representative will contact you to discuss your requirements, environment, and plans. Depending on our representative's assessment, you will be authorized to participate in an administrator demo or a proof of concept trial. For more information, see [Citrix Cloud Service Trials](#).

When you are authorized for a trial, the text on the Citrix Virtual Apps and Desktops service tile in the Citrix Cloud console changes to **Manage**.

## Order through Azure Marketplace

You can order the following Citrix offerings through Azure Marketplace:

- Citrix Virtual Apps and Desktops Standard for Azure
- Citrix Virtual Apps and Desktops Advanced edition
- Citrix Virtual Apps and Desktops Premium edition
- Workspace Premium Plus

If you plan to host your Citrix Virtual Apps and Desktops workloads on Microsoft Azure, and want to use a [Citrix Managed Azure](#) subscription, order the Citrix Azure Consumption Fund after ordering the Citrix Virtual Apps and Desktops service or Workspace Premium Plus.

With the Citrix Azure Consumption Fund, you're charged each month for your consumption, which can vary depending on the hosting resources you choose and the hours of use. You can review your consumption use through Citrix Cloud.

From the Azure Marketplace:

- You cannot combine the service and consumption fund in one order.
- The order process for the Citrix Azure Consumption Fund is essentially the same as ordering the Citrix service, but you must have previously ordered the service.

## Requirements for ordering through Azure Marketplace

- The OrgID of your Citrix Cloud account.
  - If you have a Citrix Cloud account, but don't know the OrgID, look in the upper right corner of the Citrix Cloud console. Or, look at the email you received when you created the account.
  - If you don't have a Citrix Cloud account, follow the guidance in [Sign up for Citrix Cloud](#).
- An Azure account and at least one Azure subscription in that account.

## Procedure for ordering through Azure Marketplace

Follow this procedure to order a Citrix Virtual Apps and Desktops service or Workspace Premium Plus through Azure Marketplace. (If you want to use Citrix Managed Azure, place another order for the Citrix Azure Consumption Fund, after you order the Citrix service.)

1. Sign in to the [Azure Marketplace](#) using your Azure account credentials.
2. Search for and then navigate to the Citrix offering you want to order.
3. Select **Get it now**.
4. On the **One more thing** message, fill in the required information, enable the consent check box, and then select **Continue**.
5. Review the tabs containing information about the product, plans, pricing, and usage. When you're ready, select a plan (if more than one is available), and then select **Set up + subscribe**.
6. On the **Basics** tab:
  - **Subscription:** Indicates the plan that you selected.
  - **Resource group:** Select or create a resource group.
  - **Name:** Enter a name for your subscription order so you can easily identify it later.
  - The **Plan** information shows the price for the selected plan, based on the billing term. To change the plan term, select **Change plan**. Select the term you want and select **Change plan**.
7. On the **Review + subscribe** tab, review the contact information, and update it, if needed. Review the basic subscription information. Select **Subscribe**.
8. On the **Subscription in progress** page, select **Configure account now**. (If the button is disabled, wait a moment.) You're taken to a Citrix activation page.
9. On the activation page:
  - Use the **Sign in** link to sign in to Citrix Cloud. A successful sign-in automatically populates the **Organization ID** field.

- **Quantity:** Enter the number of users. (An initial order must be at least 25.) An estimated price is displayed.
- Agree to the terms and conditions, and then select **Activate Order**.

### After ordering through Azure Marketplace

Citrix sends you an email when your service is provisioned. Provisioning can take a while. If you don't receive the email by the following day, contact [Citrix Support](#). When you receive the email from Citrix, you can begin using the service.

Fulfillment of a Citrix Azure Consumption Fund order does not take much time. When Citrix is notified of the order, a banner appears in the Citrix service console, indicating that a Citrix Managed Azure subscription will be prepared for you.

Do not delete the service resource in Azure. Deleting that resource cancels your subscription.

### What's next

After your order is fulfilled, continue with the next steps in [Plan and build a deployment](#).

For example:

- If you haven't already set up your hypervisor or cloud service, or Active Directory, see [Set up a resource location](#).
- If your host environment and Active Directory are already set up, see [Create a connection](#).

## Set up resource locations

June 18, 2021

Resource locations contain the resources required to deliver applications and desktops to users. You manage those items from Citrix Cloud and the service's **Manage** console. Typically, resources include:

- Active Directory domain controller.
- Hypervisors or cloud services, known as *hosts*.
- Virtual Delivery Agents (VDAs). VDAs are the machines containing the apps or desktop. Each machine also has a Citrix VDA installed. The term *VDA* often refers to the VDA software and the machine on which it is installed.
- Citrix Gateway (optional): To enable secure external access to the applications and desktops offered to users, add a Citrix Gateway VPX appliance to the resource location. Then set up Citrix Gateway.
- Citrix StoreFront servers (customer-managed).

- To communicate with Citrix Cloud, every resource location must contain a Citrix Cloud Connector. A minimum of two Cloud Connectors per resource location is recommended, for availability.

A resource location is considered a zone in a Citrix Virtual Apps and Desktops service environment. For more information, see [Zones](#).

To learn more about resource locations, see [Connect to Citrix Cloud](#).

## Host requirements

The hypervisor or cloud service where you provision VMs that deliver apps or desktops to users might have unique permission or setup requirements.

- If the hypervisor or cloud service requires virtual networks or other items, follow the guidance in its documentation.
- Create the appropriate virtual private cloud (VPC) or virtual networks for the machines you'll add to your resource location, if needed. For example, when using AWS, set up a VPC with public and private subnets.
- Create the appropriate rules to secure inbound and outbound internet traffic, and traffic between machines in the virtual network. For example, when using AWS, ensure the VPC's security group has the appropriate rules configured so that machines in the VPC are accessible to only the IP addresses you specify.

Review the article for the host type you're using.

- [Microsoft System Center Virtual Machine Manager virtualization environments](#)
- [Microsoft Azure Resource Manager virtualization environments](#)
- [Amazon Web Services \(AWS\) virtualization environments](#)
- [Citrix Hypervisor virtualization environments](#)
- [VMware virtualization environments](#)
- [Nutanix virtualization environments](#)
- [Google Cloud Platform virtualization environments](#)

Those articles also contain host-specific details you need when creating a catalog.

## Active Directory

Provision a Windows server, install Active Directory Domain Services, and promote it to a domain controller. For guidance, see the Microsoft Active Directory documentation.

- You must have at least one domain controller running Active Directory Domain Services.
- Do not install any Citrix components on a domain controller.
- Do not use a forward slash (/) when specifying Organizational Unit names in Studio.



For more information, see:

- [Active Directory functional levels](#)
- [Identity and access management](#) in Citrix Cloud.

## Cloud Connectors

The Cloud Connector is a group of services from Citrix Cloud that allow communication between the VDAs, customer-managed StoreFront, and the cloud-based Delivery Controller. You can install Cloud Connectors interactively or from the command line.

For complete Cloud Connector information, see:

- [Citrix Cloud Connector](#)
- [Technical details](#), which include system requirements
- [Proxy and firewall configuration](#)
- [Installation](#)
- [Connector updates](#)

## Size and scale considerations

When evaluating the Citrix Virtual Apps and Desktops service for sizing and scalability, consider all components. Research and test the configuration of the Cloud Connectors and the customer-managed StoreFront for your specific requirements. Undersizing the machines can impact system performance negatively.

The following articles contain size and scale testing information. They provide details of the tested maximum capacities, plus best practice recommendations for Cloud Connector machine configuration.

- [Scale and size considerations for Cloud Connectors](#)
- [Scale and size considerations for Local Host Cache](#)

## Add a resource location in Citrix Cloud

To add a resource location:

1. Sign in to [Citrix Cloud](#).
2. In the upper left menu, select **Resource Locations**.
3. If you have not already installed a Citrix Cloud Connector, you are prompted to download one.
4. After installing a Cloud Connector (and preferably at least two), in the Citrix Cloud console, enter a name for the resource location and then click **Save**. See [Name restrictions](#).

## Where to go next

- For a simple proof-of-concept deployment, [install a VDA](#) on a machine that will deliver apps or a desktop to your users.
- For a full deployment, [create a connection](#) to a resource location.
- [Review all the steps in the installation and configuration process](#)

## Microsoft Azure Resource Manager virtualization environments

September 8, 2021

Follow the guidance in this article when using the Microsoft Azure Resource Manager to provision virtual machines in your Citrix Virtual Apps or Citrix Virtual Desktops service deployment.

We assume you are familiar with the following:

- Azure Active Directory: <https://azure.microsoft.com/en-us/documentation/articles/active-directory-howto-tenant/>
- Consent framework: <https://azure.microsoft.com/en-us/documentation/articles/active-directory-integrating-applications/>
- Service principal: <https://azure.microsoft.com/en-us/documentation/articles/active-directory-application-objects/>

### Note:

Azure supports disk encryption at rest by default, using Azure-managed encryption keys. This form of encryption is used by all catalogs in MCS and requires no user configuration.

## Azure on-demand provisioning

When you use MCS to create machine catalogs in the Azure Resource Manager, the Azure on-demand provisioning feature:

- Reduces your storage costs
- Provides faster catalog creation
- Provides faster virtual machine (VM) power operations

For administrators, on-demand provisioning introduces no differences in the Full Configuration interface procedures for creating host connections and MCS machine catalogs. The differences lie in how and when resources are created and managed in Azure, and VM visibility in the Azure portal.

Before Citrix Virtual Apps and Desktops service integrated Azure on-demand provisioning, MCS behavior was different. When MCS created a catalog, VMs were created in Azure during the provisioning process.

With Azure on-demand provisioning, VMs are created only when Citrix Virtual Apps and Desktops initiates a power-on action, after the provisioning completes. A VM is visible in the Azure portal only when it is running. In the **Manage > Full Configuration** interface, VMs are visible, if they're running.

When you create an MCS catalog, the Azure portal displays the resource groups, network security group, storage accounts, network interfaces, base images, and identity disks. The Azure portal does not show a VM until Citrix Virtual Apps and Desktops initiates a power-on action for it. Then, the VM's status in the Full Configuration interface changes to **On**.

- For a pooled machine, the operating system disk and write-back cache exist only when the VM exists. The cache can result in significant storage savings if you routinely shut down machines (for example, outside of working hours).
- For a dedicated machine, the operating system disk is created the first time the VM is powered on. It remains in storage until the machine identity is deleted.

When Citrix Virtual Apps and Desktops initiates a power-off action for a VM, that machine identity is deleted in Azure. It no longer appears in the Azure portal. (The VM's status in the Full Configuration interface changes to **Off**.)

### **Catalogs created before on-demand provisioning**

Machine catalogs created before Citrix Virtual Apps and Desktops service supported the Azure on-demand provisioning feature (mid-2017) behaved differently. VMs in those catalogs are visible in the Azure portal if they're running. You cannot convert an image in a region other than where MCS provisions the catalog. The image is copied to a VHD in a conventional storage account in the catalog's region. It is then converted back to a managed disk.

On the **Storage and License Types** page of the catalog creation wizard, you can select a check box to use conventional storage accounts instead of managed disks. This check box is disabled when you are provisioning in an Azure region that does not support managed disks.

### **Create a connection to Azure Resource Manager**

[Create and manage connections](#) describes the wizards that create a connection. The following information covers details specific to Azure Resource Manager connections.

Considerations:

- Service principals must have been granted contributor role for the subscription.
- When creating the first connection, Azure prompts you to grant it the necessary permissions. For future connections you must still authenticate, but Azure remembers your previous consent and does not display the prompt again.
- Accounts used for authentication must be a co-administrator of the subscription.

- The account used for authentication must be a member of the subscription's directory. There are two types of accounts to be aware of: 'Work or School' and 'personal Microsoft account.' See [CTX219211](#) for details.
- While you can use an existing Microsoft account by adding it as a member of the subscription's directory, there can be complications if the user was previously granted guest access to one of the directory's resources. In this case, they might have a placeholder entry in the directory that does not grant them the necessary permissions, and an error is returned.

Rectify this by removing the resources from the directory and add them back explicitly. However, exercise this option carefully, because it has unintended effects for other resources that account can access.

- There is a known issue where certain accounts are detected as directory guests when they are actually members. Configurations like this typically occurs with older established directory accounts. Workaround: add an account to the directory, which takes the proper membership value.
- Resource groups are simply containers for resources, and they can contain resources from regions other than their own region. This can potentially be confusing if you expect resources displayed in a resource group's region to be available.
- Ensure that your network and subnet are large enough to host the number of machines you require. This requires some foresight, but Microsoft helps you specify the right values, with guidance about the address space capacity.

You can establish a host connection to Azure in two ways:

- Authenticate to Azure to create a service principal.
- Use the details from a previously created service principal to connect to Azure.

### Create a service principal

#### **Important:**

This feature is not yet available for Azure China and Azure Germany subscriptions.

Before you start, authenticate to Azure. Ensure:

- You have a user account in your subscription's Azure Active Directory tenant.
- The Azure AD user account is also a co-administrator for the Azure subscription you want to use for provisioning resources.
- Authentication requires global administrator permissions.

When you authenticate to Azure to create a service principal, an application is registered in Azure. A secret key (client secret) is created for the registered application.

To view the application ID, sign in to the Azure portal and navigate to **Azure Active Directory > App registrations > All applications**. The name of the application registered through the console (**Manage** tab) has a prefix `citrix-xd-`.

Display name	Application (client) ID	Created on	Certificates & secrets
ci CitrixCloud-StudioPool		9/5/2018	⚠ Expiring soon
xi xingltest		6/20/2021	✅ Current
ci christian24app		2/17/2020	❌ Expired
ps P2P Server		4/6/2020	-
lf leixpb2		2/17/2020	✅ Current
do Dome9-Connect-studiopool-integration		5/21/2019	❌ Expired
u lisa-demo		8/25/2020	✅ Current
do Dome9-Connect-CSO-Labs		5/22/2019	❌ Expired
aa Azure AD Domain Services Sync		8/30/2018	✅ Current
ci citrix-xd-		8/25/2020	✅ Current

To view the application's client secret, click the application on the **All Applications** tab and navigate to **Certificates & secrets > Client secrets**.

Description	Expires	Value	Secret ID
No description	8/25/2023	Hidden	

The client secret created through the console expires after two years. The registered application uses the client secret to authenticate to Azure AD. Be sure to change the client secret before it expires. To change it, complete the following steps:

1. In Azure, select **Azure Active Directory**.
2. From **App registrations** in Azure AD, select your application.
3. Select **Certificates & secrets**.
4. Select **Client secrets > New client secret**.
5. Provide a description of the secret and specify a duration. When you are done, select **Add**.

**Note:**

Be sure to save the client secret because you cannot retrieve it later.

6. Copy the client secret value.
7. In the console, edit the corresponding connection and replace the content in the **Application secret** field with the value you copied.

To authenticate to Azure to create a service principal, complete the following steps in the **Add Connection and Resources** wizard:

1. On the **Connection** page, select **Create a new connection**, the **Microsoft Azure** connection type, and your Azure environment.
2. Select which tools to use to create the virtual machines and then select **Next**.
3. On the **Connection Details** page, enter your Azure subscription ID and a name for the connection. After you enter the subscription ID, the **Create new** button is enabled.

**Note:**

The connection name can contain 1–64 characters, and cannot contain only blank spaces nor the characters `\ / ; : ## . * ? = < > | [ ] { } " ' ( ) ' .`

4. Select **Create new** and then enter the Azure Active Directory account user name and password.
5. Select **Sign in**.
6. Select **Accept** to give Citrix Virtual Apps and Desktops the listed permissions. Citrix Virtual Apps and Desktops creates a service principal that allows it to manage Azure resources on behalf of the specified user.
7. After you select **Accept**, you return to the **Connection** page in the wizard.

**Note:**

After you successfully authenticate to Azure, the **Create new** and **Use existing** buttons disappear. The **Connection successful** text appears, with a green check mark, indicating the successful connection to your Azure subscription.

8. On the **Connection Details** page, select **Next**.

**Note:**

You cannot proceed to the next page until you successfully authenticate to Azure and consent to giving the required permissions.

9. Configure resources for the connection. Resources comprise the region and the network.
  - On the **Region** page, select a region.
  - On the **Network** page, do the following:
    - Type a 1–64 character resource name to help identify the region and network combination. A resource name cannot contain only blank spaces nor the characters `\ / ; : ## . * ? = < > | [ ] { } " ' ( ) ' .`
    - Select a virtual network/resource group pair. (If you have more than one virtual network with the same name, pairing the network name with the resource group provides unique combinations.) If the region you selected on the previous page does not

have any virtual networks, return to that page and select a region that has virtual networks.

10. On the **Summary** page, view a summary of settings and select **Finish** to complete your setup.

## Use the details from a previously created service principal to connect to Azure

To create a service principal manually, connect to your Azure Resource Manager subscription and use the PowerShell cmdlets provided in the following sections.

Prerequisites:

- **\$SubscriptionId:** Azure Resource Manager `SubscriptionID` for the subscription where you want to provision VDAs.
- **\$AADUser:** Azure AD user account for your subscription's AD tenant. Make the `$AADUser` the co-administrator for your subscription.
- **\$ApplicationName:** Name for the application to be created in Azure AD.
- **\$ApplicationPassword:** Password for the application. You use this password as the application secret when creating the host connection.

To create a service principal:

1. Connect to your Azure Resource Manager subscription.

```
Connect-AzAccount
```

2. Select the Azure Resource Manager subscription where you want to create the service principal.

```
Select-AzSubscription -SubscriptionID $SubscriptionId;
```

3. Create the application in your AD tenant.

```
$AzureADApplication = New-AzADApplication -DisplayName $ApplicationName  
-HomePage "https://localhost/$ApplicationName"-IdentifierUri https://  
$ApplicationName -Password $ApplicationPassword
```

4. Create a service principal.

```
New-AzADServicePrincipal -ApplicationId $AzureADApplication.ApplicationId
```

5. Assign a role to the service principal.

```
New-AzRoleAssignment -RoleDefinitionName Contributor -ServicePrincipalName  
$AzureADApplication.ApplicationId -scope /subscriptions/$SubscriptionId
```

6. From the output window of the PowerShell console, note the `ApplicationId`. You provide that ID when creating the host connection.

In the **Add Connection and Resources** wizard:

1. On the **Connection** page, select **Create a new connection**, the **Microsoft Azure** connection type, and your Azure environment.
2. Select which tools to use to create the virtual machines and then select **Next**.
3. On the **Connection Details** page, enter your Azure subscription ID and a name for the connection.

**Note:**

The connection name can contain 1–64 characters, and cannot contain only blank spaces nor the characters `\ / ; : ## . * ? = < > | [ ] { } " ' ( ) ' .`

4. Select **Use existing**. In the **Existing Service Principal Details** window, enter the following settings for the existing service principal. After you enter the details, the **Save** button is enabled. Select **Save**. You cannot progress beyond this page until you provide valid details.
  - **Subscription ID**. Enter your Azure subscription ID. To obtain your subscription ID, sign in to the Azure portal and navigate to **Subscriptions > Overview**.
  - **Active Directory ID** (tenant ID). Enter the Directory (tenant) ID of the application you registered with Azure AD.
  - **Application ID**. Enter the Application (client) ID of the application you registered with Azure AD.
  - **Application secret**. Create a secret key (client secret). The registered application uses the key to authenticate to Azure AD. We recommend that you change keys regularly for security purposes. Be sure to save the key because you cannot retrieve the key later.
  - **Authentication URL**. This field is automatically populated and is not editable.
  - **Management URL**. This field is automatically populated and is not editable.
  - **Storage suffix**. This field is automatically populated and is not editable.
5. After selecting **Save**, you return to the **Connection Details** page. Select **Next** to proceed to the next page.
6. Configure resources for the connection. Resources comprise the region and the network.
  - On the **Region** page, select a region.
  - On the **Network** page, do the following:
    - Type a 1–64 character resource name to help identify the region and network combination. A resource name cannot contain only blank spaces nor the characters `\ / ; : ## . * ? = < > | [ ] { } " ' ( ) ' .`
    - Select a virtual network/resource group pair. (If you have more than one virtual network with the same name, pairing the network name with the resource group provides unique combinations.) If the region you selected on the previous page does not



have any virtual networks, return to that page and select a region that has virtual networks.

7. On the **Summary** page, view a summary of settings and select **Finish** to complete your setup.

## Create a machine catalog using an Azure Resource Manager image

This information is a supplement to the guidance in [Create machine catalogs](#).

An image is the template that is used to create the VMs in a machine catalog. Before creating the machine catalog, create an image in Azure Resource Manager. For general information about images, see [Create machine catalogs](#).

In the machine catalog creation wizard:

- The **Operating System** and **Machine Management** pages do not contain Azure-specific information. Follow the guidance in the [Create machine catalogs](#) article.
- On the **Master Image** page, select a resource group and then navigate (drill down) through the containers to the Azure VHD or the Shared Image Gallery or the Azure ImageVersion you want to use as the image. The VHD or the ImageVersion must have a Citrix VDA installed on it. If the VHD is attached to a VM, stop the VM.

If you want to use a machine profile, enable the **Use a machine profile** check box and then select a machine profile from the list. (In Azure, you created VMs as machine profile VMs.) The list displays only machine profiles that are from the same resource group where the selected image resides. The image can inherit the following configurations from the selected machine profile:

- Accelerated networking
  - Boot diagnostics
  - Host disk caching (relating to OS and MCSIO disks)
  - Machine size (unless otherwise specified)
  - Tags placed on the VM
- The **Storage and License Types** page appears only when using an Azure Resource Manager image.

The screenshot shows the 'Machine Catalog Setup' wizard. The left sidebar lists steps 1 through 11, with step 5 'Storage and License Types' highlighted. The main content area is titled 'Storage and License Types' and contains the following text: 'Select the type of Locally Redundant Storage (LRS) to use for this Machine Catalog. LRS makes multiple synchronous copies of your data within a single data center. The LRS type you select affects the machine sizes offered later in this wizard.' Below this text are three radio button options: 'Premium SSD (supports I/O-intensive workloads with significantly high throughput and low latency)' (selected), 'Standard SSD', and 'Standard HDD'. A paragraph follows: 'You can use Windows volume licenses to provision VMs in Azure at the base compute rate. To verify that your volume licensing agreement with Microsoft qualifies for the Azure base compute rate, consult Microsoft.' Below this are three radio button options: 'Use my Windows Server licenses' (selected), 'Use my Windows 10 licenses', and 'Use Azure Windows Server licenses'. At the bottom of the main area is a checkbox labeled 'Place image in Azure Shared Image Gallery' with a help icon. At the bottom of the wizard are three buttons: 'Back', 'Next', and 'Cancel'.

You have the following storage types to use for the machine catalog:

- **Premium SSD.** Offers a high-performance, low-latency disk storage option suitable for VMs with I/O-intensive workloads.
- **Standard SSD.** Offers a cost-effective storage option that is suitable for workloads that require consistent performance at lower IOPS levels. An Azure identity disk is always created using Standard SSD.
- **Standard HDD.** Offers a reliable, low-cost disk storage option suitable for VMs that run latency-insensitive workloads.

The storage type determines which machine sizes are offered on the **Virtual Machines** page of the wizard. MCS configures premium and standard disks to use Locally Redundant Storage (LRS). LRS makes multiple synchronous copies of your disk data within a single data center. For details about Azure storage types and storage replication, see the following:

- <https://docs.microsoft.com/en-us/azure/storage/common/storage-introduction/>
- <https://azure.microsoft.com/en-us/documentation/articles/storage-premium-storage/>
- <https://azure.microsoft.com/en-us/documentation/articles/storage-redundancy/>

Select whether to use existing Windows licenses. Using Windows licenses along with Windows images (Azure platform support images or custom images) lets you run Windows VMs in Azure

at a reduced cost. There are two types of licenses:

- **Windows Server license.** Lets you use your Windows Server or Azure Windows Server licenses, allowing you to use Azure Hybrid Benefits. For details, see <https://azure.microsoft.com/en-us/pricing/hybrid-benefit/>. Azure Hybrid Benefit reduces the cost of running VMs in Azure to the base compute rate, waiving the cost of extra Windows Server licenses from the Azure gallery.
- **Windows Client license.** Lets you bring your Windows 10 licenses to Azure, allowing you to run Windows 10 VMs in Azure without the need for extra licenses. For details, see [Client Access Licenses and Management Licenses](#).

**Note:**

The Windows Client license option varies depending on the operating system you select during machine catalog setup. If you select **Multi-session OS**, the option appears as **Use my Windows 10 licenses**. If you select **Single-session OS**, the option appears as **Use my Windows Client licenses**.

You can verify that the provisioned VM is using the licensing benefit by running the following PowerShell command: `Get-AzM -ResourceGroup MyResourceGroup -Name MyVM`.

- For the Windows Server license type, verify that the license type is **Windows\_Server**. More instructions are available at <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/hybrid-use-benefit-licensing/>.
- For the Windows Client license type, verify that the license type is **Windows\_Client**. More instructions are available at <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/windows-desktop-multitenant-hosting-deployment/>.

Alternatively, you can use the `Get-Provscheme` PowerShell SDK to perform the verification. For example: `Get-Provscheme -ProvisioningSchemeName "My Azure Catalog"`. For more information about this cmdlet, see <https://developer-docs.citrix.com/projects/citrix-virtual-apps-desktops-sdk/en/latest/MachineCreation/Get-ProvScheme/>.

Azure Managed Disks are used for VMs in the catalog by default. If you want to use regular storage accounts instead, select the check box at the bottom of the page. The **Use unmanaged disks instead of Azure Managed Disks for VMs in this catalog** check box is available only in the *legacy console*.

Azure Shared Image Gallery (SIG) is a repository for managing and sharing images. It lets you make your images available throughout your organization. We recommend that you store an image in SIG when creating large non-persistent machine catalogs because doing that enables faster reset of VDA OS disks. The **Place image in Azure Shared Image Gallery** check box is available only in the *web-based console*.

- On the **Virtual Machines** page, indicate how many VMs you want to create. You must specify at least one. Select a machine size. After you create a catalog, you cannot change the machine size. If you later want a different size, delete the catalog and then create a catalog that uses the same image and specifies the desired machine size.

If you select the **Place image in Azure Shared Image Gallery** check box on the **Storage and License Types** page, the Azure Shared Image Gallery settings section appears, letting you specify more SIG settings:

- **Ratio of virtual machines to image replicas.** Lets you specify the ratio of virtual machines to image replicas that you want Azure to keep. By default, Azure keeps a single image replica for every 40 non-persistent machines. For persistent machines, that number defaults to 1,000.
- **Maximum replica count.** Lets you specify the maximum number of image replicas that you want Azure to keep. The default is 10.

Virtual machine names cannot contain non-ASCII or special characters.

- (When using MCS) On the **Resource Groups** page, choose whether to create resource groups or use existing groups.
  - If you choose to create resource groups, select **Next**.
  - If you choose to use existing resource groups, select groups from the **Available Provisioning Resource Groups** list. **Remember:** Select enough groups to accommodate the machines you're creating in the catalog. A message appears if you choose too few. You might want to select more than the minimum required if you plan to add more VMs to the catalog later. You can't add more resource groups to a catalog after the catalog is created.

For more information, see Azure resource groups.

- The **Network Cards**, **Computer Accounts**, and **Summary** pages do not contain Azure-specific information. Follow the guidance in the [Create Machine Catalogs](#) article.

Complete the wizard.

## Azure throttling

Azure Resource Manager throttles requests for subscriptions and tenants, routing traffic based on defined limits, tailored to the specific needs of the provider. See [Throttling Resource Manager requests](#) on the Microsoft site for more information. Limits exist for subscriptions and tenants, where managing many machines can become problematic. For example, a subscription containing many machines might experience performance problems related to power operations.

**Tip:**

For more information, see [Improving Azure performance with Machine Creation Services](#).

To help mitigate these issues, Citrix Virtual Apps and Desktops service allows you to remove MCS internal throttling to use more of the available request quota from Azure.

Citrix recommends the following optimal settings when powering VMS on/off in large subscriptions, for example, those containing 1000 VMs:

- Absolute simultaneous operations - 500
- Maximum new operations per minute - 2000
- Max concurrency of operations - 500

Use the Full Configuration interface to configure Azure operations for a given host connection:

1. From **Manage > Full Configuration**, select **Hosting** in the left pane.
2. Select an Azure-related connection to edit it.
3. In the **Edit Connection** screen, select **Advanced**.
4. In the **Advanced** screen, use the configuration options to specify the number of simultaneous actions and maximum new actions per minute, and any additional connection options.

**Edit Connection '819azure'**

Connection Properties  
**Advanced**  
 Scopes

**Advanced**

Use these settings to specify a maximum number of simultaneous actions (or concurrent machines) per hosting connection. For simultaneous actions, specify both settings. The lower value overrides the higher value. [Learn more](#)

	Absolute	Percentage(%)
Simultaneous actions (all types): ?	500	100
Maximum new actions per minute:	2000	

Connection options:

Use this setting only when Citrix Technical Support or the product documentation makes the recommendation.

Save Cancel Apply

MCS supports 500 maximum concurrent operations by default. You can use the PowerShell information to set the maximum number of concurrent operations.

Use the **PowerShell** property, `MaximumConcurrentProvisioningOperations`, to specify the maximum number of concurrent Azure provisioning operations. When using this property in the provisioning scheme, consider:

- MCS supports 500 maximum concurrent operations by default, using the custom properties parameter `MaximumConcurrentProvisioningOperations`.
- Configure the `MaximumConcurrentProvisioningOperations` parameter using the PowerShell command `Set-ProvScheme`.

### Azure resource groups

Azure provisioning resource groups provide a way to provision the VMs that provide applications and desktops to users. You can add existing empty Azure resource groups when you create an MCS machine catalog, or have new resource groups created for you. For information about Azure resource groups, see the [Microsoft documentation](#).

### Azure Resource Group Usage

There is no limit on the number of virtual machines, managed disks, snapshots, and images per Azure Resource Group. (The limit of 240 VMs per 800 managed disks per Azure Resource Group has been removed.)

- When using a full scope service principal to create a machine catalog, MCS creates only one Azure Resource Group and uses that group for the catalog.
- When using a narrow scope service principal to create a machine catalog, you must supply an empty, pre-created Azure Resource Group for the catalog.

### Azure ephemeral disks

An [Azure ephemeral disk](#) allows you to repurpose the cache disk to store the OS disk for an Azure-enabled virtual machine. This functionality is useful for Azure environments that require a higher performant SSD disk over a standard HDD disk. To use ephemeral disks, you must set the custom property `UseEphemeralOsDisk` to **true** when running `New-ProvScheme`.

#### Note:

If the custom property `UseEphemeralOsDisk` is set to **false** or a value is not specified all provisioned VDAs continue to use a provisioned OS disk.

The following is an example set of custom properties to use in the provisioning scheme:

```
1  "CustomProperties": [  
2      {  
3  
4          "Name": "UseManagedDisks",  
5          "Value": "true"  
6      }  
7  ,  
8      {  
9  
10         "Name": "StorageAccountType",  
11         "Value": "Standard_LRS"  
12     }  
13  ,  
14     {  
15  
16         "Name": "UseSharedImageGallery",  
17         "Value": "true"  
18     }  
19  ,  
20     {  
21  
22         "Name": "SharedImageGalleryReplicaRatio",  
23         "Value": "40"  
24     }  
25  ,  
26     {  
27  
28         "Name": "SharedImageGalleryReplicaMaximum",  
29         "Value": "10"  
30     }  
31  ,  
32     {  
33  
34         "Name": "LicenseType",  
35         "Value": "Windows_Server"  
36     }  
37  ,  
38     {  
39  
40         "Name": "UseEphemeralOsDisk",  
41         "Value": "true"  
42     }  
43  
44  ],
```

## How to create machines using ephemeral OS disks

Ephemeral OS disks are controlled based on the `UseEphemeralOsDisk` property in the `CustomProperties` parameter.

## Important considerations for ephemeral disks

To use provision ephemeral OS disks using `New-ProvScheme`, consider the following constraints:

- The VM size used for the catalog (the service offering) must support ephemeral OS disks.
- The cache disk size represented by the VM size must be greater than or equal to the size of the OS disk.

Also consider these issues when:

- Creating the provisioning scheme.
- Modifying the provisioning scheme.
- Updating the image.

## Azure server side encryption

Citrix Virtual Apps and Desktops service supports customer-managed encryption keys for Azure managed disks through Azure Key Vault. With this support you can manage your organizational and compliance requirements by encrypting the managed disks of your machine catalog using your own encryption key. For more information, see [Server-side encryption of Azure Disk Storage](#).

When using this feature for managed disks:

- To change the key that the disk is encrypted with, you change the current key in the `DiskEncryptionSet`. All resources associated with that `DiskEncryptionSet` change to be encrypted with the new key.
- When you disable or delete your key, any VMs with disks using that key automatically shut down. After shutting down, the VMs are not usable unless the key is enabled again or you assign a new key. Any catalog using the key cannot be powered on, and you cannot add VMs to it.

## Important considerations when using customer-managed encryption keys

Consider the following when using this feature:

- All resources related to your customer-managed keys (Azure Key Vaults, disk encryption sets, VMs, disks, and snapshots) must reside in the same subscription and region.



- Once you have enabled the customer-managed encryption key you cannot disable it later. If you want to disable or remove the customer-managed encryption key, copy all the data to a different managed disk that is not using the customer-managed encryption key.
- Disks created from encrypted custom images using server-side encryption and customer-managed keys must be encrypted using the same customer-managed keys. These disks must be in the same subscription.
- Snapshots created from disks that are encrypted with server-side encryption and customer-managed keys must be encrypted with the same customer-managed keys.
- Disks, snapshots, and images encrypted with customer-managed keys cannot move to another resource group and subscription.
- Managed disks currently or previously encrypted using Azure Disk Encryption cannot be encrypted using customer-managed keys.
- You can only create up to 50 disk encryption sets per region, per subscription.

**Note:**

See [Quickstart: Create a Key Vault using the Azure portal](#) for information on configuring Azure server side encryption.

## Azure shared image gallery

Use Azure Shared Image Gallery as a published image repository for MCS provisioned machines in Azure. You can store a published image in the gallery to accelerate the creation and hydration of OS disks, improving boot and application launch times for non-persistent VMs. Shared image gallery contains the following three elements:

- Gallery. Images are stored here. MCS creates one gallery for each machine catalog.
- Gallery Image Definition. This definition includes information (operating system type and state, Azure region) about the published image. MCS creates one image definition for each image created for the catalog.
- Gallery Image Version. Each image in a Shared Image Gallery can have multiple versions, and each version can have multiple replicas in different regions. Each replica is a full copy of the published image. Citrix Virtual Apps and Desktops service creates one Standard\_LRS image version (version 1.0.0) for each image with the appropriate number of replicas in the catalog's region, based on the number of machines in the catalog, the configured replica ratio, and the configured replica maximum.

**Note:**

Shared Image Gallery functionality only works with managed disks. It is not available for legacy

machine catalogs.

For more information, see [Azure shared image gallery overview](#).

## Configure Shared Image Gallery

Use the `New-ProvScheme` command to create a provisioning scheme with Shared Image Gallery support. Use the `Set-ProvScheme` command to enable or disable this feature for a provisioning scheme and to change the replica ratio and replica maximum values.

Three custom properties were added to provisioning schemes to support the Shared Image Gallery feature:

### `UseSharedImageGallery`

- Defines whether to use the Shared Image Gallery to store the published images. If set to **True**, the image is stored as a Shared Image Gallery image, otherwise the image is stored as a snapshot.
- Valid values are **True** and **False**.
- If the property is not defined, the default value is **False**.

### `SharedImageGalleryReplicaRatio`

- Defines the ratio of machines to gallery image version replicas.
- Valid values are integer numbers greater than 0.
- If the property is not defined, default values are used. The default value for persistent OS disks is 1000 and the default value for non-persistent OS disks is 40.

### `SharedImageGalleryReplicaMaximum`

- Defines the maximum number of replicas for each gallery image version.
- Valid values are integer numbers greater than 0.
- If the property is not defined, the default value is 10.
- Azure currently supports up to 10 replicas for a gallery image single version. If the property is set to a value greater than that supported by Azure, MCS attempts to use the specified value. Azure generates an error, which MCS logs then leaves the current replica count unchanged.

#### **Tip:**

When using Shared Image Gallery to store a published image for MCS provisioned catalogs, MCS sets the gallery image version replica count based on the number of machines in the catalog, the replica ratio, and the replica maximum. The replica count is calculated by dividing the number of machines in the catalog by the replica ratio (rounding up to the nearest integer value) and then capping the value at the maximum replica count. For example, with a replica ratio of 20

and a maximum of 5, 0–20 machines have one replica created, 21–40 have 2 replicas, 41–60 have 3 replicas, 61–80 have 4 replicas, 81+ have 5 replicas.

### Use case: Updating the Shared Image Gallery replica ratio and replica max

The existing machine catalog uses Shared Image Gallery. Use the `Set-ProvScheme` command to update the custom properties for all existing machines in the catalog and any future machines:

```
1 Set-ProvScheme -ProvisioningSchemeName catalog-name -CustomProperties '  
  <CustomProperties xmlns="http://schemas.citrix.com/2014/xd/  
  machinecreation" xmlns:xsi="http://www.w3.org/2001/XMLSchema-  
  instance"> <Property xsi:type="StringProperty" Name="  
  StorageAccountType" Value="Standard_LRS"/> <Property xsi:type="  
  StringProperty" Name="UseManagedDisks" Value="True"/> <Property xsi:  
  type="StringProperty" Name="UseSharedImageGallery" Value="True"/> <  
  Property xsi:type="IntProperty" Name="SharedImageGalleryReplicaRatio  
  " Value="30"/> <Property xsi:type="IntProperty" Name="  
  SharedImageGalleryReplicaMaximum" Value="20"/></CustomProperties> '  
2 <!--NeedCopy-->
```

### Use Case: Converting a snapshot catalog to a Shared Image Gallery catalog

For this use case:

1. Run `Set-ProvScheme` with the `UseSharedImageGallery` flag set to **True**. Optionally include the `SharedImageGalleryReplicaRatio` and `SharedImageGalleryReplicaMaximum` properties.
2. Update the catalog.
3. Power cycle the machines to force an update.

For example:

```
1 Set-ProvScheme -ProvisioningSchemeName catalog-name -CustomProperties '  
  <CustomProperties xmlns="http://schemas.citrix.com/2014/xd/  
  machinecreation" xmlns:xsi="http://www.w3.org/2001/XMLSchema-  
  instance"> <Property xsi:type="StringProperty" Name="  
  StorageAccountType" Value="Standard_LRS"/> <Property xsi:type="  
  StringProperty" Name="UseManagedDisks" Value="True"/> <Property xsi:  
  type="StringProperty" Name="UseSharedImageGallery" Value="True"/> <  
  Property xsi:type="IntProperty" Name="SharedImageGalleryReplicaRatio  
  " Value="30"/> <Property xsi:type="IntProperty" Name="  
  SharedImageGalleryReplicaMaximum" Value="20"/></CustomProperties> '  
2 <!--NeedCopy-->
```

**Tip:**

The parameters `SharedImageGalleryReplicaRatio` and `SharedImageGalleryReplicaMaximum` are not required. After the `Set-ProvScheme` command completes the Shared Image Gallery image has not yet been created. Once the catalog is configured to use the gallery, the next catalog update operation stores the published image in the gallery. The catalog update command creates the gallery, the gallery image, and the image version. Power cycling the machines updates them, at which point the replica count is updated, if appropriate. From that time, all existing non-persistent machines are reset using the Shared Image Gallery image and all newly provisioned machines are created using the image. The old snapshot is cleaned up automatically within a few hours.

**Use Case: Converting a Shared Image Gallery Catalog to a snapshot catalog**

For this use case:

1. Run `Set-ProvScheme` with the `UseSharedImageGallery` flag set to **False** or not defined.
2. Update the catalog.
3. Power cycle the machines to force an update.

For example:

```
1 Set-ProvScheme -ProvisioningSchemeName catalog-name -CustomProperties '  
  <CustomProperties xmlns="http://schemas.citrix.com/2014/xd/  
  machinecreation" xmlns:xsi="http://www.w3.org/2001/XMLSchema-  
  instance"> <Property xsi:type="StringProperty" Name="  
  StorageAccountType" Value="Standard_LRS"/> <Property xsi:type="  
  StringProperty" Name="UseManagedDisks" Value="True"/> <Property xsi:  
  type="StringProperty" Name="UseSharedImageGallery" Value="False"/></  
  CustomProperties>'  
2 <!--NeedCopy-->
```

**Tip:**

Unlike updating from a snapshot to a Shared Image Gallery catalog, the custom data for each machine is not yet updated to reflect the new custom properties. Run the following command to see the original Shared Image Gallery custom properties: `Get-ProvVm -ProvisioningSchemeName catalog-name`. After the `Set-ProvScheme` command completes the image snapshot has not yet been created. Once the catalog is configured to not use the gallery, the next catalog update operation stores the published image as a snapshot. From that time, all existing non-persistent machines are reset using the snapshot and all newly provisioned machines are created from the snapshot. Power cycling the machines updates them, at which point the custom machine data is updated to reflect that `UseSharedImageGallery` is

set to **False**. The old Shared Image Gallery assets (gallery, image, and version) are automatically cleaned up within a few hours.

## Provision machines into specified availability zones

You can provision machines into specific Availability Zones in Azure environments. You can achieve that using the Full Configuration interface or PowerShell.

(If using the legacy console, you must use PowerShell.)

### Note:

If no zones are specified, MCS lets Azure place the machines within the region. If more than one zone is specified, MCS randomly distributes the machines across them.

## Configuring Availability Zones in the Full Configuration interface

When creating a machine catalog, you can specify Availability Zones into which you want to provision machines. On the **Virtual Machines** page, select one or more Availability Zones where you want to create machines.

There are two reasons that no Availability Zones are available: The region has no Availability Zones or the selected machine size is unavailable.

## Configuring Availability Zones through PowerShell

Using PowerShell, you can view the service offering inventory items by using `Get-Item`. For example, to view the *Eastern US region Standard\_B1ls* service offering:

```
1 $serviceOffering = Get-Item -path "XDHyp:\Connections\my-connection-  
   name\East US.region\serviceoffering.folder\Standard_B1ls.  
   serviceoffering"  
2 <!--NeedCopy-->
```

To view the zones, use the `AdditionalData` parameter for the item:

```
$serviceOffering.AdditionalData
```

If Availability Zones are not specified, there is no change in how machines are provisioned.

To configure Availability Zones through PowerShell, use the **Zones** custom property available with the `New-ProvScheme` operation. The **Zones** property defines a list of Availability Zones to provision machines into. Those zones can include one or more Availability Zones. For example, `<Property xsi:type="StringProperty"Name="Zones"Value="1, 3"/>` for Zones 1 and 3.

Use the `Set-ProvScheme` command to update the zones for a provisioning scheme.

If an invalid zone is provided, the provisioning scheme is not updated, and an error message appears providing instructions on how to fix the invalid command.

**Tip:**

If you specify an invalid custom property, the provisioning scheme is not updated and a relevant error message appears.

## Azure ephemeral disk

[Azure ephemeral disks](#) allow you to repurpose the cache disk to store the OS disk for an Azure-enabled virtual machine. This functionality is useful for Azure environments that require a higher performant SSD disk over a standard HDD disk.

**Note:**

Persistent catalogs do not support ephemeral OS disks. Also, when using this feature, consider that the extra performant disk incurs an extra cost. It's beneficial to reuse the cache disk to store the OS disk instead of paying for an extra managed disk.

Ephemeral OS disks require that your provisioning scheme use managed disks and a Shared Image Gallery. For more information, see [Azure shared image gallery](#).

## Using PowerShell to configure an ephemeral disk

To configure an Azure ephemeral OS disk for a catalog, use the `UseEphemeralOsDisk` parameter in `Set-ProvScheme`. Set the value of the `UseEphemeralOsDisk` parameter to **true**.

**Note:**

To use this feature, you must also enable the parameters `UseManagedDisks` and `UseSharedImageGallery`.

For example:

```
1 Set-ProvScheme -ProvisioningSchemeName catalog-name -CustomProperties <
   CustomProperties xmlns="http://schemas.citrix.com/2014/xd/
   machinecreation" xmlns:xsi="http://www.w3.org/2001/XMLSchema-
   instance">
2 <Property xsi:type="StringProperty" Name="UseManagedDisks" Value="true"
   />
3 <Property xsi:type="StringProperty" Name="UseSharedImageGallery" Value=
   "true" />
4 <Property xsi:type="StringProperty" Name="UseEphemeralOsDisk" Value="
   true" />
5 </CustomProperties>'
```

```
6 <!--NeedCopy-->
```

## Storage types

Select different storage types for virtual machines in Azure environments that use MCS. For target VMs, MCS supports:

- OS disk: premium SSD, SSD or HDD
- Write back cache disk: premium SSD, SSD, or HDD

When using these storage types, consider the following:

- Ensure that your VM supports the selected storage type.
- If your configuration uses an Azure ephemeral disk, enabling write back cache always uses standard HDD storage. For implementations using write back cache, enable the `WBCDiskStorageType` parameter in `CustomProperties`. If you do not specify the `WBCDiskStorageType` parameter, the `StorageAccountType` uses the default value.

### Tip:

`StorageAccountType` is configured for an OS type and storage account. `WBCDiskStorageType` is configured for write back cache storage type. For a normal catalog, `StorageAccountType` is required. If `WBCDiskStorageType` is not configured, the `StorageAccountType` is used as the default for `WBCDiskStorageType`.

`WBCDiskStorageType` is not configured, then `StorageAccountType` will be used as the default for `WBCDiskStorageType`

## Configuring storage types

To configure storage types for VM, use the `StorageAccountType` parameter in `New-ProvScheme`. Set the value of the `StorageAccountType` parameter to one of the supported storage types.

The following is an example set of the `CustomProperties` parameter in a provisioning scheme:

```
1 Set-ProvScheme -ProvisioningSchemeName catalog-name -CustomProperties '  
    <CustomProperties xmlns="http://schemas.citrix.com/2014/xd/  
    machinecreation" xmlns:xsi="http://www.w3.org/2001/XMLSchema-  
    instance">  
2 <Property xsi:type="StringProperty" Name="UseManagedDisks" Value="true"  
    />  
3 <Property xsi:type="StringProperty" Name="StorageAccountType" Value="  
    Premium_LRS" />  
4 <Property xsi:type="StringProperty" Name="LicenseType" Value="  
    Windows_Client" />
```

```
5 </CustomProperties>'
6 <!--NeedCopy-->
```

## Requirements

- If you want the Citrix Virtual Apps and Desktops service to create resource groups for each MCS catalog, the Azure service principal associated with the host connection must have permission to create and delete resource groups. If you want the Citrix Virtual Apps and Desktops service to use existing empty resource groups, the Azure service principal associated with the host connection must have Contributor permission on those empty resource groups.
- When you create a host connection using the **Create new** option, the created service principal has subscription scope contribute permissions. Alternatively, you can use the **Use existing** option to create the connection, and provide the details of an existing subscription scope service principal. If you use the **Create new** option and create the Service Principal in the Full Configuration interface, it has the needed permissions to create and delete new resource groups or provision into existing empty resource groups.
- Narrow scope service principals must be created using PowerShell. Also, when using a narrow scope service principal, you must use PowerShell or the Azure portal to create empty resource groups in the same region as your host connection for each catalog where MCS provisions VMs. For instructions, see the blog post <https://www.citrix.com/blogs/2016/11/09/azure-role-based-access-control-in-xenapp-xendesktop/>.)

If you are using a narrow scope service principal for the host connection and don't see your image resource group on the **Master Image** page of the catalog creation wizard, it is probably because the narrow scope service principal you are using doesn't have the permission `Microsoft.Resources/subscriptions/resourceGroups/read` to list the image resource group. Close the wizard, update the service principal with the permission (see the blog post for instructions), and then restart the wizard. The update in Azure can take up to 10 minutes to appear in the Full Configuration interface.

### Tip:

The Azure service principal requires contributor rights. The contributor right can either be a full (subscription) scope-wide contributor where MCS automatically creates an Azure Resource Group for a machine catalog. Or, the contributor right can be narrow scope, where an empty Azure Resource Group must be created in advance and contributor rights granted accordingly.

## Configure resource groups for a machine catalog in the Full Configuration interface

The **Resource Groups** page in the catalog creation wizard allows you to choose whether to create resource groups or use existing groups. See [Create a machine catalog using an Azure Resource Manager image](#).



**What happens to resource groups when you delete a machine catalog.** If you let the Citrix Virtual Apps and Desktops service create resource groups when you create the machine catalog, and then later delete the catalog, those resource groups, and other resources in those resource groups are also deleted.

If you use existing resource groups when you create the machine catalog, and then later delete the catalog, all resources in those resource groups are deleted, but the resource groups are not deleted.

### Considerations and limitations

When you use existing resource groups, the list of available resource groups on the Resource Groups page in the catalog creation wizard does not auto-refresh. So, if you have that wizard page open and create or add permissions to resource groups in Azure, the changes are not reflected in the wizard's list. To see the latest changes, go back to the **Machine Management** page in the wizard and reselect the resources associated with the host connection. Or, close and restart the wizard. It can take up to 10 minutes for changes made in Azure to appear in the Full Configuration interface.

If your connection uses a service principal that can access empty resource groups in various regions, they appear in the available list. Be sure to choose resource groups in the same region where you're creating the machine catalog.

### Troubleshooting

- Resource groups don't appear in the list on the Resource Groups page of the catalog creation wizard.

The service principal must have appropriate permissions applied to the resource groups you want to appear in the list. See Requirements.

### More information

- [Connections and resources](#)
- [Create machine catalogs](#)
- [CTX219211](#): Set up a Microsoft Azure Active Directory account
- [CTX219243](#): Grant XenApp and XenDesktop access to your Azure subscription
- [CTX219271](#): Deploy hybrid cloud using site-to-site VPN

## Citrix Hypervisor virtualization environments

September 2, 2021

## Create a connection to Citrix Hypervisor

When you create a connection to Citrix Hypervisor (formerly XenServer), you must provide the credentials for a VM Power Admin or higher-level user.

Citrix recommends using HTTPS to secure communications with Citrix Hypervisor. To use HTTPS, you must replace the default SSL certificate installed on Citrix Hypervisor; see [CTX128656](#).

You can configure high availability if it is enabled on the Citrix Hypervisor server. Citrix recommends that you select all servers in the pool (from Edit High Availability) to allow communication with the Citrix Hypervisor server if the pool master fails.

You can select a GPU type and group, or pass through, if the Citrix Hypervisor supports vGPU. The display indicates if the selection has dedicated GPU resources.

When using local storage on one or more Citrix Hypervisor hosts for temporary data storage, make sure that each storage location in the pool has a unique name. (To change a name in XenCenter, right-click the storage and edit the name property.)

## Use IntelliCache for Citrix Hypervisor connections

Using IntelliCache, hosted VDI deployments are more cost-effective because you can use a combination of shared storage and local storage. This enhances performance and reduces network traffic. The local storage caches the master image from the shared storage, which reduces the amount of reads on the shared storage. For shared desktops, writes to the differencing disks are written to local storage on the host and not to shared storage.

- Shared storage must be NFS when using IntelliCache.
- Citrix recommends that you use a high performance local storage device to ensure the fastest possible data transfer.

To use IntelliCache, you must enable it in both this product and Citrix Hypervisor.

- When installing Citrix Hypervisor, select **Enable thin provisioning (Optimized storage for Citrix Virtual Desktops)**. Citrix does not support mixed pools of servers that have IntelliCache enabled and servers that do not. For more information, see the Citrix Hypervisor documentation.
- In Citrix Virtual Apps and Desktops, IntelliCache is disabled by default. You can change the setting only when creating a Citrix Hypervisor connection; you cannot disable IntelliCache later.

When you add a Citrix Hypervisor connection:

- Select **Shared** as the storage type.
- Select the **Use IntelliCache** check box.

## Create a machine catalog using a Citrix Hypervisor connection

GPU-capable machines require a dedicated master image. Those VMs require video card drivers that support GPUs. Configure GPU-capable machines to allow the VM to operate with software that uses the GPU for operations.

1. In XenCenter, create a VM with standard VGA, networks, and vCPU.
2. Update the VM configuration to enable GPU use (either Passthrough or vGPU).
3. Install a supported operating system and enable RDP.
4. Install Citrix VM Tools and NVIDIA drivers.
5. Turn off the Virtual Network Computing (VNC) Admin Console to optimize performance, and then restart the VM.
6. You are prompted to use RDP. Using RDP, install the VDA and then restart the VM.
7. Optionally, create a snapshot for the VM as a baseline template for other GPU master images.
8. Using RDP, install customer-specific applications that are configured in XenCenter and use GPU capabilities.

### More information

- [Connections and resources](#)
- [Create machine catalogs](#)

## Microsoft System Center Virtual Machine Manager virtualization environments

July 7, 2021

Follow this guidance if you use Hyper-V with Microsoft System Center Virtual Machine Manager (VMM) to provide virtual machines.

See [System requirements](#) for a list of supported VMM versions.

You can use Machine Creation Services or Citrix Provisioning (formerly Provisioning Services) to provision:

- Generation 1 Desktop or Server OS VMs
- Generation 2 Windows Server 2012 R2, Windows Server 2016, and Windows 10 VMs (with or without Secure Boot)

### Install and configure a hypervisor

Install Microsoft Hyper-V server and VMM on your servers.

Verify the following account information:

In Manage > Full Configuration, the account you specify when creating a connection must be a VMM administrator or VMM delegated administrator for the relevant Hyper-V machines. If this account has only the delegated administrator role in VMM, the storage data is not listed in the Full Configuration interface during the connection creation process.

Your user account must also be a member of the administrators local security group on each Hyper-V server to support VM lifecycle management (such as VM creation, update, and deletion).

### **Install the VMM console**

Install a System Center Virtual Machine Manager console on each server that contains a Citrix Cloud Connector.

The console version must match the management server version. Although an earlier console can connect to the management server, provisioning VDAs fails if the versions differ.

### **Create a master VM**

- Install a VDA on the master VM, and select the option to optimize the desktop. This improves performance.
- Take a snapshot of the master VM to use as a backup.
- Create virtual desktops.

### **Create a connection**

If you used MCS to provision VMs, in the connection creation wizard:

- Enter the address as a fully qualified domain name of the host server.
- Enter credentials for the administrator account you set up earlier. This account must have permission to create new VMs.
- In the Host Details dialog box, select the cluster or standalone host to use when creating VMs.  
**Important:** Browse for a select a cluster or standalone host even if you are using a single Hyper-V host deployment.

### **MCS on SMB 3 file shares**

For machine catalogs created with MCS on SMB 3 file shares for VM storage, credentials must meet the following requirements to ensure that calls from the Citrix Hypervisor Communications Library (HCL) connect successfully to SMB storage.

- VMM user credentials must include full read write access to the SMB storage.

- Storage virtual disk operations during VM lifecycle events are performed through the Hyper-V server using the VMM user credentials.

When using VMM 2012 SP1 with Hyper-V on Windows Server 2012: When using SMB as storage, enable the Authentication Credential Security Support Provider (CredSSP) from the Cloud Connector to individual Hyper-V machines. For more information, see [CTX 137465](#).

Using a standard PowerShell V3 remote session, the HCL in the Cloud Connector uses CredSSP to open a connection to the Hyper-V machine. This feature passes Kerberos-encrypted user credentials to the Hyper-V machine, and the PowerShell commands in the session on the remote Hyper-V machine run with the credentials provided (in this case, those of the VMM user), so that communication commands to storage work correctly.

The following tasks use PowerShell scripts that originate in the HCL. The scripts are then sent to the Hyper-V machine to act on the SMB 3.0 storage.

**Consolidate Master Image:** An image creates a new MCS provisioning scheme (machine catalog). It clones and flattens the master VM ready for creating new VMs from the new disk created (and removes dependency on the original master VM).

ConvertVirtualHardDisk on the root\virtualization\v2 namespace

Example:

```
$ims = Get-WmiObject -class $class -namespace "root\virtualization\v2";  
$result = $ims.ConvertVirtualHardDisk($diskName, $vhdaText)  
$result
```

**Create difference disk:** Creates a difference disk from the image generated by consolidating the image. The difference disk is then attached to a new VM.

CreateVirtualHardDisk on the root\virtualization\v2 namespace

Example:

```
$ims = Get-WmiObject -class $class -namespace "root\virtualization\v2";  
$result = $ims.CreateVirtualHardDisk($vhdaText);  
$result
```

**Upload identity disks:** The HCL cannot directly upload the identity disk to SMB storage. Therefore, the Hyper-V machine must upload and copy the identity disk to the storage. Because the Hyper-V machine cannot read the disk from the Cloud Connector, the HCL must first copy the identity disk through the Hyper-V machine as follows.

1. The HCL uploads the Identity to the Hyper-V machine through the administrator share.
2. The Hyper-V machine copies the disk to the SMB storage through a PowerShell script running in the PowerShell remote session. A folder is created on the Hyper-V machine and the permissions on that folder are locked for the VMM user only (through the remote PowerShell connection).

3. The HCL deletes the file from the administrator share.
4. When the HCL completes the identity disk upload to the Hyper-V machine, the remote PowerShell session copies the identity disks to SMB storage and then deletes it from the Hyper-V machine.

The identity disk folder is recreated if it is deleted so that it is available for reuse.

**Download identity disks:** As with uploads, the identity disks pass through the Hyper-V machine to the HCL. The following process creates a folder that has only VMM user permissions on the Hyper-V server if it does not exist.

1. The HyperV machine copies the disk from the SMB storage to local Hyper-V storage through a PowerShell script running in the PowerShell V3 remote session.
2. HCL reads the disk from the Hyper-V machine's administrator share into memory.
3. HCL deletes the file from the administrator share.

### More information

- [Connections and resources](#)
- [Create machine catalogs](#)

## VMware virtualization environments

July 14, 2021

Follow this guidance if you use VMware to provide virtual machines.

Install vCenter Server and the appropriate management tools. (No support is provided for vSphere vCenter Linked Mode operation.)

If you plan to use Machine Creation Services (MCS), do not disable the Datastore Browser feature in vCenter Server (described in [https://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=2101567](https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2101567)). If you disable this feature, MCS does not work correctly.

### Required privileges

Create a VMware user account and one or more VMware roles with a set or all of the privileges listed below. Base the roles' creation on the specific level of granularly required over the user's permissions to request the various Citrix Virtual Apps or Citrix Virtual Desktops operations at any time. To grant the user specific permissions at any point, associate them with the respective role, at the data center level at a minimum.

The following tables show the mappings between Citrix Virtual Apps and Desktops operations and the minimum required VMware privileges.

### Add connections and resources

SDK	User interface
System.Anonymous, System.Read, and System.View	Added automatically. Can use the built-in read-only role.

### Provision machines (Machine Creation Services)

SDK	User interface
Datastore.AllocateSpace	Datastore > Allocate space
Datastore.Browse	Datastore > Browse datastore
Datastore.FileManagement	Datastore > Low level file operations
Network.Assign	Network > Assign network
Resource.AssignVMToPool	Resource > Assign virtual machine to resource pool
VirtualMachine.Config.AddExistingDisk	Virtual machine > Configuration > Add existing disk
VirtualMachine.Config.AddNewDisk	Virtual machine > Configuration > Add new disk
VirtualMachine.Config.AdvancedConfig	Virtual machine > Configuration > Advanced
VirtualMachine.Config.RemoveDisk	Virtual machine > Configuration > Remove disk
VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power Off
VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power On
VirtualMachine.Inventory.CreateFromExisting	Virtual machine > Inventory > Create from existing
VirtualMachine.Inventory.Create	Virtual machine > Inventory > Create new
VirtualMachine.Inventory.Delete	Virtual machine > Inventory > Remove
VirtualMachine.Provisioning.Clone	Virtual machine > Provisioning > Clone virtual machine

SDK	User interface
VirtualMachine.State.CreateSnapshot	vSphere 5.0, Update 2, vSphere 5.1, Update 1, and vSphere 6.x, Update 1: Virtual machine > State > Create snapshot; vSphere 5.5: Virtual machine > Snapshot management > Create snapshot

If you want the VMs you create to be tagged, add the following permissions for the user account.

To ensure that you use a clean base image for creating VMs, tag VMs created with Machine Creation Services to exclude them from the list of VMs available to use as base images.

SDK	User interface
Global.ManageCustomFields	Global > Manage custom attributes
Global.SetCustomField	Global > Set custom attribute

### Provision machines (Citrix Provisioning)

All privileges from **Provision machines (Machine Creation Services)** and the following.

SDK	User interface
VirtualMachine.Config.AddRemoveDevice	Virtual machine > Configuration > Add or remove device
VirtualMachine.Config.CPUCount	Virtual machine > Configuration > Change CPU Count
VirtualMachine.Config.Memory	Virtual machine > Configuration > Memory
VirtualMachine.Config.Settings	Virtual machine > Configuration > Settings
VirtualMachine.Provisioning.CloneTemplate	Virtual machine > Provisioning > Clone template
VirtualMachine.Provisioning.DeployTemplate	Virtual machine > Provisioning > Deploy template

### Power management



SDK	User interface
VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power Off
VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power On
VirtualMachine.Interact.Reset	Virtual machine > Interaction > Reset
VirtualMachine.Interact.Suspend	Virtual machine > Interaction > Suspend

### Image update and rollback

SDK	User interface
Datastore.AllocateSpace	Datastore > Allocate space
Datastore.Browse	Datastore > Browse datastore
Datastore.FileManagement	Datastore > Low level file operations
Network.Assign	Network > Assign network
Resource.AssignVMToPool	Resource > Assign virtual machine to resource pool
VirtualMachine.Config.AddExistingDisk	Virtual machine > Configuration > Add existing disk
VirtualMachine.Config.AddNewDisk	Virtual machine > Configuration > Add new disk
VirtualMachine.Config.AdvancedConfig	Virtual machine > Configuration > Advanced
VirtualMachine.Config.RemoveDisk	Virtual machine > Configuration > Remove disk
VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power Off
VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power On
VirtualMachine.Interact.Reset	Virtual machine > Interaction > Reset
VirtualMachine.Inventory.CreateFromExisting	Virtual machine > Inventory > Create from existing
VirtualMachine.Inventory.Create	Virtual machine > Inventory > Create new
VirtualMachine.Inventory.Delete	Virtual machine > Inventory > Remove
VirtualMachine.Provisioning.Clone	Virtual machine > Provisioning > Clone virtual machine

### Delete provisioned machines

SDK	User interface
Datastore.Browse	Datastore > Browse datastore
Datastore.FileManagement	Datastore > Low level file operations
VirtualMachine.Config.RemoveDisk	Virtual machine > Configuration > Remove disk
VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power Off
VirtualMachine.Inventory.Delete	Virtual machine > Inventory > Remove

## Securing connections to the VMware environment

Using [HTTPS/SSL](#) connections to vCenter requires that the connection is trusted by Citrix Virtual Apps and Desktops service.

There are two options:

- Each cloud connector trusts the vCenter certificate, and services on the connector reuses this trust. This trust can be from a:
  - vCenter certificate, issued by the Certificate Authority and trusted by windows, resulting in established trust between Windows and vCenter.
  - vCenter certificate installed on Windows, resulting in established trust between Windows and vCenter.
- Alternatively the Citrix Virtual Apps and Desktops database has the SSL thumbprint installed. This thumbprint is used by Citrix Virtual Apps and Desktops service on each cloud connector to trust connections to vCenter.

## Obtain and import a certificate

To protect vSphere communications, Citrix recommends that you use HTTPS rather than HTTP. HTTPS requires digital certificates. Citrix recommends you use a digital certificate issued from a certificate authority in accordance with your organization's security policy.

If you are unable to use a digital certificate issued from a certificate authority, and your organization's security policy permits it, you can use the VMware-installed self-signed certificate. Add the VMware vCenter certificate to each Cloud Connector.

1. Add the fully qualified domain name (FQDN) of the computer running vCenter Server to the hosts file on that server, located at %SystemRoot%/WINDOWS/system32/Drivers/etc/. This step is required only if the FQDN of the computer running vCenter Server is not already present in the domain name system.

2. Obtain the vCenter certificate using any of the following three methods:

**From the vCenter server:**

- a) Copy the file rui.crt from the vCenter server to a location accessible on your Cloud Connectors.
- b) On the Cloud Connector, navigate to the location of the exported certificate and open the rui.crt file.

**Download the certificate using a web browser:** If you are using Internet Explorer, depending on your user account, you may must right-click on Internet Explorer and choose **Run as Administrator** to download or install the certificate.

- a) Open your web browser and make a secure web connection to the vCenter server (for example <https://server1.domain1.com>).
- b) Accept the security warnings.
- c) Click the address bar displaying the certificate error.
- d) View the certificate and click the Details tab.
- e) Select **Copy to file and export in .CER format**, providing a name when prompted to do so.
- f) Save the exported certificate.
- g) Navigate to the location of the exported certificate and open the .CER file.

**Import directly from Internet Explorer running as an administrator:**

- a) Open your web browser and make a secure web connection to the vCenter server (for example <https://server1.domain1.com>).
- b) Accept the security warnings.
- c) Click the address bar displaying the certificate error.
- d) View the certificate.

3. Import the certificate into the certificate store on each Cloud Connector.

- a) Click **Install certificate**, select **Local Machine**, and then click **Next**.
- b) Select **Place all certificates in the following store**, and then click **Browse**. On a later supported version: Select **Trusted People** and then click **OK**. Click **Next** and then click **Finish**.

**Important:**

If you change the name of the vSphere server after installation, you must generate a new self-signed certificate on that server before importing the new certificate.

### VMware SSL thumbprint

The VMware SSL thumbprint feature addresses a frequently reported error when creating a host connection to a VMware vSphere hypervisor. Previously, administrators had to manually create a trust

relationship between the Citrix-managed Delivery Controllers in the Site and the hypervisor's certificate before creating a connection. The VMware SSL thumbprint feature removes that manual requirement: the untrusted certificate's thumbprint is stored on the Site database so that the hypervisor can be continuously identified as trusted by Citrix Virtual Apps or Citrix Virtual Desktops, even if not by the Controllers.

When creating a vSphere host connection, a dialog box allows you to view the certificate of the machine you are connecting to. You can then choose whether to trust it.

The VMware SSL thumbprint can be updated at a later time using PowerShell SDK `Set-Item -LiteralPath "<FullPath_to_connection>"-username $cred.username - Securepassword $cred.password -SslThumbprint "<New ThumbPrint>"-hypervisorAddress <vcenter URL>`.

**Tip:**

The certificate thumbprint has to be written in capital letters.

## Create a master VM

Use a master VM to provide user desktops and applications in a machine catalog. On your hypervisor:

1. Install a VDA on the master VM, selecting the option to optimize the desktop, which improves performance.
2. Take a snapshot of the master VM to use as a back-up.

## Create a connection

In the connection creation wizard:

- Select the VMware connection type.
- Specify the address of the access point for the vCenter SDK.
- Specify the credentials for a VMware user account you set up earlier that has permissions to create new VMs. Specify the user name in the form domain/username.

## More information

- [Connections and resources](#)
- [Create machine catalogs](#)

## Amazon Web Services virtualization environments

August 31, 2021

This article walks you through setting up your Amazon Web Services (AWS) account as a resource location you can use with the Citrix Virtual Apps and Desktops service. The resource location includes a basic set of components, ideal for a proof-of-concept or other deployment that does not require resources spread over multiple availability zones. After you complete these tasks, you can install VDAs, provision machines, create machine catalogs, and create Delivery Groups.

When you complete the tasks in this article, your resource location includes the following components:

- A virtual private cloud (VPC) with public and private subnets inside a single availability zone.
- An instance that runs as both an Active Directory domain controller and DNS server, located in the private subnet of the VPC.
- Two domain-joined instances on which the Citrix Cloud Connector is installed, located in the private subnet of the VPC.
- An instance that acts as a bastion host, located in the public subnet of your VPC. This instance is used to initiate RDP connections to the instances in the private subnet for administration purposes. After you finish setting up your resource location, you can shut down this instance so it is no longer readily accessible. When you must manage other instances in the private subnet, such as VDA instances, you can restart the bastion host instance.

### Task overview

**Set up a virtual private cloud (VPC) with public and private subnets.** When you complete this task, AWS deploys a NAT instance with an Elastic IP address in the public subnet, which enables instances in the private subnet to access the Internet. Instances in the public subnet are accessible to inbound public traffic while instances in the private subnet are not.

**Configure security groups.** Security groups act as virtual firewalls that control traffic for the instances in your VPC. You add rules to your security groups that allow instances in your public subnet to communicate with instances in your private subnet. You will also associate these security groups with each instance in your VPC.

**Create a DHCP options set.** With an Amazon VPC, DHCP and DNS services are provided by default, which affects how you configure DNS on your Active Directory domain controller. Amazon's DHCP cannot be disabled and Amazon's DNS can be used only for public DNS resolution, not Active Directory name resolution. To specify the domain and name servers that should be handed to instances via DHCP, you create a DHCP options set. The set assigns the Active Directory domain suffix and specifies the DNS server for all instances in your VPC. To ensure Host (A) and Reverse Lookup (PTR) records are

automatically registered when instances join the domain, you configure the network adapter properties for each instance you add to the private subnet.

**Add a bastion host, domain controller, and Citrix Cloud Connectors to the VPC.** Through the bastion host, you can log on to instances in the private subnet to set up the domain, join instances to the domain, and install the Citrix Cloud Connector.

### Task 1: Set up the VPC

1. From the AWS management console, select **VPC**.
2. From the VPC Dashboard, select **Start VPC Wizard**.
3. Select **VPC with Public and Private Subnets** and then select **Select**.
4. Enter a VPC name and change the IP CIDRE block and Public and Private subnet IP ranges, if necessary.
5. If a NAT gateway is selected, select **Use a NAT Instance instead**.
6. For the NAT instance, specify the instance type and the key pair you want to use. The key pair enables you to securely connect to the instance later.
7. In Enable DNS host names, leave **Yes** selected.
8. Select **Create VPC**. AWS creates the public and private subnets, Internet gateway, route tables, and default security group. Also, a NAT instance is created and assigned an Elastic IP address.

#### Note:

Changing the name of an AWS Virtual Private Cloud (VPC) in the AWS console breaks the existing hosting unit in Citrix Cloud. When the hosting unit is broken, you cannot create catalogs or add machines to existing catalogs. From Known Issue: PMCS-7701

### Task 2: Configure security groups

This task creates and configures the following security groups for your VPC:

- A security group for the NAT instance.
- A public security group, with which instances in your Public subnet will be associated.
- A private security group, with which instances in your Private subnet will be associated.

To create the security groups

1. From the VPC Dashboard, select **Security Groups**.
2. Create a security group for the NAT instance. Select **Create Security Group** and enter a name tag and description for the group. In VPC, select the VPC you created earlier. Select **Yes, Create**.
3. Repeat Step 2 to create a public security group and a private security group.

**Configure the NAT security group**

1. From the security group list, select the NAT security group.
2. Select the **Inbound Rules** tab and select **Edit** to create the following rules:

Type	Source
ALL Traffic	Select the Private security group.
22 (SSH)	0.0.0.0/0

3. When finished, select **Save**.

**Configure the Public security group**

1. From the security group list, select the Public security group.
2. Select the **Inbound Rules** tab and select Edit to create the following rules:

Type	Source
ALL Traffic	Select the Private security group.
ALL Traffic	Select the Public security group.
ICMP	0.0.0.0/0
22 (SSH)	0.0.0.0/0
80 (HTTP)	0.0.0.0/0
443 (HTTPS)	0.0.0.0/0
1494 (ICA/HDX)	0.0.0.0/0
2598 (Session Reliability)	0.0.0.0/0
3389 (RDP)	0.0.0.0/0

3. When finished, select **Save**.
4. Select the **Outbound Rules** tab and select **Edit** to create the following rules:

Type	Destination
ALL Traffic	Select the Private security group.
ALL Traffic	0.0.0.0/0

Type	Destination
ICMP	0.0.0.0/0

- When finished, select **Save**.

### Configure the private security group

- From the security group list, select the Private security group.
- Select the **Inbound Rules** tab and select **Edit** to create the following rules:

Type	Source
ALL Traffic	Select the NAT security group.
ALL Traffic	Select the Private security group.
ALL Traffic	Select the Public security group.
ICMP	Select the Public security group.
TCP 53 (DNS)	Select the Public security group.
UDP 53 (DNS)	Select the Public security group.
80 (HTTP)	Select the Public security group.
TCP 135	Select the Public security group.
TCP 389	Select the Public security group.
UDP 389	Select the Public security group.
443 (HTTPS)	Select the Public security group.
TCP 1494 (ICA/HDX)	Select the Public security group.
TCP 2598 (Session Reliability)	Select the Public security group.
3389 (RDP)	Select the Public security group.
TCP 49152–65535	Select the Public security group.

- When finished, select **Save**.
- Select the **Outbound Rules** tab and select **Edit** to create the following rules:



---

Type	Destination
ALL Traffic	Select the Private security group.
ALL Traffic	0.0.0.0/0
ICMP	0.0.0.0/0
UDP 53 (DNS)	0.0.0.0/0

---

5. When finished, select **Save**.

### Task 3: Associate the NAT instance with the NAT security group

1. From the AWS management console, select **EC2**.
2. From the EC2 Dashboard, select **Instances**.
3. Select the NAT instance and then select **Actions > Networking > Change Security Groups**.
4. Clear the default security group check box.
5. Select the NAT security group you created earlier and then select **Assign Security Groups**.

### Task 4: Launch instances

The following steps create four EC2 instances and decrypt the default Administrator password that Amazon generates.

1. From the AWS management console, select **EC2**.
2. From the EC2 Dashboard, select **Launch Instance**.
3. Select a Windows Server machine image and instance type.
4. On the Configure Instance Details page, enter a name for the instance and select the VPC you set up earlier.
5. In **Subnet**, make the following selections for each instance:
  - Bastion host: Select the Public subnet.
  - Domain controller and Connectors: Select the Private subnet.
6. In **Auto-assign Public IP address**, make the following selections for each instance:
  - Bastion host: Select **Enable**.
  - Domain controller and Connectors: Select **Use default setting** or **Disable**.
7. In **Network Interfaces**, enter a primary IP address within the IP range of your private subnet for the domain controller and Cloud Connector instances.
8. On the Add Storage page, modify the disk size, if necessary.

9. On the Tag Instance page, enter a friendly name for each instance.
10. On the Configure Security Groups page, select **Select an existing security group** and then make the following selections for each instance:
  - Bastion host: Select the Public security group.
  - Domain controller and Cloud Connectors: Select the Private security group.
11. Review your selections and then select **Launch**.
12. Create a new key pair or select an existing one. If you create a new key pair, download your private key (.pem) file and keep it in safe place. You must supply your private key when you acquire the default Administrator password for the instance.
13. Select **Launch Instances**. select **View Instances** to display a list of your instances. Wait until the newly launched instance has passed all status checks before accessing it.
14. Acquire the default Administrator password for each instance:
  - a) From the instance list, select the instance and then select **Connect**.
  - b) Select **Get Password** and supply your private key (.pem) file when prompted.
  - c) Select **Decrypt Password**. AWS displays the default password.
15. Repeat Steps 2–14 until you have created four instances: a bastion host instance in your public subnet and three instances in your private subnet that for use as a domain controller and two Cloud Connectors.

### **Task 5: Create a DHCP options set**

1. From the VPC Dashboard, select **DHCP Options Sets**.
2. Enter the following information:
  - Name tag: Enter a friendly name for the set.
  - Domain name: Enter the fully qualified domain name you use when you configure the domain controller instance.
  - Domain name servers: Enter the private IP address you assigned to the domain controller instance and the string **AmazonProvidedDNS**, separated by commas.
  - NTP servers: Leave this field blank.
  - NetBIOS name servers: Enter the private IP address of the domain controller instance.
  - NetBIOS node type: Enter **2**.
3. Select **Yes, Create**.
4. Associate the new set with your VPC:
  - a) From the VPC Dashboard, select **Your VPCs** and then select the VPC you set up earlier.
  - b) Select **Actions > Edit DHCP Options Set**.

- c) When prompted, select the new set you created and then select **Save**.

### Task 6: Configure the instances

1. Using an RDP client, connect to the public IP address of the bastion host instance. When prompted, enter the credentials for the Administrator account.
2. From the bastion host instance, launch Remote Desktop Connection and connect to the private IP address of the instance you want to configure. When prompted, enter the Administrator credentials for the instance.
3. For all instances in the private subnet, configure the DNS settings:
  - a) Select **Start > Control Panel > Network and Internet > Network and Sharing Center > Change adapter settings**. Double-click the network connection displayed.
  - b) Select **Properties**, select **Internet Protocol Version 4 (TCP/IPv4)**, and then select **Properties**.
  - c) Select **Advanced** and then select the **DNS** tab. Ensure that the following settings are enabled and select **OK**:
    - Register this connection's addresses in DNS
    - Use this connection's DNS suffix in DNS registration
4. To configure the domain controller:
  - a) Using Server Manager, add the Active Directory Domain Services role with all default features.
  - b) Promote the instance to a domain controller. During promotion, enable DNS and use the domain name you specified when you created the DHCP options set. Restart the instance when prompted.
5. To configure the first Cloud Connector:
  - a) Join the instance to the domain and restart when prompted. From the bastion host instance, reconnect to the instance using RDP.
  - b) Sign in to Citrix Cloud. Select **Resource Locations** from the upper left menu.
  - c) Download the Cloud Connector.
  - d) When prompted, run the `cwconnector.exe` file and supply your Citrix Cloud credentials. Follow the wizard.
  - e) When finished, select **Refresh** to display the Resource Locations page. When the Cloud Connector is registered, the instance appears on the page.
6. Repeat Step 5 to configure the second Cloud Connector.

## Create a connection

When you create a connection from the Full Configuration interface:

- You must provide the API key and secret key values. You can export the key file containing those values from AWS and then import them. You must also provide the region, availability zone, VPC name, subnet addresses, domain name, security group names, and credentials.
- The credentials file for the root AWS account (retrieved from the AWS console) is not formatted the same as credentials files downloaded for standard AWS users. Therefore, Citrix Virtual Apps and Desktops management cannot use the file to populate the API key and secret key fields. Ensure that you are using AWS Identity Access Management (IAM) credentials files.

### Note:

After you create a connection, attempts to update the API key and secret key might fail. To resolve the issue, check your proxy server or firewall restrictions and ensure that the following address is contactable: [https://\\*.amazonaws.com](https://*.amazonaws.com).

## Defining IAM permissions

Use the information in this section to define IAM permissions for the Citrix Virtual Apps and Desktops service on AWS. Amazon's IAM service permits accounts having multiple users, which can be further organized into groups. These users can possess different permissions to control their ability to perform operations associated with the account. For more information about IAM permissions, see [IAM JSON policy reference](#).

To apply IAM permissions policy to a new group of users:

1. Log into the AWS management console and select the **IAM service** from the drop-down list.
2. Select **Create a New Group of Users**.
3. Type a name for the new user group and select **Continue**.
4. On the **Permissions** page, choose **Custom Policy**. Select **Select**.
5. Type a name for the **Permissions policy**.
6. In the **Policy Document** section, enter relevant permissions.

After entering the policy information, select **Continue** to complete the group of users. Users in the group are granted permissions to perform only those actions that are required for the Citrix Virtual Apps and Desktops service.

### Important:

Use the policy text provided in the example above to list the actions that a Citrix Virtual Apps and Desktops service uses to perform actions within an AWS account without restricting those

actions to specific resources. Citrix recommends that you use the example for testing purposes. For production environments, you might choose to add further restrictions on resources.

## AWS tenancy

AWS provides the following tenancy options: shared tenancy (the default type) and dedicated tenancy. Shared tenancy means that multiple Amazon EC2 instances from different customers might reside on the same piece of physical hardware. Dedicated tenancy means that your EC2 instances run only on hardware with other instances that you have deployed. Other customers do not use the same piece of hardware.

When you use MCS to create a catalog to provision machines in AWS, the **Machine Catalog Setup > Security** page presents the following options:

- **Use shared hardware.** This setting is suitable for most deployments. Multiple customers share pieces of hardware even though they do not interact with each other. Using shared hardware is the least expensive option for running your Amazon EC2 instances.
- **Use dedicated host.** An Amazon EC2 dedicated host is a physical server with EC2 instance capacity that is fully dedicated, letting you use existing per-socket or per-VM software licenses. Dedicated hosts have preset utilization based on instance type. For example, a single allocated dedicated host of C4 Large instance types is limited to running 16 instances. See the [AWS site](#) for more information.

The requirements for provisioning to AWS hosts include:

- An imported BYOL (bring your own license) image (AMI). With dedicated hosts, use and manage your existing licenses.
- An allocation of dedicated hosts with sufficient utilization to satisfy provisioning requests.
- Enabling **auto-placement**.

This setting is suitable for deployments with licensing restrictions or security requirements that need your use of a dedicated host. With a dedicated host, you own an entire physical host and are billed on an hourly basis. Owning that host lets you spin up as many EC2 instances as that host permits, without more charges.

Alternatively, you can provision AWS dedicated hosts through PowerShell. To do that, use the `New-ProvScheme` cmdlet with the parameter `TenancyType` set to `Host`. See the [Citrix Developer Documentation](#) for more information.

- **Use dedicated instance.** This setting is more suitable for deployments with specific security or compliance requirements. With a dedicated instance, you still enjoy the benefits of having a host separate from other AWS customers but you do not pay for the entire host. You do not need to worry about the capacity of the host but you are charged at a higher rate for the instances.

## AWS instance property capturing

When you create a catalog to provision machines using Machine Creation Services (MCS) in AWS, you select an AMI to represent the golden image of that catalog. From that AMI, MCS uses a snapshot of the disk. In previous releases, if you wanted roles and/or tags on your machines you would use the AWS console to set them individually.

To improve this process, MCS reads properties from the instance from which the AMI was taken and applies the Identity Access Management (IAM) role and tags of the machine to the machines provisioned for a given catalog. When using this optional feature, the catalog creation process finds the selected AMI source instance, reading a limited set of properties. These properties are then stored in an AWS Launch Template, which is used to provision machines for that catalog. Any machine in the catalog inherits the captured instance properties.

Captured properties include:

- IAM roles – applied to provisioned instances
- Tags - applied to provisioned instances, their disk, and NICs. These tags are applied to transient Citrix resources, including: S3 bucket and objects, volume and worker resources, and AMIs, snapshots, and launch templates.

### Tip:

The tagging of transient Citrix resources is optional and is configurable using the custom property `AwsOperationalResourcesTagging`.

## Capturing the AWS instance property

You can use this feature by specifying a custom property, `AwsCaptureInstanceProperties`, when creating a provisioning scheme for an AWS hosting connection:

```
New-ProvScheme -CustomProperties "AwsCaptureInstanceProperties,true" ...<
standard provscheme parameters
```

To use this feature, you must specify a broader set of permissions for the AWS service key. These permissions include:

- ec2:AssociateIamInstanceProfile
- ec2:CreateLaunchTemplate
- ec2>DeleteLaunchTemplate
- ec2>DeleteTags
- ec2:DisassociateIamInstanceProfile
- ec2:DescribeIamInstanceProfileAssociations
- ec2:DescribeLaunchTemplates
- ec2:DescribeLaunchTemplateVersions

- ec2:DescribeSnapshots
- ec2:DescribeTags
- iam:PassRole
- s3:PutBucketTagging
- s3:PutObjectTagging

Refer to the [Citrix Developer Documentation](#) for more information.

## Applying AWS instance properties and tagging operational resources in the Full Configuration interface

When creating a catalog to provision machines in AWS by using MCS, you can control whether to apply the IAM role and tag properties to those machines. You can also control whether to apply machine tags to operational resources. You have the following two options:

**Machine Catalog Setup**

- Introduction
- Operating System
- Machine Managem...
- 4 Machine Template**
- 5 Security
- 6 Virtual Machines
- 7 Network Cards
- 8 Computer Accounts
- 9 Domain Credentials
- 10 Scopes
- 11 Summary

### Machine Template

Select the machine template that the virtual machines will be based upon.

Name	Description ↓
<input checked="" type="radio"/> AlertLogic Scan Appliance 1568213815 (ami-...	-
<input type="radio"/> Alertlogic TMC - P13 (ami-5934df24)	-
<input type="radio"/> Alex-VDA (ami-91d10efa)	-
<input type="radio"/> aliceAMI (ami-0ef03f44d11f8b3dc)	-
<input type="radio"/> Citrix.XenDesktop.VolumeWorker.500759a5-...	-
<input type="radio"/> citrix-ebs-1586205002 (ami-0582405326cc...	-
<input type="radio"/> citrix-ebs-1591292792 (ami-09e27db17c0741...	-

Select the minimum functional level for this catalog: ?

7.6 (or newer) ↓

Machines will require the selected VDA version (or newer) in order to register in Delivery Groups that reference this machine catalog. [Learn more](#)

Apply machine template properties to virtual machines ?

Apply machine tags to operational resources ?

Back Next Cancel

- **Apply machine template properties to virtual machines**
  - Controls whether to apply the IAM role and tag properties associated with the selected machine template to virtual machines in this catalog.

- **Apply machine tags to operational resources**

- Controls whether to apply machine tags to every item created in your AWS environment that facilitates provisioning of machines. Operational resources are created as byproducts of catalog creation. They include both temporary and persistent resources, such as preparation VM instance and AMI.

## **AWS operational resource tagging**

An Amazon Machine Image (AMI) represents a type of virtual appliance used to create a virtual machine within the Amazon Cloud environment, commonly referred to as EC2. You use an AMI to deploy services that use the EC2 environment. When you create a catalog to provision machines using MCS for AWS, you select the AMI to act as the golden image for that catalog.

### **Important:**

Creating catalogs by capturing an instance property and launch template is required for using operational resource tagging. For details, see the preceding section [AWS instance property capturing](#).

To create an AWS catalog, you must first create an AMI for the instance you want to be the golden image. MCS reads the tags from that instance and incorporates them into the launch template. The launch template tags are then applied to all Citrix resources created in your AWS environment, including:

- Virtual Machines
- VM disks
- VM network interfaces
- S3 buckets
- S3 objects
- Launch templates
- AMIs

## **Tagging an operational resource**

To use PowerShell to tag resources:

1. Open a PowerShell window from the DDC host.
2. Run the command `asnp citrix` to load Citrix-specific PowerShell modules.

To tag a resource for a provisioned VM, use the new custom property `AwsOperationalResourcesTagging`. The syntax for this property is:

```
New-ProvScheme -CustomProperties "AwsCaptureInstanceProperties,true;  
AwsOperationalResourcesTagging,true" ...<standard provscheme parameters>
```

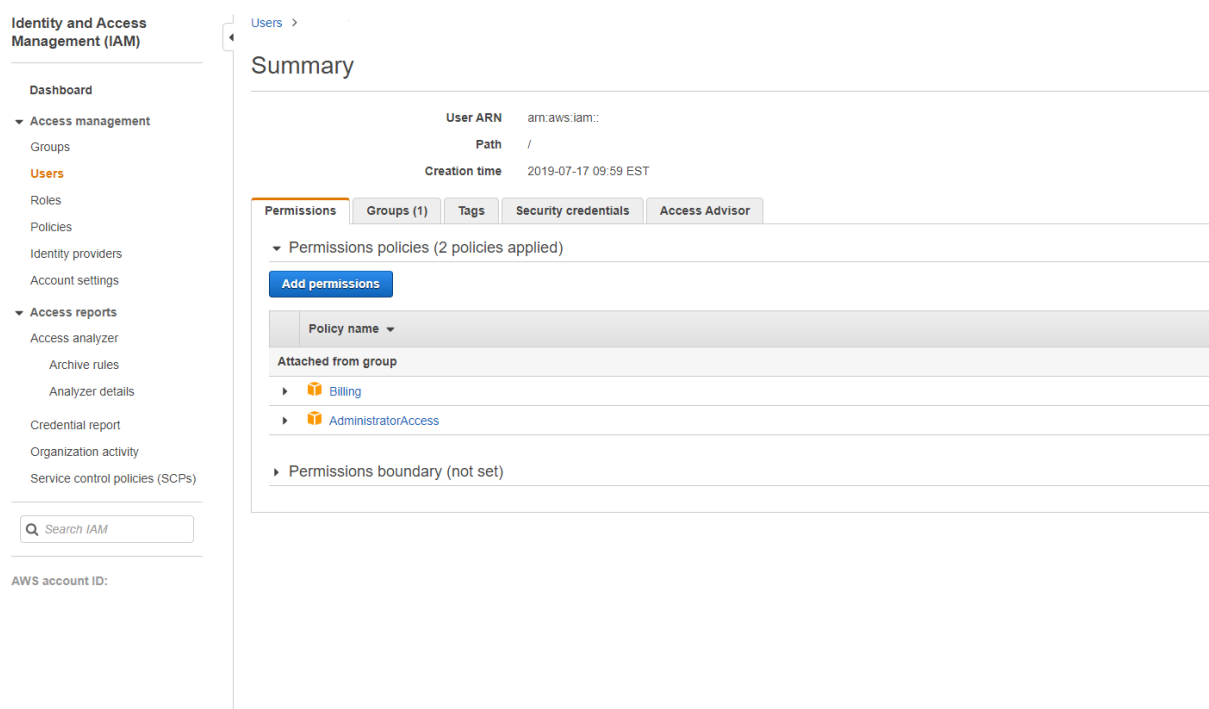


To use the `AwsOperationalResourcesTagging` custom property, ensure that the following new permissions exist for the AWS service key:

- `ec2:CreateTags`
- `ec2>DeleteTags`
- `ec2:DescribeTags`
- `s3:PutBucketTagging`
- `s3:PutObjectTagging`

Set these permissions in the **IAM** section of the AWS Management Console:

1. In the **Summary** panel, select the **Permissions** tab.
2. Select **Add permissions**.



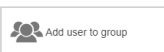


In the **Add Permissions to** screen, grant permissions:

## Citrix Virtual Apps and Desktops service

Add permissions to

### Grant permissions

Use IAM policies to grant permissions. You can assign an existing policy or create a new one.

Create policy

Filter policies	Search	Policy name	Type	Used as
<input type="checkbox"/>		 AdministratorAccess	Job function	Permissions policy (8)
<input type="checkbox"/>		 AlexaForBusinessDeviceSetup	AWS managed	None
<input type="checkbox"/>		 AlexaForBusinessFullAccess	AWS managed	None
<input type="checkbox"/>		 AlexaForBusinessGatewayExecution	AWS managed	None
<input type="checkbox"/>		 AlexaForBusinessPolyDelegatedAccessPolicy	AWS managed	None
<input type="checkbox"/>		 AlexaForBusinessReadOnlyAccess	AWS managed	None
<input type="checkbox"/>		 AmazonAPIGatewayAdministrator	AWS managed	None
<input type="checkbox"/>		 AmazonAPIGatewayInvokeFullAccess	AWS managed	None

Use the following as an example in the **JSON** tab:

### Create policy

1 2

A policy defines the AWS permissions that you can assign to a user, group, or role. You can create and edit a policy in the visual editor and using JSON. [Learn more](#)

Visual editor

JSON

Import managed policy

```
1 {
2   "Version": "2012-10-17",
3   "Statement": [
4     {
5       "Sid": "VisualEditor0",
6       "Effect": "Allow",
7       "Action": [
8         "ec2:CreateTags",
9         "ec2:DeleteTags",
10        "ec2:DescribeTags",
11        "ec2:PutObjectTagging",
12        "ec2:PutBucketTagging"
13      ],
14      "Resource": "*"
15    },
16    {
17      "Sid": "VisualEditor1",
18      "Effect": "Allow",
19      "Action": "iam:PassRole",
20      "Resource": "arn:aws:iam::*:role/*"
21    }
22  ]
23 }
```

Character count: 304 of 6,144.

Cancel

Review policy

### Tip:

The noted JSON example may not include all the permissions for your environment. See [How to Define Identity Access Management Permissions Running Citrix Virtual Apps and Desktops on Amazon Web Services](#) for more information.

## About AWS permissions

This section contains the complete list of AWS permissions.

```
1 {
2
3   "Version": "2012-10-17",
4   "Statement": [
5     {
6
7       "Action": [
8         "ec2:AttachVolume",
9         "ec2:AuthorizeSecurityGroupEgress",
10        "ec2:AuthorizeSecurityGroupIngress",
11        "ec2:CreateImage",
12        "ec2:CreateLaunchTemplate",
13        "ec2:CreateNetworkInterface",
14        "ec2:CreateSecurityGroup",
15        "ec2:CreateTags",
16        "ec2:CreateVolume",
17        "ec2>DeleteLaunchTemplate",
18        "ec2>DeleteNetworkInterface",
19        "ec2>DeleteSecurityGroup",
20        "ec2>DeleteSnapshot",
21        "ec2>DeleteTags",
22        "ec2>DeleteVolume",
23        "ec2:DeregisterImage",
24        "ec2:DescribeAccountAttributes",
25        "ec2:DescribeAvailabilityZones",
26        "ec2:DescribeImages",
27        "ec2:DescribeInstances",
28        "ec2:DescribeInstanceTypes",
29        "ec2:DescribeLaunchTemplates",
30        "ec2:DescribeNetworkInterfaces",
31        "ec2:DescribeRegions",
32        "ec2:DescribeSecurityGroups",
33        "ec2:DescribeSubnets",
34        "ec2:DescribeTags",
35        "ec2:DescribeVolumes",
36        "ec2:DescribeVpcs",
37        "ec2:DetachVolume",
38        "ec2:RebootInstances",
39        "ec2:RevokeSecurityGroupEgress",
40        "ec2:RevokeSecurityGroupIngress",
41        "ec2:RunInstances",
```

```
42         "ec2:StartInstances",
43         "ec2:StopInstances",
44         "ec2:TerminateInstances",
45         "s3:PutBucketTagging",
46         "s3:PutObjectTagging"
47     ],
48     "Effect": "Allow",
49     "Resource": "*"
50 }
51 ,
52 {
53
54     "Action": [
55         "s3:CreateBucket",
56         "s3>DeleteBucket",
57         "s3>DeleteObject",
58         "s3:GetObject",
59         "s3:PutObject"
60     ],
61     "Effect": "Allow",
62     "Resource": "arn:aws:s3:::citrix*"
63 }
64 ,
65 {
66
67     "Effect": "Allow",
68     "Action": "iam:PassRole",
69     "Resource": "arn:aws:iam::*:role/*"
70 }
71
72 ]
73 }
74
75 <!--NeedCopy-->
```

## More information

- [Connections and resources](#)
- [Create machine catalogs](#)

## Nutanix virtualization environments

January 26, 2021

Follow this guidance when using Nutanix Acropolis to provide virtual machines in your Citrix Virtual Apps or Citrix Virtual Desktops deployment. The setup process includes the following tasks:

- Install and register the Nutanix plug-in in your Citrix Virtual Apps or Citrix Virtual Desktops environment.
- Create a connection to the Nutanix Acropolis hypervisor.
- Create a machine catalog that uses a snapshot of a master image you created on the Nutanix hypervisor.

For more information, see the Nutanix Acropolis MCS plug-in Installation Guide, available at the [Nutanix Support Portal](#).

### Important:

The Nutanix virtualization AHV MCS plug-in for the Citrix Cloud Connector must be installed on all Cloud Connectors registered with the Citrix Cloud tenant. They must be registered even if they serve a resource location without the AHV.

## Install and register the Nutanix plug-in

Complete the following procedure to install and register the Nutanix plug-in on the Cloud Connectors. Use the **Manage** functions in Citrix Cloud to create a connection to the Nutanix hypervisor. Then create a machine catalog that uses a snapshot of a master image you created in the Nutanix environment.

### Tip:

Citrix recommends that you stop, then restart the Citrix Host Service, the Citrix Broker Service, and the Machine Creation Service when you install or update the Nutanix plug-in.

For information about installing the Nutanix plug-in, see the [Nutanix Documentation site](#).

## Create a connection to Nutanix

In the Add Connection and Resources wizard, select the **Nutanix** connection type on the **Connection** page, and then specify the hypervisor address and credentials, plus a name for the connection. On the **Network** page, select a network for the hosting unit.

## Create a machine catalog using a Nutanix snapshot

This information is a supplement to the guidance in [Create machine catalogs](#). It describes only the fields that are unique to Nutanix.

The snapshot you select is the template that is used to create the VMs in the catalog. Before creating the catalog, create images and snapshots in Nutanix. See the Nutanix documentation for guidance.

In the catalog creation wizard:

- The **Operating System** and **Machine Management** pages do not contain Nutanix-specific information.
- On the **Container** page, which is unique to Nutanix, select the container where the VMs' disks are placed.
- On the **Master Image** page, select the image snapshot. Acropolis snapshot names must be prefixed with "XD\_" to be used in Citrix Virtual Apps and Desktops. Use the Acropolis console to rename your snapshots, if needed. If you rename snapshots, restart the catalog creation wizard to see a refreshed list.
- On the **Virtual Machines** page, indicate the number of virtual CPUs and the number of cores per vCPU.
- The **Network Cards**, **Computer Accounts**, and **Summary** pages do not contain Nutanix-specific information.

### More information

- [Connections and resources](#)
- [Create machine catalogs](#)

## Google Cloud Platform virtualization environments

July 26, 2021

The Citrix Virtual Apps and Desktops service lets you provision and manage machines on Google Cloud Platform (GCP). This article walks you through using Machine Creation Services (MCS) to provision virtual machines in your Citrix Virtual Apps or Citrix Virtual Desktops service deployment.

### Requirements

- Citrix Cloud account. The feature described in this article is available only in Citrix Cloud.
- Citrix Virtual Apps and Desktops service subscription. For details, see [Get started](#).

- A GCP project. The project stores all compute resources associated with the machine catalog. It can be an existing project or a new one.
- Enable four APIs in your Google Cloud project. For details, see [Enable Google Cloud APIs](#).
- GCP service account. The service account authenticates to Google Cloud to enable access to the project. For details, see [Configure the Google Cloud service account](#).
- Enable Google private access. For details, see [Enable-private-google-access](#).

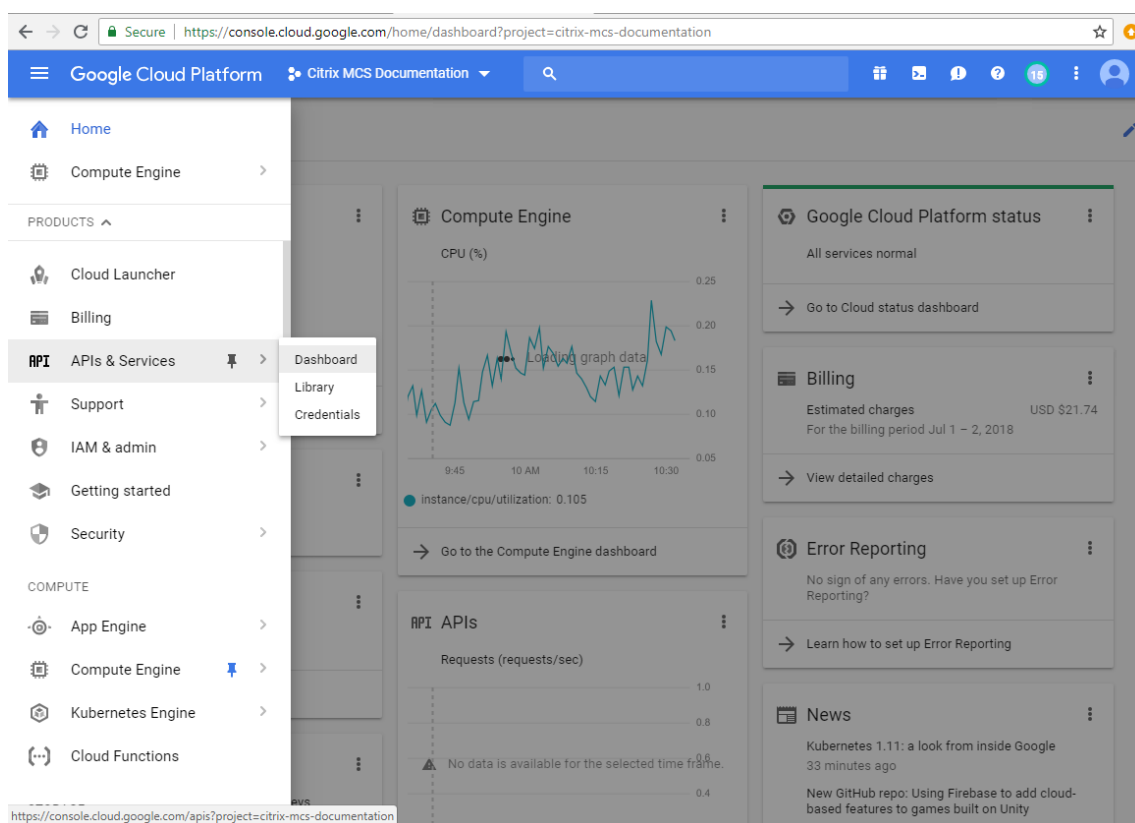
### Enable Google Cloud APIs

To use the Google Cloud functionality through the Citrix Virtual Apps and Desktops Full Configuration interface, enable these APIs in your Google Cloud project:

- Compute Engine API
- Cloud Resource Manager API
- Identity and Access Management (IAM) API
- Cloud Build API


From the GCP console, complete these steps:

1. In the upper left menu, select **APIs and Services > Dashboard**.



2. On the **Dashboard** screen, ensure that Compute Engine API is enabled. If not, follow these steps:

- a) Navigate to **APIs and Services > Library**.



- b) In the search box, type *Compute Engine*.
  - c) From the search results, select **Compute Engine API**.
  - d) On the **Compute Engine API** page, select **Enable**.
3. Enable Cloud Resource Manager API.
    - a) Navigate to **APIs and Services > Library**.
    - b) In the search box, type *Cloud Resource Manager*.
    - c) From the search results, select **Cloud Resource Manager API**.
    - d) On the **Cloud Resource Manager API** page, select **Enable**. The API's status appears.
  4. Similarly, enable **Identity and Access Management (IAM) API** and **Cloud Build API**.

## Configure the Google Cloud service account

A Google Cloud service account lets you create and manage resources inside GCP projects. A Google Cloud service account is required to provision and manage machines, as described in this article. The Google Cloud account authenticates to Citrix Cloud using a [key](#) generated by Google Cloud. Each account (personal or service) contains various roles defining the management of the project.

We recommend that you create a service account. To do so, follow these steps:



1. In the GCP console, navigate to **IAM & Admin > Service accounts**.
2. On the **Service accounts** page, select **CREATE SERVICE ACCOUNT**.
3. On the **Create service account** page, type the required information and then select **CREATE**.

When creating the service account, consider the following:

- You can select **CANCEL** to save and exit the **Service account details** page without completing the **Grant this service account access to project** and the **Grant users access to this service account** pages. You can do these optional steps later.
- If you choose to skip these optional configuration steps, the newly created service account does not display in the **IAM & Admin > IAM** page.
- To display roles associated with a service account, add the roles without skipping the optional steps. This process ensures that roles appear for the configured service account.

When creating a service account, there is an option to create a key for the account. You need this key when creating a connection in the Citrix service. The key is contained in a credential file (.json). The file is automatically downloaded and saved to the **Downloads** folder after you create the key. When you create the key, be sure to set the key type to JSON. Otherwise, the Citrix Full Configuration interface cannot parse it.

**Tip:**

Create keys using the **Service accounts** page in the GCP console. We recommend that you change keys regularly for security purposes. You can provide new keys to the Citrix Virtual Apps and Desktops application by editing an existing GCP connection.

Also, you need to grant your service account the necessary permissions to access your GCP project:

1. In the GCP console, navigate to **IAM & Admin > IAM**. On the **IAM** page, locate the service account you created and then select the pencil icon to edit the service account.
2. On the **Edit permissions** page, select **ADD ANOTHER ROLE** to add the following roles to your service account one by one and then select **SAVE**.
  - Compute Admin
  - Storage Admin
  - Cloud Build Editor
  - Service Account User
  - Cloud Datastore User
3. Update the roles assigned to your project's Cloud Build service account:
  - a) In the GCP console, navigate to **IAM & Admin > IAM**.
  - b) On the **IAM** page, locate the Cloud Build service account and then select the pencil icon to edit the service account. You can identify the Cloud Build service account by its

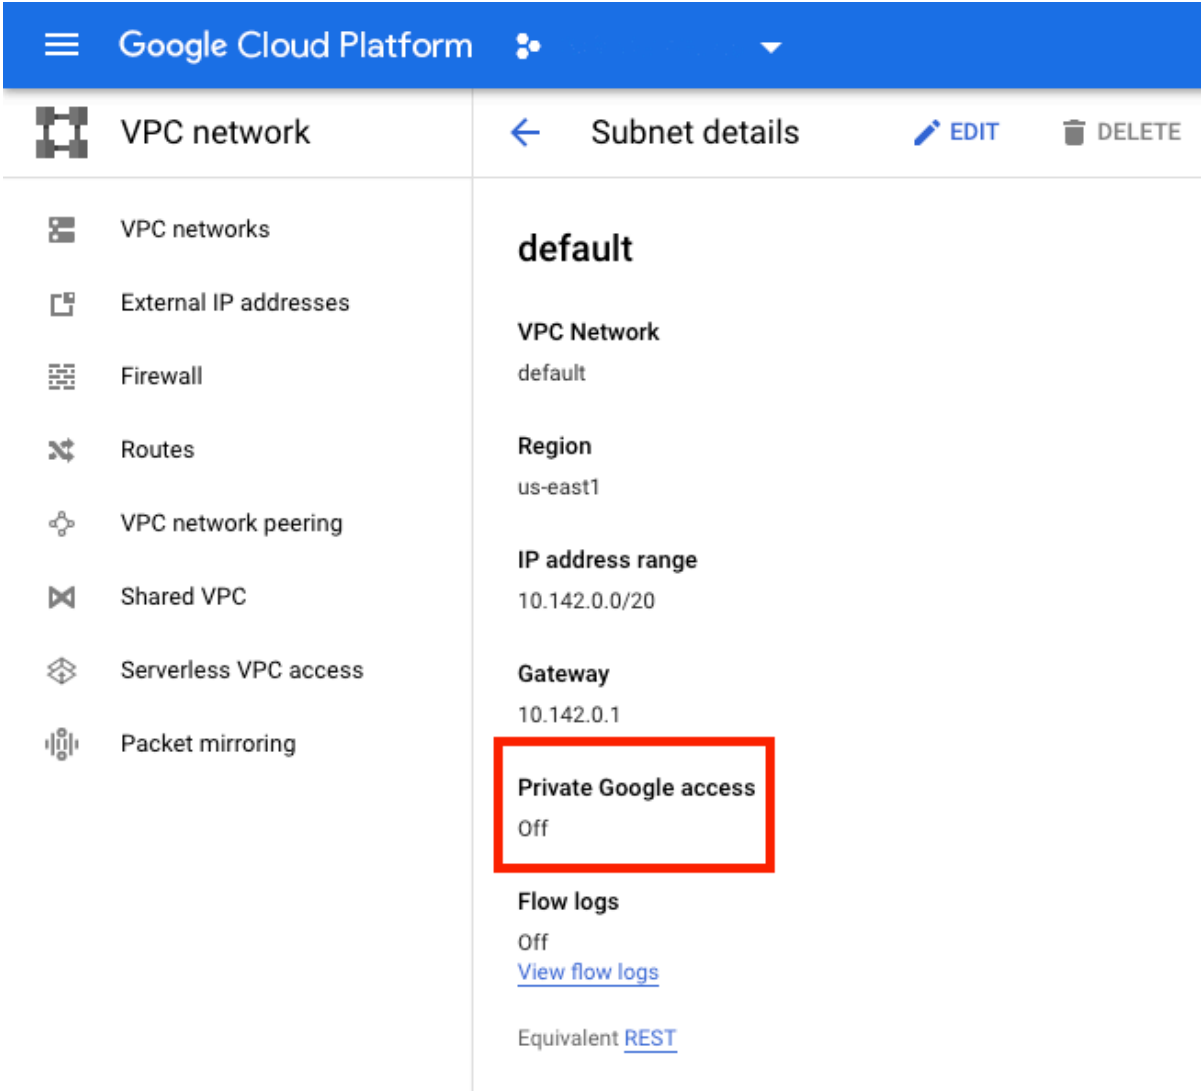
user name, which is in this format: <your\_gcp\_project\_ID\_number>@cloudbuild.gserviceaccount.com.

- c) On the **Edit permissions** page, select **ADD ANOTHER ROLE** to add the following roles to your Cloud Build service account one by one and then select **SAVE**.
- Cloud Build Service Account
  - Compute Instance Admin
  - Service Account User

### **Enable private Google access**

When a VM lacks an external IP address assigned to its network interface, packets are only sent to other internal IP addresses destinations. When you enable private access, the VM connects to the set of external IP addresses used by the Google API and associated services. To ensure that a VM in your subnet can access the Google APIs without a public IP address for MCS provisioning:

1. In GCP, access the **VPC network configuration**.
2. In the Subnet details screen, turn on **Private Google access**.



The screenshot shows the Google Cloud Platform console interface. At the top, there is a blue header with the Google Cloud Platform logo and a navigation menu. Below the header, the left sidebar contains a list of network-related services: VPC networks, External IP addresses, Firewall, Routes, VPC network peering, Shared VPC, Serverless VPC access, and Packet mirroring. The main content area displays the 'Subnet details' for a subnet named 'default'. The details include the VPC Network (default), Region (us-east1), IP address range (10.142.0.0/20), and Gateway (10.142.0.1). The 'Private Google access' toggle is highlighted with a red box and is currently set to 'Off'. Other details include 'Flow logs' (Off) with a link to 'View flow logs' and 'Equivalent REST'.

For more information, see [Configuring Private Google Access](#).

**Important:**

If your network is configured to prevent VM access to the Internet, ensure that your organization assumes the risks associated with enabling Private Google access for the subnet to which the VM is connected.

### Add a connection

In the Full Configuration interface, follow the guidance in [Create a connection and resources](#). The following description guides you through setting up a hosting connection:

1. From **Manage > Full Configuration**, select **Hosting** in the left pane.
2. Select **Add Connection and Resources** in the action bar.

3. On the **Connection** page, select **Create a new Connection** and **Studio tools**, and then select **Next**.
  - **Connection type.** Select **Google Cloud Platform** from the menu.
  - **Service account key.** Import the key contained in your Google credential file (.json). To do so, locate your credential file, open the file with Notepad (or any text editor), and then copy the content. After that, return to the **Connection** page, select **Import key**, paste the content, and then select **Save**.
  - **Service account ID.** The field automatically populates with the information from the imported key.
  - **Connection name.** Type a name for the connection.
4. On the **Region** page, select a project name from the menu, select a region containing the resources you want to use, and then select **Next**.
5. On the **Network** page, type a name for the resources, select a virtual network from the menu, select a subset, and then select **Next**. The resource name helps identify the region and network combination. Virtual networks with the (*Shared*) suffix appended to their name represent shared VPCs. If you configure a subnet-level IAM role for a shared VPC, only specific subnets of the shared VPC appear on the subnet list.

**Note:**

  - The resource name can contain 1–64 characters, and cannot contain only blank spaces or the characters \ / ; : ## . \* ? = < > | [ ] { } " ' ( ) ' ).
6. On the **Summary** page, confirm the information and then select **Finish** to exit the **Add Connection and Resources** window.

After creating the connection and resources, the connection and resources you created are listed. To configure the connection, select the connection and then select the applicable option in the action bar.

Similarly, you can delete, rename, or test the resources created under the connection. To do so, select the resource under the connection and then select the applicable option in the action bar.

## Prepare a master VM instance and a persistent disk

**Tip:**

Persistent disk is the GCP term for virtual disk.

To prepare your master VM instance, create and configure a VM instance with properties that match the configuration you want for the cloned VDA instances in your planned machine catalog. The configuration does not apply only to the instance size and type. It also includes instance attributes such as metadata, tags, GPU assignments, network tags, and service account properties.

As part of the mastering process, MCS uses your master VM instance to create the GCP *instance template*. The instance template is then used to create the cloned VDA instances that comprise the machine catalog. Cloned instances inherit the properties (except the VPC, subnet, and persistent disk properties) of the master VM instance from which the instance template was created.

After configuring the properties of the master VM instance to your specifics, start the instance and then prepare the persistent disk for the instance.

We recommend that you manually create a snapshot of the disk. Doing so lets you use a meaningful naming convention to track versions, gives you more options to manage earlier versions of your master image, and saves time for machine catalog creation. If you do not create your own snapshot, MCS creates one for you. You can use it to create the custom image in your GCP image library.

## Create a machine catalog

### Note:

Create your resources before you create a machine catalog. Use the naming conventions established by GCP when configuring machine catalogs. See [Bucket and object naming guidelines](#) for more information.

Follow the guidance in [Create machine catalogs](#). The following description is unique to GCP catalogs.

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select **Create Machine Catalog** in the action bar.
3. On the **Operating System** page, select **Multi-session OS** and then select **Next**.
  - The Citrix Virtual Apps and Desktops service also supports single-session OS.
4. On the **Machine Management** page, select the **Machines that are power managed** and the **Citrix Machine Creation Services** options and then select **Next**. If there are multiple resources, select one from the menu.
5. On the **Master Image** page, select a VM and the minimum functional level for the catalog and then select **Next**. If you want to use the sole tenancy functionality, be sure to select an image whose node group property is correctly configured. See [Enable zone selection](#).
6. On the **Virtual Machines** page, specify how many VMs you want to create, view the detailed specification of the VMs, and then select **Next**. If you use sole tenant node groups for machine catalogs, be sure to select **only** the zones where reserved sole tenant nodes are available. See [Enable zone selection](#).
7. On the **Disk Settings** page, choose whether to use your own key to protect disk contents. To use the feature, you must first create your own Customer Managed Encryption Keys (CMEKs). For more information, see [Using Customer Managed Encryption Keys \(CMEK\)](#).

**Note:**

This feature is available as a preview. It is available only in the **Manage > Full Configuration** interface.

After creating the keys, you can select one of those keys from the list. You cannot change the key after you create the catalog. Google Cloud Platform does not support rotating keys on existing persistent disks or images. Therefore, after you provision a catalog, the catalog is tied to a specific version of the key. If that key is disabled or destroyed, the instances and disks encrypted with it become unusable until the key is reenabled or restored.

8. On the **Machine Identities** page, select an Active Directory account and then select **Next**.
  - If you select **Create new Active Directory accounts**, select a domain and then enter the sequence of characters representing the naming scheme for the provisioned VM computer accounts created in Active Directory. The account naming scheme can contain 1–64 characters, and cannot contain blank spaces, or non-ASCII or special characters.
  - If you select **Use existing Active Directory accounts**, select **Browse** to navigate to the existing Active Directory computer accounts for the selected machines.
9. On the **Domain Credentials** page, select **Enter credentials**, type the user name and password, select **Save**, and then select **Next**.
  - The credential you type must have permissions to perform Active Directory account operations.
10. On the **Scopes** page, select scopes for the machine catalog and then select **Next**.
  - You can select optional scopes or select **custom scope** to customize scopes as needed.
11. On the **Summary** page, confirm the information, specify a name for the catalog, and then select **Finish**.

**Note:**

The catalog name can contain 1–39 characters, and cannot contain only blank spaces or the characters `\ / ; : ## . * ? = < > | [ ] { } "' ( )' )`.

Machine catalog creation might take a long time to complete. When it completes, the catalog is listed. You can verify that the machines are created on the target node groups in the GCP console.

## Add machines to a catalog

To add machines to a catalog, follow these steps:

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select the machine catalog to which you want to add machines.

3. Select **Add Machines** in the action bar.
4. On the **Virtual Machines** page, specify the number of machines you want to add and then select **Next**.
5. On the **Computer Accounts** page, select an Active Directory account and then select **Next**.
6. On the **Domain Credentials** page, select **Enter credentials**, type the user name and password, select **Save**, and then select **Next**.
7. On the **Summary** page, confirm the information and then select **Finish**.

## Update machines

This feature can be useful in cases where you want to update your master image or the minimum functional level.

To update machines, follow these steps:

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select the machine catalog that contains machines you want to update.
3. Select **Update Machines** in the action bar.
4. On the **Master Image** page, select a VM and the minimum functional level for the catalog and then select **Next**.
5. On the **Rollout Strategy** page, specify when you want to update the machines and then select **Next**.
6. On the **Summary** page, confirm the information and then select **Finish**.

To roll back a machine update, follow these steps:

### Important:

Do not rename, delete, or move master images. Otherwise you cannot roll back the update.

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select the machine catalog where you want to roll back the machine update.
3. Select **Rollback Machine Update** in the action bar.
4. On the **Overview** page, confirm the information and then select **Next**.
5. On the **Rollout Strategy** page, configure the rollout strategy and then select **Next**.
6. On the **Summary** page, confirm the information and then select **Finish**.

## Power management

The Citrix Virtual Apps and Desktops service lets you power manage GCP machines. Use the **Search** node in the navigation pane to locate the machine you want to power manage. The following power actions are available:

- Delete
- Start
- Restart
- Force Restart
- Shut Down
- Force Shutdown
- Add to Delivery Group
- Manage Tags
- Turn On Maintenance Mode

You can also power manage GCP machines by using Autoscale. To do so, add the GCP machines to a Delivery Group and then enable Autoscale for that Delivery Group. For more information about Autoscale, see [Autoscale](#).

## Import manually created GCP machines

You can *create a connection to GCP* and then *create a catalog containing GCP machines*. Then, you can manually power cycle GCP machines through Citrix Virtual Apps and Desktops service. With this feature, you can:

- Import manually created GCP multi-session OS machines into a Citrix Virtual Apps and Desktops machine catalog.
- Remove manually created GCP multi-session OS machines from a Citrix Virtual Apps and Desktops catalog.
- Use existing Citrix Virtual Apps and Desktops power management capabilities to power manage GCP Windows multi-session OS machines. For example, set a restart schedule for those machines.

This functionality does not require changes to an existing Citrix Virtual Apps and Desktops provisioning workflow, nor the removal of any existing feature. We recommend that you use MCS to provision machines in the Citrix service's Full Configuration interface instead of importing manually created GCP machines.

## Shared Virtual Private Cloud

Shared Virtual Private Clouds (VPCs) comprise a host project, from which the shared subnets are made available, and one or more service projects that use the resource. Shared VPCs are desirable options



for larger installations because they provide centralized control, usage, and administration of shared corporate Google cloud resources. For more information, see the [Google Documentation site](#).

With this feature, Machine Creation Services (MCS) supports provisioning and managing machine catalogs deployed to Shared VPCs. This support, which is functionally equivalent to the support currently provided in local VPCs, differs in two areas:

1. You must grant extra permissions to the Service Account used to create the Host Connection. This process allows MCS to access and use Shared VPC Resources.
2. You must create two firewall rules, one each for ingress and egress. These firewall rules are used during the image mastering process.

### New permissions required

A GCP service account with specific permissions is required when creating the host connection. These additional permissions must be granted to any service accounts used to create Shared VPC based host connections.

#### Tip:

These additional permissions are not new to the Citrix Virtual Apps and Desktops service. They are used to facilitate the implementation of local VPCs. With Shared VPCs, these additional permissions allow access to other shared VPC resources.

A maximum of four extra permissions must be granted to the service account associated with the host connection to support Shared VPC:

1. **compute.firewalls.list** - This permission is mandatory. It allows MCS to retrieve the list of firewall rules present on the Shared VPC.
2. **compute.networks.list** - This permission is mandatory. It allows MCS to identify the Shared VPC networks available to the service account.
3. **compute.subnetworks.list** - This permission is optional depending on how you use VPCs. It allows MCS to identify the subnets within the visible Shared VPCs. This permission is already required when using local VPCs but must also be assigned in the Shared VPC host project.
4. **compute.subnetworks.use** - This permission is optional depending on how you use VPCs. It is necessary to use subnet resources in the provisioned machine catalogs. This permission is already required for using local VPCs but must also be assigned in the Shared VPC host project.

When using these permissions, consider that there are different approaches based on the type of permission used to create the machine catalog:

- Project-level permission:
  - Allows access to all Shared VPCs within the host project.
  - Requires the permissions #3 and #4 must be assigned to the service account.
- Subnet-level permission:

- Allows access to specific subnets within the Shared VPC.
- Permissions #3 and #4 are intrinsic to the subnet level assignment and therefore do not need to be assigned directly to the service account.

Select the approach that matches your organizational needs and security standards.

**Tip:**

For more information about the differences between project-level and subnet-level permissions, see the [Google Cloud documentation](#).

## Firewall Rules

During the preparation of a machine catalog, a machine image is prepared to serve as the master image system disk for the catalog. When this process occurs, the disk is temporarily attached to a virtual machine. This VM must run in an isolated environment that prevents all inbound and outbound network traffic. This is accomplished through a pair of deny-all firewall rules; one for ingress and one for egress traffic. When using GCP local VPCs, MCS creates this firewall in the local network and applies it to the machine for mastering. After mastering completes, the firewall rule is removed from the image.

We recommend keeping the number of new permissions required to use Shared VPCs to a minimum. Shared VPCs are higher-level corporate resources and typically have more rigid security protocols in place. For this reason, create a pair of firewall rules in the host project on the shared VPC resources, one for ingress and one for egress. Assign the highest priority to them. Apply a new target tag to each of these rules, using the following value:

```
citrix-provisioning-quarantine-firewall
```

When MCS creates or updates a machine catalog, it searches for firewall rules containing this target tag. It then examines the rules for correctness and applies them to the machine used to prepare the master image for the catalog. If the firewall rules are not found, or the rules are found but the rules or their priorities are incorrect, a message similar to the following appears:

```
"Unable to find valid INGRESS and EGRESS quarantine firewall rules for VPC <name> in project <project>. "Please ensure you have created 'deny all' firewall rules with the network tag 'citrix-provisioning-quarantine-firewall' and proper priority." "Refer to Citrix Documentation for details."
```

## Configuring the shared VPC

Before adding the Shared VPC as a host connection in the Citrix service's Full Configuration interface, complete the following steps to add service accounts from the project you intend to provision into:

1. Create an IAM role.

2. Add the service account used to create a CVAD host connection to the Shared VPC host project IAM role.
3. Add the Cloud Build service account from the project you intend to provision into to the Shared VPC host project IAM role.
4. Create firewall rules.

### Create an IAM role

Determine the role's access level — *project level access* or a more restricted model using *subnet level access*.

**Project level access for IAM role.** For the project level IAM role, include the following permissions:

- `compute.firewalls.list`
- `compute.networks.list`
- `compute.subnetworks.list`
- `compute.subnetworks.use`

To create a project level IAM role:

1. In the GCP console, navigate to **IAM & Admin > Roles**.
2. On the **Roles** page, select **CREATE ROLE**.
3. On the **Create Role** page, specify the role name. Select **ADD PERMISSIONS**.
  - a) On the **Add permissions** page, add permissions to the role, individually. To add a permission, type the name of the permission in the **Filter table** field. Select the permission and then select **ADD**.
  - b) Select **CREATE**.

**Subnet-level IAM role.** This role omits the addition of the permissions `compute.subnetworks.list` and `compute.subnetworks.use` after selecting **CREATE ROLE**. For this IAM access level, the permissions `compute.firewalls.list` and `compute.networks.list` must be applied to the new role.

To create a subnet level IAM role:

1. In the GCP console, navigate to **VPC network > Shared VPC**. The **Shared VPC** page appears, displaying the subnets of the Shared VPC networks that the host project contains.
2. On the **Shared VPC** page, select the subnet that you want to access.
3. In the top-right corner, select **ADD MEMBER** to add a service account.
4. On the **Add members** page, complete these steps:
  - a) In the **New members** field, type the name of your service account and then select your service account in the menu.
  - b) Select the **Select a role** field and then **Compute Network User**.
  - c) Select **SAVE**.

5. In the GCP console, navigate to **IAM & Admin > Roles**.
6. On the **Roles** page, select **CREATE ROLE**.
7. On the **Create Role** page, specify the role name. Select **ADD PERMISSIONS**.
  - a) On the **Add permissions** page, add permissions to the role, individually. To add a permission, type the name of the permission in the **Filter table** field. Select the permission, and then select **ADD**.
  - b) Select **CREATE**.

### Add a service account to the host project IAM role

After creating an IAM role, perform the following steps to add a service account for the host project:


1. In the GCP console, navigate to the host project and then to **IAM & Admin > IAM**.
2. On the **IAM** page, select **ADD** to add a service account.
3. On the **Add members** page:
  - a) In the **New members** field, type the name of your service account and then select your service account in the menu.
  - b) Select the **Select a role** field, type the IAM role you created, and then select the role in the menu.
  - c) Select **SAVE**.

The service account is now configured for the host project.

### Add the cloud build service account to the shared VPC

Every Google Cloud subscription has a service account that is named after the project ID number, followed by `cloudbuild.gserviceaccount`. For example: `705794712345@cloudbuild.gserviceaccount`.

You can determine what the project ID number is for your project by selecting **Home** and **Dashboard** in the Google Cloud console:



Find the **Project Number** below the **Project Info** area of the screen.

Perform the following steps to add the Cloud Build service account to the Shared VPC:

1. In the Google Cloud console, navigate to the host project and then to **IAM & Admin > IAM**.
2. On the **Permissions** page, select **ADD** to add an account.
3. On the **Add members** page, complete these steps:
  - a) In the **New members** field, type the name of the Cloud Build service account and then select your service account in the menu.
  - b) Select the **Select a role** field, type `Computer Network User`, and then select the role in the menu.
  - c) Select **SAVE**.

### Create firewall rules

As part of the mastering process, MCS copies the selected machine image and uses it to prepare the master image system disk for the catalog. During mastering, MCS attaches the disk to a temporary virtual machine, which then runs preparation scripts. This VM must run in an isolated environment that prohibits all inbound and outbound network traffic. To create an isolated environment, MCS requires two *deny all* firewall rules (an ingress rule and an egress rule). Therefore, create two firewall rules in the *Host Project* as follows:

1. In the GCP console, navigate to the host project and then to **VPC network > Firewall**.
2. On the **Firewall** page, select **CREATE FIREWALL RULE**.
3. On the **Create a firewall rule** page, complete the following:
  - **Name**. Type a name for the rule.
  - **Network**. Select the Shared VPC network to which the ingress firewall rule applies.
  - **Priority**. The smaller the value is, the higher the priority of the rule is. We recommend a small value (for example, 10).
  - **Direction of traffic**. Select **Ingress**.
  - **Action on match**. Select **Deny**.
  - **Targets**. Use the default, **Specified target tags**.
  - **Target tags**. Type `citrix-provisioning-quarantine-firewall`.
  - **Source filter**. Use the default, **IP ranges**.
  - **Source IP ranges**. Type a range that matches all traffic. Type `0.0.0.0/0`.
  - **Protocols and ports**. Select **Deny all**.
4. Select **CREATE** to create the rule.
5. Repeat steps 1–4 to create another rule. For **Direction of traffic**, select **Egress**.

### Add a connection

After adding the network interfaces to the Cloud Connector instance, [add a connection](#).

## Enable zone selection

The Citrix Virtual Apps and Desktops service supports zone selection. With zone selection, you specify the zones where you want to create VMs. With zone selection, administrators can place sole tenant nodes across zones of their choice. To configure sole tenancy, you must complete the following on GCP

- Reserve a Google Cloud Platform sole-tenant node
- Create the VDA master image

## Reserving a Google Cloud Platform sole-tenant node

To reserve a sole-tenant node, refer to the Google Cloud Platform [documentation](#).

### Important:

A node template is used to indicate performance characteristics of the system that is reserved in the node group. Those characteristics include the number of vGPUs, the amount of memory allocated to the node, and the machine type used for machines created on the node. For more information see the Google Cloud Platform [documentation](#).

## Creating the VDA master image

To deploy machines on the sole-tenant node successfully, you need to take extra steps when creating a master VM image. Machine instances on GCP have a property called *node affinity labels*. Instances used as master images for catalogs deployed to the sole-tenant node require a *node affinity label* that matches the name of the **target node group**. To achieve this, keep the following in mind:

- For a new instance, set the label in the Google Cloud console when creating an instance. For details, see [Set a node affinity label when creating an instance](#).
- For an existing instance, set the label by using the **gcloud** command line. For details, see [Set a node affinity label for an existing instance](#).

### Note:

If you intend to use sole tenancy with a shared VPC, see [Shared Virtual Private Cloud](#).

## Set a node affinity label when creating an instance

To set the node affinity label:

1. In the GCP console, navigate to **Compute Engine > VM instances**.
2. On the **VM instances** page, select **Create instance**.

3. On the **Instance creation** page, type or configure the required information and then select **management, security, disks, networking, sole tenancy** to open the settings panel.
4. On the **Sole tenancy** tab, select **Browse** to view the available node groups in the current project. The **Sole-tenant node** page appears, displaying a list of available node groups.
5. On the **Sole-tenant node** page, select the applicable node group from the list and then select **Select** to return to the **Sole tenancy** tab. The node affinity labels field populates with the information you selected. This setting ensures that machine catalogs created from the instance will be deployed to the selected node group.
6. Select **Create** to create the instance.

### Set a node affinity label for an existing instance

To set the node affinity label:

1. In the Google Cloud Shell terminal window, use the `gcloud compute instances` command to set a node affinity label. Include the following information in the **gcloud** command:
  - **Name of the VM.** For example, use an existing VM named `s*2019-vda-base.*`
  - **Name of the node group.** Use the node group name you previously created. For example, `mh-sole-tenant-node-group-1`.
  - **The zone where the instance resides.** For example, the VM resides in the `*us-east-1b* zone`.

For example, type the following command in the terminal window:

```
gcloud compute instances set-scheduling "s2019-vda-base"--node-group="mh-sole-tenant-node-group-1"--zone="us-east1-b"
```

For more information about the `gcloud compute instances` command, see the Google Developer Tools documentation at <https://cloud.google.com/sdk/gcloud/reference/beta/compute/instances/set-scheduling>.

2. Navigate to the **VM instance details** page of the instance and verify that the **Node Affinities** field populates with the label.

### Create a machine catalog

After setting the node affinity label, [configure the machine catalog](#).

### Preview: Using Customer Managed Encryption Keys (CMEK)

You can use Customer Managed Encryption Keys (CMEK) for MCS catalogs. When using this functionality, you assign the Google Cloud Key Management Service `CryptoKey Encrypter/Decrypter` role

to the Compute Engine Service Agent. Refer to [Helping to protect resources by using Cloud KMS keys](#) for more information.

Your Compute Engine Service Agent is in the following form: `service-<Project _Number>@compute-system.iam.gserviceaccount.com`. This form is different than the default Compute Engine Service Account.

**Note:**

This Compute Engine Service Account might not appear in the Google Console **IAM Permissions** display. In such cases, use the `gcloud` command as described in [Helping to protect resources by using Cloud KMS keys](#).

### Assign permissions to the Citrix Virtual Apps and Desktops Service account

GCP Cloud KMS permissions can be configured in various ways. You can either provide *project level* KMS permissions or *resource level* KMS permissions. See [Permissions and roles](#) for more information.

#### Project level permissions

One option is to provide the Citrix Virtual Apps and Desktops Service account with project-level permissions to browse Cloud KMS resources. To do this, create a custom role, and add the following permissions:

- `cloudkms.keyRings.list`
- `cloudkms.keyRings.get`
- `cloudkms.cryptokeys.list`
- `cloudkms.cryptokeys.get`

Assign this custom role to your Citrix Virtual Apps and Desktops Service account. This allows you to browse regional keys in the relevant project in the inventory.

#### Resource Level Permissions

For the other option, resource level permissions, in the GCP console, browse to the `cryptoKey` you use for MCS provisioning. Add the Citrix Virtual Apps and Desktops Service account to a key ring or a key that you use for catalog provisioning.

**Tip:**

With this option, you cannot browse regional keys for your project in the inventory because the Citrix Virtual Apps and Desktops Service account does not have project-level list permissions on the Cloud KMS resources. However, you can still provision a catalog using CMEK by specifying the correct `cryptoKeyId` in the `ProvScheme` custom properties, described below.



## Provisioning with CMEK using custom properties

When creating your Provisioning Scheme via PowerShell, specify a `CryptoKeyId` property in `ProvScheme CustomProperties`. For example:

```
1 '<CustomProperties xmlns="http://schemas.citrix.com/2014/xd/
   machinecreation" xmlns:xsi="http://www.w3.org/2001/XMLSchema-
   instance">
2   <Property xsi:type="StringProperty" Name="CryptoKeyId" Value="<
   yourCryptoKeyId>" />
3 </CustomProperties>'
4 <!--NeedCopy-->
```

The `cryptoKeyId` should be specified in the following format:

`projectId:location:keyRingName:cryptoKeyName`

For example, if you'd like to use the key `my-example-key` in key ring `my-example-key-ring` in the region `us-east1` and project with ID `my-example-project-1`, your `ProvScheme` custom settings would resemble:

```
1 '<CustomProperties xmlns="http://schemas.citrix.com/2014/xd/
   machinecreation" xmlns:xsi="http://www.w3.org/2001/XMLSchema-
   instance">
2   <Property xsi:type="StringProperty" Name="CryptoKeyId" Value="my-
   example-project-1:us-east1:my-example-key-ring:my-example-key"
   />
3 </CustomProperties>'
4 <!--NeedCopy-->
```

All MCS provisioned disks and images related to this provisioning scheme use this customer managed encryption key.

### Tip:

If you use global keys, the customer properties location must say `global` and not the **region** name, which in the example above is `us-east1`. For example: `<Property xsi:type="StringProperty"Name="CryptoKeyId"Value="my-example-project-1:global:my-example-key-ring:my-example-key"/>`.

## Rotating customer managed keys

GCP does not support rotating keys on existing persistent disks or images. Once a machine is provisioned it is tied to the key version in use at the time it was created. However, a new version of the key can be created and that new key is used for newly provisioned machines or resources created when a catalog is updated with a new master image.

### Important considerations about key rings

Key rings cannot be renamed or deleted. Also, you might incur unforeseen charges when configuring them. When deleting or removing a key ring, GCP displays an error message:

```
1 Sorry, you can't delete or rename keys or key rings. We were concerned
  about the security implications of allowing multiple keys or key
  versions over time to have the same resource name, so we decided to
  make names immutable. (And you can't delete them, because we wouldn't
  be able to do a true deletion--there would still have to be a
  tombstone tracking that this name had been used and couldn't be
  reused).
2 We're aware that this can make things untidy, but we have no immediate
  plans to change this.
3 If you want to avoid getting billed for a key or otherwise make it
  unavailable, you can do so by deleting all the key versions; neither
  keys nor key rings are billed for, just the active key versions
  within the keys.
4 <!--NeedCopy-->
```

#### Tip:

For more information, see [Editing or deleting a key ring from the console](#).

### More information

- [Connections and resources](#)
- [Create machine catalogs](#)

## Scale and size considerations for Cloud Connectors

September 8, 2021

When evaluating the Citrix Virtual Apps and Desktops service for sizing and scalability, consider all the components. Research and test the configuration of the Cloud Connectors and the customer-managed StoreFront for your specific requirements. Undersizing the machines can impact system performance negatively. This article provides details of the tested maximum capacities, and best practice recommendations for Cloud Connector machine configuration.

## Summary

All results in this summary are based on the findings from a test environment as configured in the detailed sections of this document. Different system configurations may yield different results.

Key results from testing:

- The Citrix Virtual Apps and Desktops service sizing and scalability
  - A set of three 4-vCPU Cloud Connectors is recommended for sites that host no more than 5,000 Workstation VDAs.
    - \* This is an N+1 High Availability configuration.
  - Provisioning 1,000 VMs takes an average of 140 minutes.
- Citrix Virtual Desktops Essentials
  - Two Cloud Connectors hosted on Azure Standard\_A2\_v2 VMs are recommended for 1,000 Windows 10 VMs.
  - Starting 1,000 sessions to Windows 10 VMs hosted in Azure takes less than 20 minutes.
  - Testing found that it takes approximately 44 seconds from when a user logs on to StoreFront until the user receives a functional VDI desktop with default settings.
  - Provisioning 1,000 Windows 10 VMs in Azure takes an average of 8 hrs.
- Citrix Cloud manages Cloud Connector services, and the customer manages the machines.
- Session launch testing for Citrix Virtual Desktops Essentials used a NetScaler appliance. All other session launch testing used connections direct to StoreFront.

## Test methodology

Tests were conducted to add load and to measure the performance of the environment components. The components were monitored by collecting performance data and procedure timing (such as logon time, machine creation time). In some cases, proprietary Citrix simulation tools were used to simulate VDAs and sessions. These tools are designed to exercise Citrix components the same way that traditional VDAs and sessions do, without the same resource requirements to host real sessions and VDAs. We executed the following tests:

- Session logon storm: a test that simulates high-volume logon periods
- VDA registration storm: a test that simulates high-volume VDA registration periods. For example, following an upgrade cycle or outage recovery.
- Machine Creation Service provisioning: a test that measure the time to perform tasks such as copying images, creating Active Directory accounts, and creating machines.

We used the data gathered from these tests to make recommendations for Cloud Connector sizing. The test execution details follow.

## **Session logon storm tests**

Sessions are started at customer-managed StoreFront servers independently. 1,000 session, 5,000 session, and 20,000 session tests were run against each environment. We collected StoreFront logon, resource enumeration, ICA file retrieval, and active desktop times. The active desktop time is the time from when the ICA file starts until the resource is fully loaded and ready to use.

For some test scenarios, we used simulation tools to facilitate testing of larger user counts. Simulation tools allow testing using less hardware than is required to run 5,000 or 20,000 real sessions. These simulated sessions go through the normal StoreFront logon, resource enumeration, and ICA file retrieval, but do not start active desktops. Instead, the simulation tool reports to the ICA stack that the session has started. All communication from the broker agent to the Broker Service is consistent with the communication of an actual session. Performance metrics are gathered from the Cloud Connectors.

To determine how the environment responded to session launches, a sustained concurrency of 25 session launches was maintained throughout the duration of the test. The measurements therefore show the results of a system under load throughout the test.

## **VDA registration storm tests**

In a VDA registration storm, hundreds or thousands of VDAs are registered all at once to simulate a site recovery. High-volume VDA registration typically happens after the upgrade cycle every two weeks, during a “Monday morning” scenario, or when the system recovers from an outage between customer managed machines and Citrix-managed services. Tests were conducted using 5,000 VDAs and the Cloud Connectors were monitored by gathering performance data during each test. Data included Perfmon counters (CPU, memory, disk utilization) and VDA registration times.

## **Machine Creation Service provisioning tests**

Provisioning tests were performed by creating catalogs of varying counts. The times for various tasks (image copy, AD account creation, and machine creation) were measured to gauge performance. We tested catalog size increases in Azure. Both Azure and customer-managed hypervisors underwent 1,000 machine provisioning testing. The testing in Azure was limited to Windows 10 VMs because Windows 10 is the only supported OS for Citrix Virtual Desktops Essentials. The customer-managed hypervisor was tested on Windows 10 and on Windows 2012 R2.

## **Test environment**

The test environment configuration included Citrix Cloud Connector, Citrix Virtual Apps and Desktops service and Citrix Virtual Apps and Desktops components. The machine and operating system specifications we used are provided here so you can compare our configuration and test results to your own configuration and requirements.

## Tools used

An internal testing tool collected performance data and metrics from the machines under test, and drove the session launches. This proprietary tool orchestrates user session launches to the Citrix Virtual Apps and Desktops environment, and provides a central location for collecting response time data and performance metrics. In essence, the test tool administers the tests and collects the results.

## Test configuration – Citrix Virtual Apps and Desktops

The following is a list of the machine and OS specifications used during Citrix Virtual Apps and Desktops testing.

- **Cloud Connectors:**
  - **Scenario One:** Two Windows 2012 R2, 2 vCPU, 4 GB memory
  - **Scenario Two:** Two Windows 2012 R2, 4 vCPU, 4 GB memory
- **StoreFront (customer-managed):** One Windows 2012 R2, 8 vCPU, 8 GB memory
- **Hypervisors:** Eight VMware vSphere ESXi 6.0 Update 1, HP ProLiant BL 460c Gen9, Two Intel E5-2620 CPU, 256 GB Memory
- **Hypervisor Storage:** 2-TB NFS share on NetApp 3250
- **VDA:** Windows 2012 R2 and Windows 10 32-bit Build 1607

## Test configuration – Citrix Virtual Desktops Essentials

Sessions were started from 100 Windows 2012 R2 client launcher machines. Sessions were authenticated against a Windows Active Directory hosted in Azure. Roaming profiles were stored on a Windows file server in Azure.

- **VDA:** 1,000 Windows 10 64-bit Build 1607, 2 vCPU, 7 GB memory (Standard\_D2\_v2 instance)
- **Client:** 100 Windows 2012 R2 Servers, 8 vCPU, 8 GB memory
- **Domain Controller:** Two Windows 2012 R2, 4 vCPU, 14 GB memory (Standard\_D3\_v2 instance)
- **File Server:** One Windows 2012 R2, DS11 instance
- **NetScaler VPX:** One NetScaler 11.0, Standard\_D3\_v2 instance that has 1,000 Platinum license
- **Cloud Connectors:**
  - **Scenario One:** Two Windows 2012 R2, 2 vCPU, 4 GB memory (Standard\_A2\_v2 instance)
  - **Scenario Two:** Two Windows 2012 R2, 4 vCPU, 7 GB memory (Standard\_A3 instance)
- **StoreFront (customer-managed):** One Windows 2012 R2, DSv2 instance

## Customer-managed machine considerations

Customer-managed machines can be in the customer office, data center, or cloud account (such as Azure or AWS). By our definition, customer-managed machine is under the complete customer control. Customer-managed machines include: Cloud Connector, StoreFront servers, RDS servers, VDI

machines, and Remote PC Access machines (not covered during testing). For the sake of brevity, we refer to RDS servers, VDI machines, and Remote PC Access machines as VDAs throughout this report.

### **StoreFront servers**

We used an 8-vCPU, 8-GB memory machine as the customer-managed StoreFront server when we tested the Citrix Virtual Apps and Desktops service. For Citrix Virtual Desktops Essentials testing, we used an Azure Standard\_DS2\_v2 (2 vCPU, 7 GB memory) for the customer-managed StoreFront server. See [Plan your StoreFront deployment](#) to size your StoreFront server properly for your environment.

### **Cloud Connectors**

We tested customer-managed Cloud Connectors hosted on VMs that had 2-vCPU and 4-GB memory in one scenario, and 4-vCPU and 4-GB memory in another. In Azure, Cloud Connectors were tested on Standard\_A2\_v2 (2 vCPU, 4 GB memory) and Standard\_A3 (4 vCPU, 7 GB memory) instances.

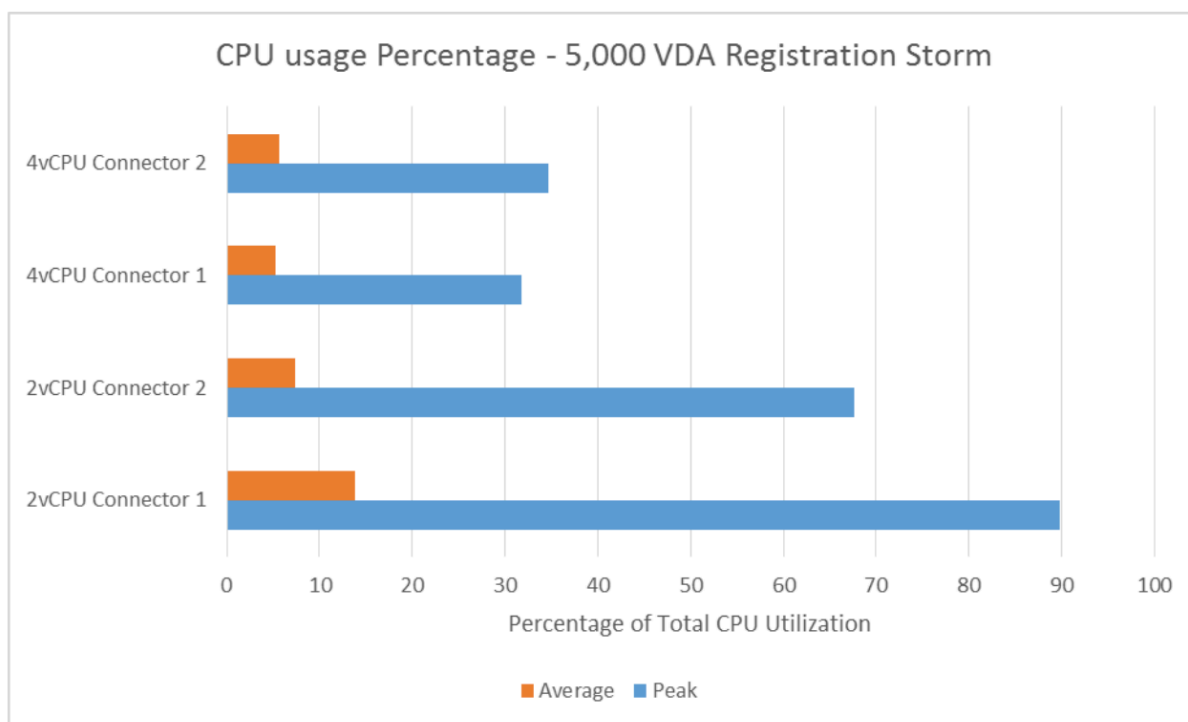
In our testing, Cloud Connectors were deployed in HA sets (**they are not load-balanced**). Although this document focuses on testing environments that have two Cloud Connectors, an N+1 set of three Cloud Connectors is recommended. The rest of this report focuses on the Cloud Connectors and how to size them for best performance.

## **Test results**

### **VDA registration storm**

The VDA Registration Storm test provides data that shows the relationship between Cloud Connector sizing and environment stability. Environment stability is tested during cases of a network outage between the customer-managed location and the Citrix-managed services. VDA registration storms can be triggered when the Delivery Controller and the Site database are upgraded, typically every two weeks.

### **Cloud Connector CPU sizing comparison 2 vCPU vs. 4 vCPU**



- The average usage is similar, but the 2-vCPU machine CPU is under strain during the test and occasional VDA de-registrations are observed.
- The use of 4-vCPU Cloud Connectors for sites that have approximately 5,000 VDAs is recommended for stability.
- The use of 2-vCPU Cloud Connectors is recommended for sites that host 2,500 VDAs.
- Cloud Connectors are a high-availability set and do not load balance.
- One reason we do not recommend the 2-vCPU Cloud Connector for sites that host 5,000 VDAs is the randomness of machine assignment. Because the Cloud Connectors are not load-balanced, you cannot predict the size of the load being funneled to either Cloud Connector. Sometimes, we found more than 60% of the load funneled to one machine.

Number of VDAs	Cloud Connectors required
<2,500	2 VMs + 1, each having 2 vCPUs
<5,000	2 VMs + 1, each having 4 vCPUs

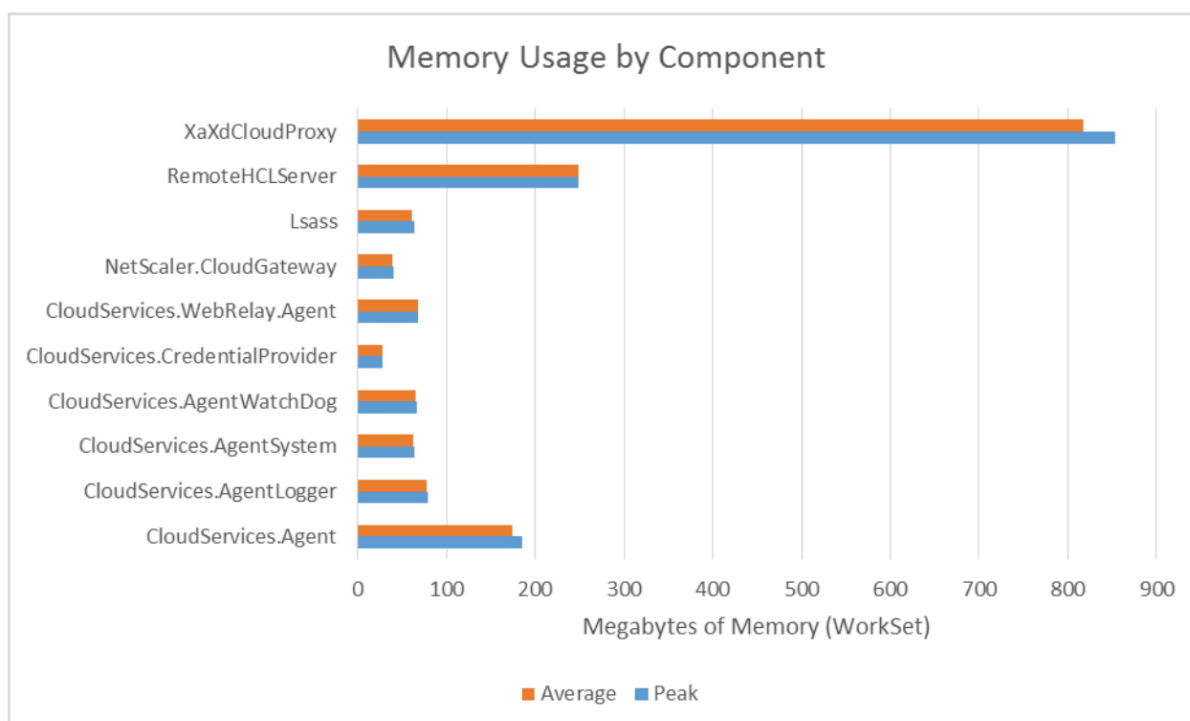
**Cloud Connector HA pair VDA registration storm timing comparison**

Cloud Connector size	VDA count	Registration time
2 VCPU	5000	11:03

Cloud Connector size	VDA count	Registration time
4 VCPU	5000	5:46

- The Cloud Connectors equipped with 4 vCPUs proved to be more stable during testing.
- VDAs registered faster when Cloud Connectors were equipped with 4 vCPUs.
- VDA re-registrations were observed during testing with the 2-vCPU Cloud Connectors.
  - Re-registrations may occur when registration attempts timeout, or VDA communication heartbeats are delayed.

### Memory usage by component on a Cloud Connector during a 5,000 VDA registration storm



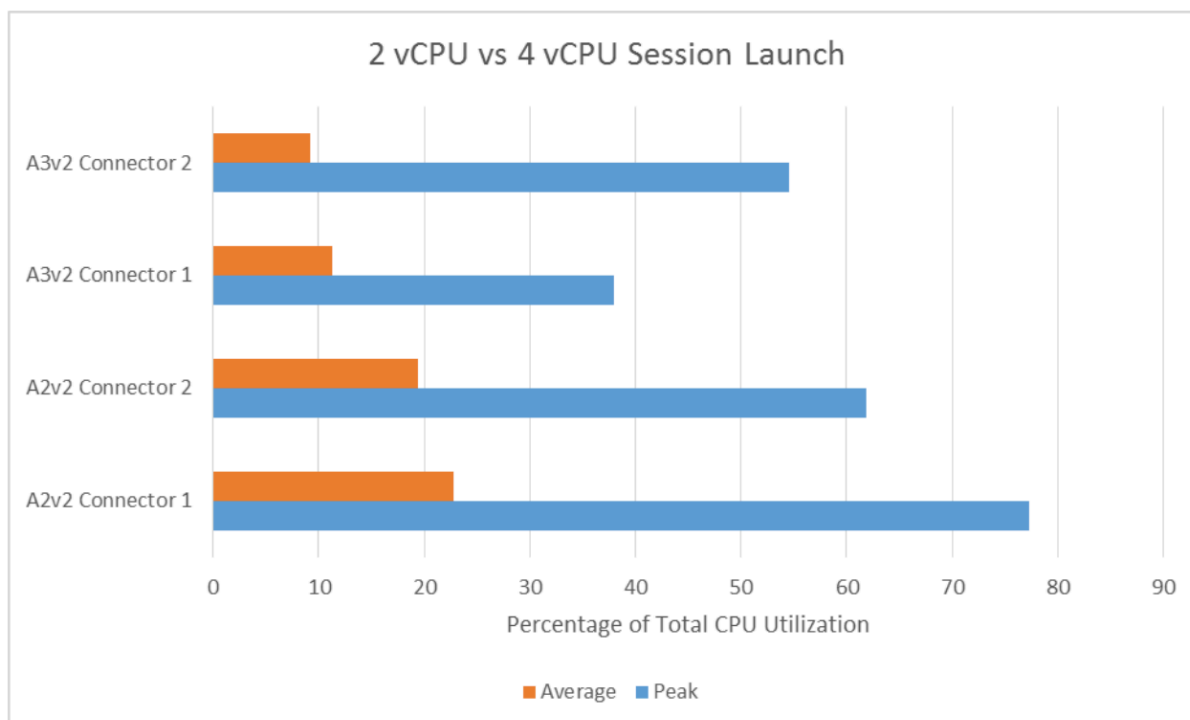
- This graph is a detailed view of the memory usage by Citrix components and Microsoft LSASS (Local Security Authority Subsystem Service), during the registration storm test.
- The LSASS process on the Cloud Connectors plays an important part in both registrations and session launches. All Active Directory authentications, made by the Citrix Cloud services, are proxied to the customer-managed Active Directory via the Cloud Connectors.
- Memory usage peaks during the VDA registration period, decreasing after all the VDAs register successfully.
- High memory utilization is observed on Cloud Connectors that have 4 GB of memory.



### Session launch (Citrix Virtual Desktops Essentials)

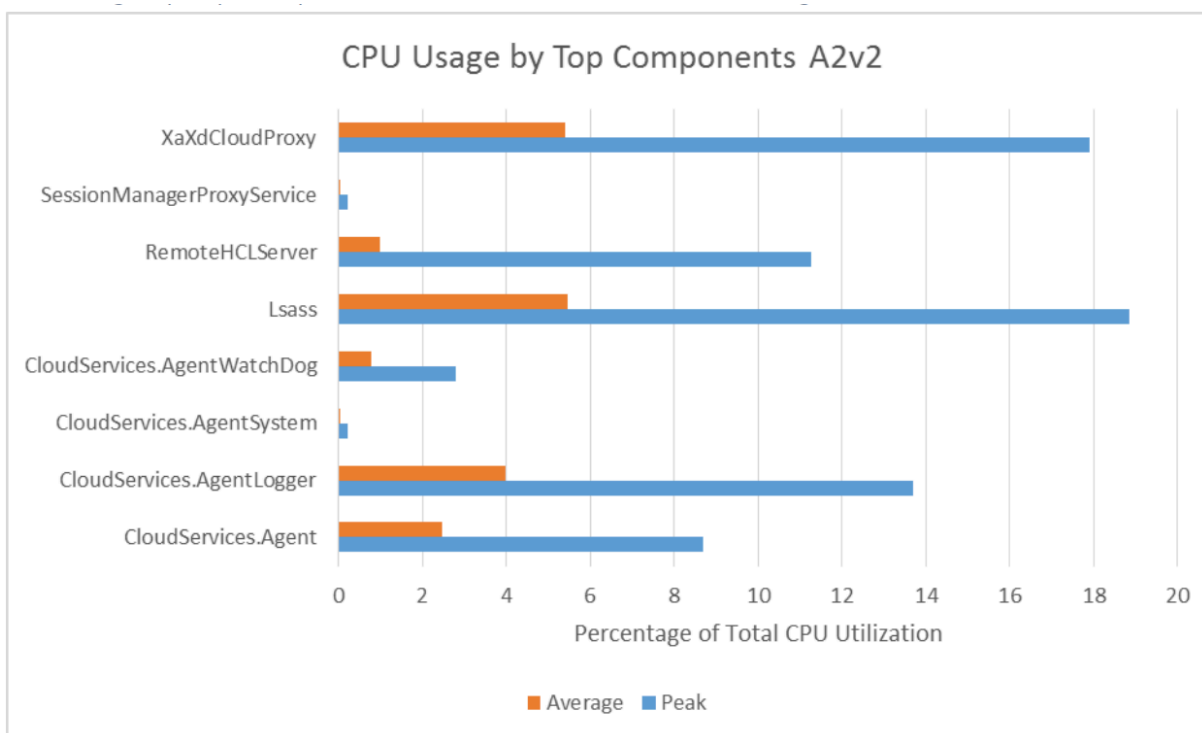
1,000 session launch tests were conducted using the Citrix Virtual Desktops Essentials platform. Testing compared different-sized Cloud Connector instances. We tested the Standard\_A2\_v2 (2 vCPU, 4 GB memory) and Standard\_A3 (4 vCPU, 7 GB memory) instances.

### Connector CPU usage with Citrix manage StoreFront during session launch test



- There is low CPU contention during the test. The Standard\_A2\_v2 instance size was more than able to handle a 1,000 machine VDI deployment during a high load session launch test.
- The Standard\_A3 instance was deemed excessive for this site size, so we continue with a breakdown of the Standard\_A2\_v2.
- Larger VDI sites might see a requirement for using the Standard\_A3.


### CPU usage by top components on A2v2 Cloud Connector during 1,000 session launch



More processes running on the Cloud Connector are not shown because they did not register meaningful metrics.

- The Citrix Remote Broker Provider (XaXdCloudProxy) handles communication between the customer-managed VDA machines and the Citrix-managed Services (Delivery Controller).
- LSASS on the Cloud Connectors processes all Active Directory authentications. The authentications made by the Citrix Cloud Services are proxied to the customer-managed Active Directory via the Cloud Connectors.
- The graph shows the usage from a single Cloud Connector that received a higher amount of load during the test. The additional Cloud Connector in the test exhibited lower CPU usage and was not included in the graph.

### Cloud Connector memory usage instance comparison



- Lower available memory on the Standard\_A2\_v2 (4 GB memory) shows high memory utilization on the Standard\_A2\_v2 VM.
- The high memory utilization is caused by the Citrix Remote HCL Server (RemoteHCLServer) process that maintains the power state of the 1,000 machines in Azure.
  - Due to Azure API rate limitations, the states cannot be queried at regular intervals.
- Changes to the Citrix Remote HCL Server (RemoteHCLServer) implemented after our testing allow the Delivery Controller to communicate machine states directly to Azure.
  - The change reduces memory usage significantly and allows the Standard\_A2\_v2 instances to manage the 1,000 VDA site without issue.

### Session launch times

#### Comparison of the Standard\_A2\_v2 and Standard\_A3

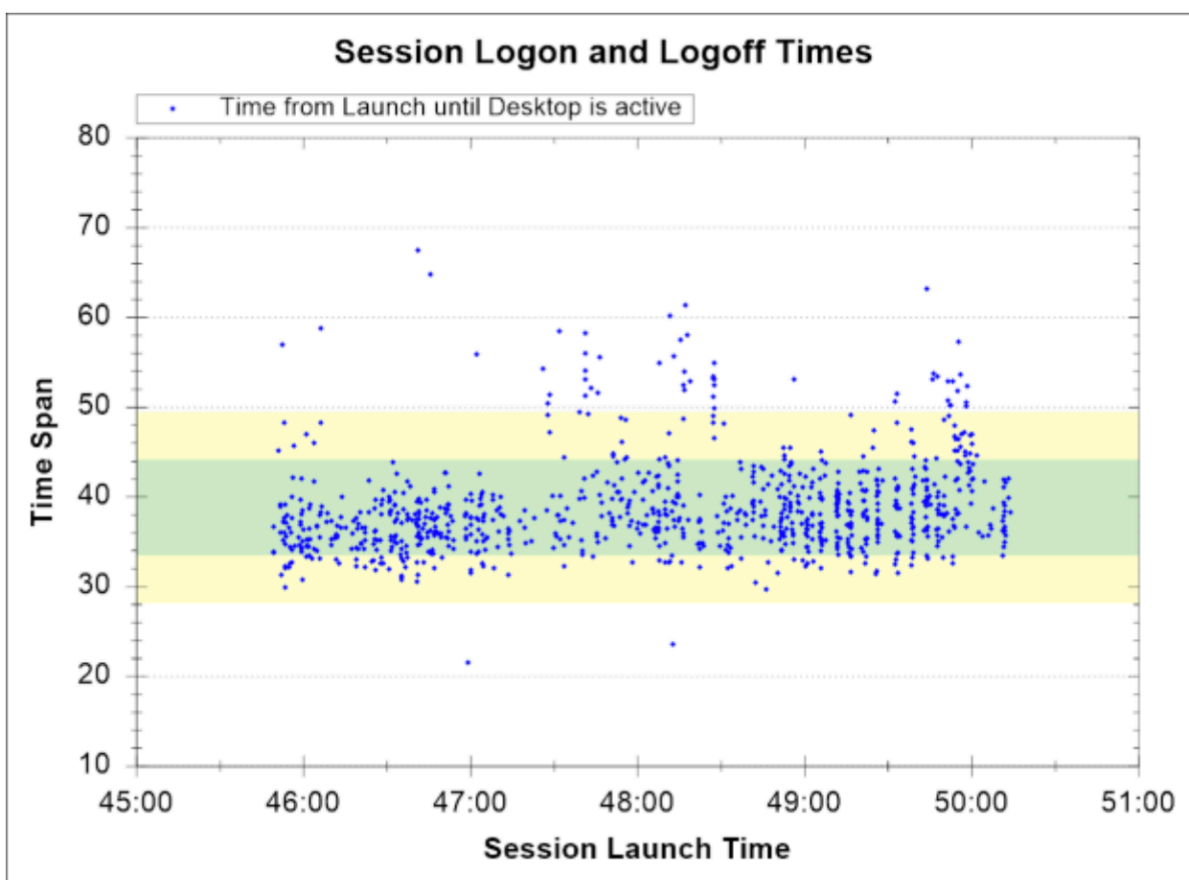
	A3	A2v2
Authenticate	561 ms	575 ms
Enumerate	1,132 ms	1,054 ms
Total login	1,693 ms	1,629 ms
Retrieve ICA file	3,464 ms	3,659 ms

	A3	A2v2
OS logon complete	38.83 seconds	41.91 seconds
Total launch	42.3 seconds	45.6 seconds

Times are the average over all test runs. Customer-managed StoreFront server in Azure:  
Standard\_DS2\_v2 (2 vCPU, 7 GB Memory)

- There were approximately 30 ms of latency between the client machines and NetScaler during testing.
- There is an average 3–4 second decrease in session launches when using a Standard\_A3 instances for Cloud Connectors when the environment is under stress.
  - The Standard\_A3 VM has twice as many CPU cores as the Standard\_A2\_v2
  - There is high memory utilization on the Standard\_A2\_v2 instance during the test.
    - \* High memory utilization was resolved when we removed the RemoteHCLServer communication from the Cloud Connectors in Azure ARM deployments.

**Session log on times for 1,000 Windows 10 sessions**



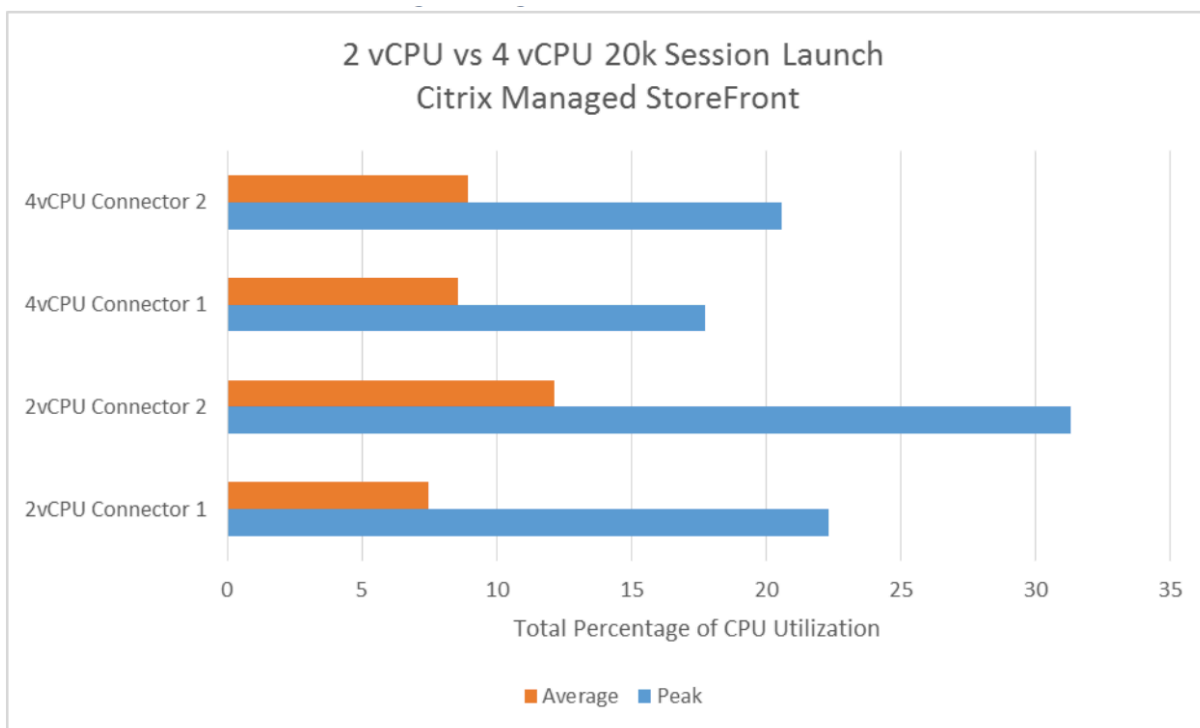
- All machines were powered on before the test.
- The test procedure started 1,000 sessions during approximately an 8-minute period.
- The average time to active desktop with a Standard\_D2\_v2 instance Windows 10 64-bit VDA was approximately 37.67 seconds.
- The graph shows individual logon times over the course of the test, from the time the ICA file is retrieved until an active usable desktop is presented.
  - The green and yellow areas denote one and two standard deviations, respectively.
- Although the session start times are consistent, there are some outliers. Momentary changes in network conditions can cause the outliers, impacting:
  - Secure Ticket Authority (STA) ticket exchange on the NetScaler being proxied via Cloud Connectors.
  - Establishment of an HDX connection over the WAN.
  - Azure Storage. Tests used standard storage.

### **Simulated session launch**

The simulated session launch test puts stress on the Cloud Connectors, Delivery Controller, and site database. Simulated session launch tests the capability of the components to handle a high number of concurrent logons and to sustain those sessions under a sustained load. Session counts of 5,000 and 20,000 were tested. This document focuses on the 20,000-session tests. The launch rate and component

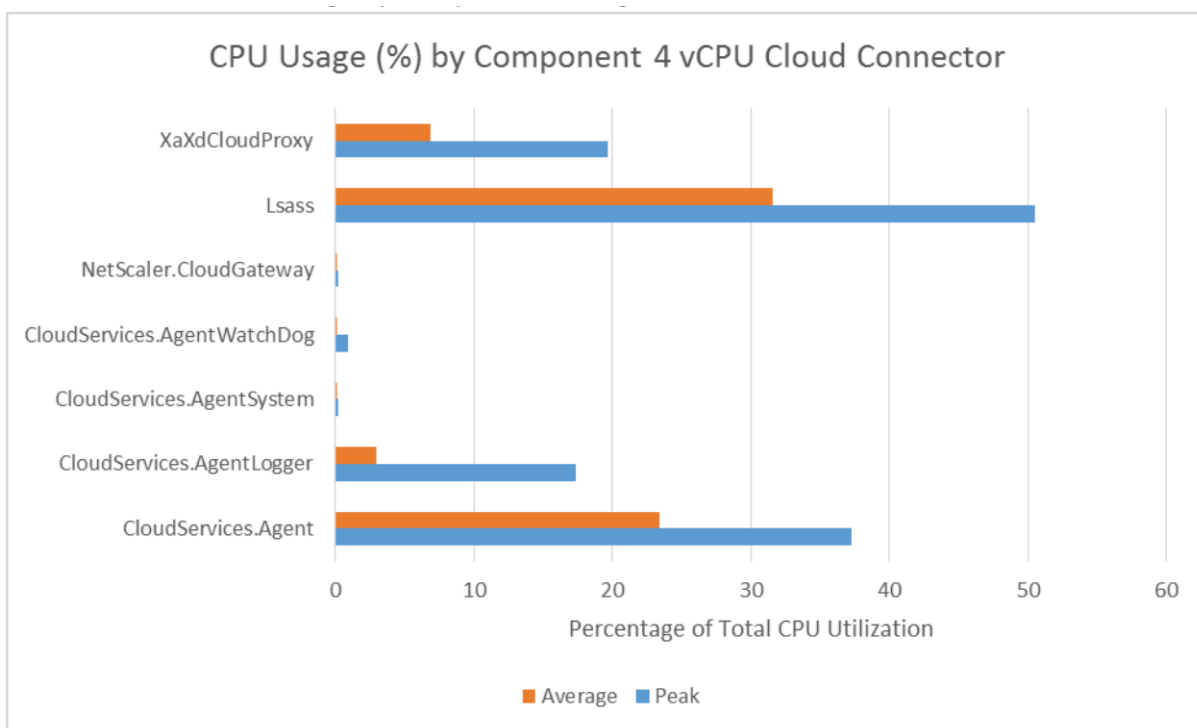
behavior are nearly identical between the two tests. The 20,000-session test runs longer and gives a broader look at the service usage over time. 25 sessions were concurrently launched as fast as possible. The setting for launching sessions as fast as possible allowed the system under test to dictate the rate at which the environment responds to connections.

### **Cloud Connector HA set CPU usage during session launch test**



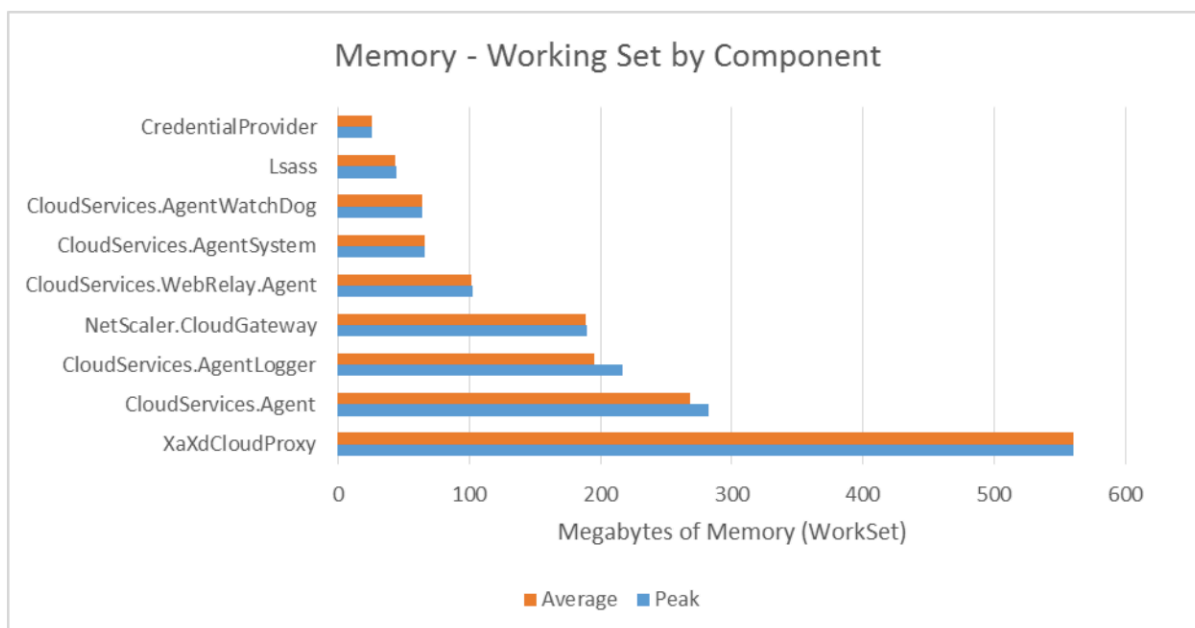
- The graph shows a comparison of Cloud Connector CPU usage during a 20,000-session launch.
- Two Cloud Connectors are deployed for stress and load testing. An N+1 deployment of three Cloud Connectors is recommended for High Availability utilization.
- No CPU contention was observed during the test.

**Cloud Connector CPU usage by component during 20,000 session launch test**



- LSASS (Local Security Authority Subsystem Service) uses CPU during session logons using Store-Front.
- All authentications from Citrix-managed services must traverse the Cloud Connectors to communicate with the customer-managed Active Directory.

**Memory usage by component during 20,000 session launch**



- Memory pressure is low during session launch.
- Memory usage by most components doesn't change throughout the test as observed by the Max and Average values being nearly equal.

### Session launch observations using StoreFront

Action	Time
Authenticate	261 ms
Enumerate	1,075 ms
Total login	1,336 ms
Retrieve ICA file	2,132 ms

### Machine Creation Services Provisioning

#### Citrix Virtual Desktops Essentials MCS testing Azure Resource Manager

Machine Creation Services allows you to create and delete virtual desktops (VDA) in Azure. The first step is to create a Windows 10 VHD and then upload the VHD to Azure. The image is created from the VHD. Citrix Virtual Desktops Essentials allows you to create virtual machines from the image.

Machine count	Image copy	AD account creation	Machine creation
10	30 mins	1 min	7 mins
100	30 mins	7 mins	50 mins
250	40 mins	8 mins	2 hours
500	55 mins	15 mins	4 hours
1000	65 mins	30 mins	8 hours

Times are approximate based on several test runs and may vary.

- We tested the machine creation process using various machine counts, to measure the time required to:
  - Copy the image
  - Create machine accounts
  - Provision the machines



- The times do not increase linearly because copies of the image must be replicated to each storage account. Replication occurs in parallel, and becomes slower with more tasks
  - There is a limit of 40 machines per storage account. The limit requires 25 storage accounts for a 1,000 VM environment.
  - There is a limit of 760 machines per resource location.
- Active Directory account creation must be proxied via the Cloud Connectors, which increases the time required to complete the task. Active Directory accounts are created at a rate of approximately 33 per minute.
- Testing used Standard\_A2\_v2 Cloud Connectors. No resource bottlenecks were observed.

### Citrix Virtual Apps and Desktops service MCS testing

MCS provisioning tests were performed on a VMware ESXi 6.0 hypervisor. There are eight vSphere hosts in the cluster and share storage is NFS on a NetApp share.

OS	Machine count	Image copy	AD account creation	Machine creation
Win 2012 R2	100	4 minutes	3 mins	4 minutes
Win 2012 R2	1,000	5 minutes	30 mins	100 minutes
Win 10 32-bit	100	4 minutes	3 mins	4 minutes

Times are approximate based on multiple test runs and may vary. Test data from these runs are averaged in the table.

- The time required for the machine creation process is similar to the time required in XenApp and XenDesktop 7.x versions. The primary difference in these tests is Active Directory account creation. In the cloud environment, account creation must be proxied via the Cloud Connectors. Active Directory accounts in the cloud environment are created at a rate of approximately 33 per minute.
- We conducted the tests using two 4-vCPU, 4-GB memory VMs for the Cloud Connectors. There was no resource contention observed during the test.

## Create and manage connections

September 8, 2021

## Introduction

Configuring a connection includes selecting the connection type from among the supported hypervisors and cloud services and the storage and network you select from the *resources* for that connection.

You must be a Full Administrator to perform connection and resource management tasks.

## Where to find information about connection types

[System requirements](#) lists the supported hypervisor and cloud service versions, and includes links to host-specific articles.

## Host storage

A storage product is supported if it can be managed by a supported hypervisor. Citrix Support assists those storage product vendors in troubleshooting and resolving issues, and documents those issues in the knowledge center, as needed.

When provisioning machines, data is classified by type:

- Operating system (OS) data, which includes images.
- Temporary data, which include all non-persistent data written to MCS-provisioned machines, Windows page files, user profile data, and any data that is synchronized with Content Collaboration (formerly ShareFile). This data is discarded each time a machine restarts.

Providing separate storage for each data type can reduce load and improve IOPS performance on each storage device, making best use of the host's available resources. It also enables appropriate storage to be used for the different data types. Persistence and resilience are more important for some data than others.

- Storage can be shared (located centrally, separate from any host, used by all hosts) or local to a hypervisor. For example, central shared storage can be one or more Windows Server 2012 clustered storage volumes (with or without attached storage), or an appliance from a storage vendor. The central storage might also provide its own optimizations such as hypervisor storage control paths and direct access through partner plug-ins.
- Storing temporary data locally avoids having to traverse the network to access shared storage, and it also reduces load (IOPS) on the shared storage device. Shared storage can be more costly, so storing data locally can lower expenses. These benefits must be weighed against the availability of sufficient storage on the hypervisor servers.

## Storage shared by hypervisors

The storage shared by hypervisors method stores data that needs longer-term persistence centrally, providing centralized backup, and management. That storage holds the OS disks.

When you select this method, you can choose whether to use local storage (on servers in the same hypervisor pool) for temporary machine data. This data does not require persistence or as much resilience as the data in the shared storage. This is called the *temporary data cache*. The local disk helps reduce traffic to the main OS storage. This disk is cleared after every machine restart. The disk is accessed through a write-through memory cache. Keep in mind that if you use local storage for temporary data, the provisioned VDA is tied to a specific hypervisor host. If that host fails, the VM cannot start.

**Exception:** If you use Clustered Storage Volumes (CSV), Microsoft System Center Virtual Machine Manager does not allow temporary data cache disks to be created on local storage.

If you store temporary data locally, you can then enable and configure nondefault values for each VM's cache disk and memory size when you create a machine catalog that uses that connection. However, the default values are tailored to the connection type, and are sufficient for most cases.

The hypervisor can also provide optimization technologies through read caching of the disk images locally. For example, Citrix Hypervisor offers IntelliCache. This can also reduce network traffic to the central storage.

## Storage local to the hypervisor

The storage local to the hypervisor method stores data locally on the hypervisor. With this method, images and other OS data are transferred to all the hypervisors used in the site, both for initial machine creation and future image updates. This results in significant traffic on the management network. Image transfers are also time-consuming, and the images become available to each host at a different time.

## Create a connection and resources

### Important:

The host resources (storage and network) in your resource location must be available before you create a connection.

1. Sign in to Citrix Cloud.
2. In the upper left menu, select **My Services > Virtual Apps and Desktops**.
3. From **Manage > Full Configuration**, select **Hosting** in the left pane.
4. Select **Add Connections and Resources** in the action bar.
5. The wizard guides you through the following pages. Specific page content depends on the selected connection type. After completing each page, select **Next** until you reach the **Summary** page.

## Step 1. Connection

### Add Connection and Resources

- 1 Connection
- 2 Storage Manageme...
- 3 Storage Selection
- 4 Network
- 5 Scopes
- 6 Summary

#### Connection

Use an existing connection

Create a new connection

Connection type:

Connection address:

User name:

Password:

Zone name:

Connection name:

Create virtual machines using:

Citrix provisioning tools (Machine Creation Services or Citrix Provisioning)

Other tools

On the **Connection** page:

- To create a new connection, select **Create a new Connection**. To create a connection based on the same host configuration as an existing connection, select **Use an existing Connection** and then choose the relevant connection.
- Select the hypervisor or cloud service you are using in the **Connection type** field.
- The connection address and credentials fields differ, depending on the selected connection type. Enter the requested information.
- Enter a connection name. This name appears in the **Manage** display.
- Choose the tool to create virtual machines: Machine Creation Services or Citrix Provisioning.

Information on the **Connection** page differs depending on the host (connection type) you're using. For example, when using Azure Resource Manager, you can use an existing service principal or create a new one. For details, see the virtualization environment page listed in [System requirements](#) for your connection type.

## Step 2. Storage management

The screenshot shows a wizard window titled "Add Connection and Resources". On the left is a vertical navigation pane with six steps: 1. Connection (checked), 2. Storage Management... (selected), 3. Storage Selection, 4. Network, 5. Scopes, and 6. Summary. The main area is titled "Storage Management" and contains the following text: "Configure virtual machine storage resources for this connection. Select an optimization method for available site storage." Below this are three radio button options: "Use storage shared by hypervisors" (selected), "Optimize temporary data on available local storage" (unchecked), and "Use storage local to the hypervisor" (unchecked). Underneath is the text "Optimization technology (optional):" followed by an unchecked checkbox for "Use intellicache to reduce load on the shared storage device". At the bottom right of the window are three buttons: "Back", "Next" (highlighted in dark blue), and "Cancel".

For information about storage management types and methods, see [Host storage](#).

If you are configuring a connection to a Hyper-V or VMware host, browse to and then select a cluster name. Other connection types do not request a cluster name.

Select a storage management method: storage shared by hypervisors or storage local to the hypervisor.

- If you choose storage shared by hypervisors, indicate if you want to keep temporary data on available local storage. (You can specify nondefault temporary storage sizes in the machine catalogs that use this connection.) **Exception:** When using Clustered Storage Volumes (CSV), Microsoft System Center Virtual Machine Manager does not allow temporary data cache disks to be created on local storage. Configuring that storage management setup in the **Manage** console fails.

If you use shared storage on a Citrix Hypervisor pool, indicate if you want to use IntelliCache to reduce the load on the shared storage device. See [Citrix Hypervisor virtualization environments](#).

### Step 3. Storage selection

## Add Connection and Resources

- ✔ Connection
- ✔ Storage Manage...
- 3 Storage Selection
- 4 Network
- 5 Scopes
- 6 Summary

### Storage Selection

When using local storage, you must select the type of data to store on each local storage device; machine operating system data, temporary data, and if not storing personal user data remotely, personal user data. At least one device must be selected for each data type.

Select data storage locations:

Name ↓	OS	Temporary
Local SSD on b02u05	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Local SSD on b02u06	<input type="checkbox"/>	<input type="checkbox"/>
Local SSD on b02u07	<input type="checkbox"/>	<input type="checkbox"/>
Local SSD on b02u08	<input type="checkbox"/>	<input type="checkbox"/>
Local SSD on b02u09	<input type="checkbox"/>	<input type="checkbox"/>
Local storage on [Test] xenserver...	<input type="checkbox"/>	<input type="checkbox"/>
Local storage on [Test] xenserver...	<input type="checkbox"/>	<input type="checkbox"/>
Local storage on [Test] xenserver...	<input type="checkbox"/>	<input type="checkbox"/>
Local storage on [Test] xenserver...	<input type="checkbox"/>	<input type="checkbox"/>

Back
Next
Cancel

For more information about storage selection, see [Host storage](#).

Select at least one host storage device for each available data type. The storage management method you selected on the previous page affects which data types are available for selection on this page. You must select at least one storage device for each supported data type before you can proceed to the next page in the wizard.

The lower portion of the **Storage Selection** page contains more configuration options if you chose storage shared by hypervisors and enabled **Optimize temporary data on available local storage**. You can select which local storage devices (in the same hypervisor pool) to use for temporary data.

The number of currently selected storage devices is shown (in the graphic, “1 storage device selected”). When you hover over that entry, the selected device names appear (unless no devices are configured).

1. Select **Select** to change the storage devices to use.
2. In the **Select Storage** dialog box, select or clear the storage device check boxes, and then select **OK**.

#### Step 4. Region

(Appears only for some host types.) The region selection indicates where VMs will be deployed. Ideally, choose a region close to where users will access their applications.

#### Step 5. Network

Enter a name for the resources. This name appears in the Manage console to identify the storage and network combination associated with the connection.

Select one or more networks that the VMs will use.

Some connection types (such as Azure Resource Manager) also list subnets that VMs will use. Select one or more subnets.

#### Step 6. Summary

Review your selections; if you want to make changes, use return to previous wizard pages. When you complete your review, select **Finish**.

**Remember:** If you store temporary data locally, you can configure nondefault values for temporary data storage when you create the catalog containing machines that use this connection.

#### Edit connection settings

Do not use this procedure to rename a connection or to create a connection. Those are different operations. Change the address only if the current host machine has a new address. Entering an address to a different machine breaks the connection's machine catalogs.

You cannot change the GPU settings for a connection, because catalogs accessing this resource must use an appropriate GPU-specific image. Instead, create a new connection.

1. From **Manage > Full Configuration**, select **Hosting** in the left pane.
2. Select the connection and then select **Edit Connection** in the action bar.
3. Follow the guidance for the settings available when you edit a connection.
4. When you are finished, select **Apply** to apply any changes you made and keep the window open, or select **OK** to apply changes and close the window.

**Connection Properties** page:

- To change the connection address and credentials, select **Edit settings** and then enter the new information.
- To specify the high-availability servers for a Citrix Hypervisor connection, select **Edit HA servers**. Citrix recommends that you select all servers in the pool to allow communication with Citrix Hypervisor if the pool master fails.

### **Advanced** page:

The throttling threshold settings enable you to specify a maximum number of power actions allowed on a connection. These settings can help when power management settings allow too many or too few machines to start at the same time. Each connection type has specific default values that are appropriate for most cases. Usually, they do not need to be changed.

- The **Simultaneous actions (all types)** and **Simultaneous Personal vDisk inventory updates** settings specify two values: a maximum absolute number that can occur simultaneously on this connection, and a maximum percentage of all machines that use this connection. You must specify both absolute and percentage values. The actual limit applied is the lower of the values.

For example, in a deployment with 34 machines, if **Simultaneous actions (all types)** is set to an absolute value of 10 and a percentage value of 10, the actual limit applied is 3 (that is, 10 percent of 34 rounded to the nearest whole number, which is less than the absolute value of 10 machines).

- The **Maximum new actions per minute** is an absolute number. There is no percentage value.

Enter information in the **Connection options** field only under the guidance of a Citrix Support representative.

### **Turn maintenance mode on or off for a connection**

Turning on maintenance mode for a connection prevents any new power action from affecting any machine stored on the connection. Users cannot connect to a machine when it is in maintenance mode. If users are already connected, maintenance mode takes effect when they log off.

1. From **Manage > Full Configuration**, select **Hosting** in the left pane.
2. Select the connection. To turn maintenance mode on, select **Turn On Maintenance Mode** in the action bar. To turn maintenance mode off, select **Turn Off Maintenance Mode**.

You can also turn maintenance mode on or off for individual machines. You can turn maintenance mode on or off for machines in machine catalogs or delivery groups.

### **Delete a connection**

#### **Caution:**

Deleting a connection can result in the deletion of large numbers of machines and loss of data. Ensure that user data on affected machines is backed up or no longer required.

Before deleting a connection, ensure that:

- All users are logged off from the machines stored on the connection.
- No disconnected user sessions are running.



- Maintenance mode is turned on for pooled and dedicated machines.
- All machines in machine catalogs used by the connection are powered off.

A machine catalog becomes unusable when you delete a connection that the catalog references. If this connection is referenced by a catalog, you can delete the catalog. Before you delete a catalog, make sure it is not used by other connections.

1. From **Manage > Full Configuration**, select **Hosting** in the left pane.
2. Select the connection and then select **Delete Connection** in the action bar.
3. If this connection has machines stored on it, you are asked whether to delete the machines. If they are to be deleted, specify what to do with the associated Active Directory computer accounts.

### Rename or test a connection

1. From **Manage > Full Configuration**, select **Hosting** in the left pane.
2. Select the connection and then select **Rename Connection** or **Test Connection** in the action bar.

### View machine details on a connection

1. From **Manage > Full Configuration**, select **Hosting** in the left pane.
2. Select the connection and then select **View Machines** in the action bar.

The upper pane lists the machines accessed through the connection. Select a machine to view its details in the lower pane. Session details are also provided for open sessions.

Use the search feature to find machines quickly. Either select a saved search from the list at the top of the window, or create a new search. You can either search by typing all or part of the machine name, or you can build an expression to use for an advanced search. To build an expression, select **Unfold**, and then select from the lists of properties and operators.

### Manage machines on a connection

1. From **Manage > Full Configuration**, select **Hosting** in the left pane.
2. Select a connection and then select **View Machines** in the action bar.
3. Select one of the following in the action bar. Some actions might not be available, depending on the machine state and the connection host type.
  - **Start:** Starts the machine if it is powered off or suspended.
  - **Suspend:** Pauses the machine without shutting it down, and refreshes the list of machines.

- **Shut down:** Requests the operating system to shut down.
- **Force shut down:** Forcibly powers off the machine, and refreshes the list of machines.
- **Restart:** Requests the operating system to shut down and then start the machine again. If the operating system cannot comply, the desktop remains in its current state.
- **Enable maintenance mode:** Temporarily stops connections to a machine. Users cannot connect to a machine in this state. If users are connected, maintenance mode takes effect when they log off. (You can also turn maintenance mode on or off for all machines accessed through a connection, as described earlier.)
- **Remove from Delivery Group:** Removing a machine from a Delivery Group does not delete it from the machine catalog that the Delivery Group uses. You can remove a machine only when no user is connected to it. Turn on maintenance mode to temporarily prevent users from connecting while you are removing the machine.
- **Delete:** When you delete a machine, users no longer have access to it, and the machine is deleted from the machine catalog. Before deleting a machine, ensure that all user data is backed up or no longer required. You can delete a machine only when no user is connected to it. Turn on maintenance mode to temporarily stop users from connecting while you are deleting the machine.

For actions that involve machine shutdown, if the machine does not shut down within 10 minutes, it is powered off. If Windows attempts to install updates during shutdown, there is a risk that the machine will be powered off before the updates are complete.

## Edit storage

You can display the status of servers that are used to store operating system, temporary, and personal (PvD) data for VMs that use a connection. You can also specify which servers to use for storage of each data type.

1. From **Manage > Full Configuration**, select **Hosting** in the left pane.
2. Select the connection and then select **Edit Storage** in the action bar.
3. In the left pane, select the data type: operating system or temporary.
4. Select or clear the check boxes for one or more storage devices for the selected data type.
5. Select **OK**.

Each storage device in the list includes its name and storage status. Valid storage status values are:

- **In use:** The storage is being used for creating machines.
- **Superseded:** The storage is being used only for existing machines. No new machines are added in this storage.
- **Not in use:** The storage is not being used for creating machines.

If you clear the check box for a device that is currently **In use**, its status changes to **Superseded**. Existing machines will continue to use that storage device (and can write data to it). So, that location can

become full even after it stops being used for creating machines.

### Delete, rename, or test resources

1. From **Manage > Full Configuration**, select **Hosting** in the left pane.
2. Select the resource and then select the appropriate entry in the action bar: **Delete Resources**, **Rename Resources**, or **Test Resources**.

### Connection timers

You can use Citrix policy settings to configure three connection timers:

- **Maximum connection timer:** Determines the maximum duration of an uninterrupted connection between a user device and a virtual desktop. Use the **Session connection timer** and **Session connection timer interval** policy settings.
- **Connection idle timer:** Determines how long an uninterrupted user device connection to a virtual desktop is maintained if there is no input from the user. Use the **Session idle timer** and **Session idle timer interval** policy settings.
- **Disconnect timer:** Determines how long a disconnected, locked virtual desktop can remain locked before the session is logged off. Use the **Disconnected session timer** and **Disconnected session timer interval** policy settings.

When you update any of these settings, ensure they are consistent across your deployment.

See the policy settings documentation for more information.

### Where to go next

If you're in the initial deployment process, [create a machine catalog](#).

## Install VDAs

September 8, 2021

### Introduction

This article begins with a description of Windows VDAs and the available VDA installers. The remainder of the article describes the steps in the VDA installation wizard. Command-line equivalents are provided. For details, see [Install VDAs using the command line](#).

For information about Linux VDAs, see [Linux Virtual Delivery Agent](#).

View an introduction to VDAs.

[This is an embedded video. Click the link to watch the video](#)

## Installation considerations

The [Citrix Virtual Apps and Desktops service](#) article describes what VDAs are and what they do. Here's more information.

- **Analytics collection:** Analytics are collected automatically when you install or upgrade components. By default, that data is uploaded to Citrix automatically when the installation completes. Also, when you install components, you are automatically enrolled in the [Citrix Customer Experience Improvement Program \(CEIP\)](#), which uploads anonymous data. Also, during an installation or upgrade, you're offered the opportunity to enroll in Call Home.

If a VDA installation fails, an MSI analyzer parses the failing MSI log, displaying the exact error code. The analyzer suggests a CTX article, if it is a known issue. The analyzer also collects anonymized data about the failure error code. This data is included with other data collected by CEIP. If you end enrollment in CEIP, the collected MSI analyzer data is no longer sent to Citrix.

For information about these programs, see [Citrix Insight Services](#).

- **Citrix Workspace app:** Citrix Workspace app for Windows is not installed by default when you install a VDA. You can download and install or upgrade Citrix Workspace app for Windows and other Citrix Workspace apps from the Citrix website. Alternatively, you can make those Citrix Workspace apps available from the Workspace or a StoreFront server.
- **Print Spooler Service:** The Print Spooler Service is enabled by default on supported Windows Servers. If you disable this service, you cannot successfully install a VDA for Multi-session OS, so ensure that this service is enabled before installing a VDA.
- **Microsoft Media Foundation:** Most supported Windows editions come with Media Foundation already installed. If the machine on which you're installing a VDA does not have Microsoft Media Foundation (such as N editions), several multimedia features are not installed and do not work.
  - Flash Redirection
  - Windows Media Redirection
  - HTML5 Video Redirection
  - HDX RealTime Webcam Redirection

You can acknowledge the limitation, or end the VDA installation and restart it later, after installing Media Foundation. In the graphical interface, this choice is presented in a message. In the command line, you can use the `/no_mediafoundation_ack` option to acknowledge the limitation.

- **Local user group:** When you install the VDA, a new local user group called Direct Access Users is created automatically. For a single-session OS VDA, this group applies only to RDP connections. For a multi-session OS VDA, this group applies to ICA and RDP connections.
- **Cloud Connector address requirement:** The VDA must have at least one valid Cloud Connector address (in the same resource location) with which to communicate. Otherwise, sessions cannot be established. You specify Cloud Connector addresses when you install the VDA. For information about other ways to specify Cloud Connector addresses where VDAs can register, see [VDA registration](#).
- **Operating system considerations:**
  - Review the [System requirements](#) for supported platforms, operating systems, and versions.
  - Ensure that each operating system maintains the latest updates.
  - Ensure that VDAs have synchronized system clocks. The Kerberos infrastructure that secures communication between the machines requires synchronization.
  - Optimization guidance for Windows 10 machines is available in [CTX216252](#).
  - If you try to install (or upgrade to) a Windows VDA on an OS that is not supported for that VDA version, a message describes your options. For example, if you try to install the latest VDA on a Windows 7 machine, a message guides you to [CTX139030](#). For more information, see [Earlier operating systems](#).
- **Installed MSIs:** Several MSIs are installed automatically when you install a VDA. You can prevent the installation of some of them on the **Additional Components** page of the graphical interface or with the `/exclude` option in the CLI. For others, the only way to prevent their installation is with the `/exclude` CLI option.
- **Domain-joined:** Ensure that the machine is domain-joined before installing the VDA software.

### VDA supportability tools

Each VDA installer includes a supportability MSI that contains Citrix tools for checking the VDA's performance, such as its overall health and the quality of connections. Enable or disable installation of this MSI on the **Additional Components** page of the VDA installer's graphical interface. From the command line, disable installation with the `/exclude "Citrix Supportability Tools"` option.

By default, the supportability MSI is installed in `C:\Program Files (x86)\Citrix\Supportability Tools\`. You can change this location on the **Components** page of the VDA installer's graphical interface, or with the `/installdir` command-line option. Keep in mind that changing the location changes it for all installed VDA components, not just the supportability tools.

Current tools in the supportability MSI:

- Citrix Health Assistant: For details, see [CTX207624](#).

- VDA Cleanup Utility: For details, see [CTX209255](#).

If you do not install the tools when you install the VDA, the CTX article contains a link to the current download package.

### Restarts during VDA installation

A restart is required at the end of the VDA installation. That restart occurs automatically by default.

To minimize the number of other restarts needed during VDA installation:

- Ensure that a supported Microsoft .NET Framework version is installed before beginning the VDA installation.
- For Windows multi-session OS machines, install and enable the RDS role services before installing the VDA.

If you do not install those prerequisites before installing the VDA:

- If you are using the graphical interface or the command line interface without the `/noreboot` option, the machine restarts automatically after installing the prerequisite.
- If you are using the command line interface with the `/noreboot` option, you must initiate the restart.

After each restart, the VDA installation continues. If you're installing from the command line, you can prevent the automatic resumption with the `/noresume` option.

When upgrading a VDA to version 7.17 or a later supported version, a restart occurs during the upgrade. This restart cannot be avoided.

### Restore on install or upgrade failure

#### Note:

This feature is available only for single-session VDAs.

If a single-session VDA installation or upgrade fails, and the “restore on failure” feature is enabled, the machine is returned to a restore point that was set before the installation or upgrade began.

When a single-session VDA installation or upgrade starts with this feature enabled, the installer creates a system restore point before beginning the actual install or upgrade. If the VDA installation or upgrade fails, the machine is returned to the restore point state. The `%temp%/Citrix` folder contains deployment logs and other information about the restore.

By default, this feature is disabled.

If you plan to enable this feature, make sure that system restore is not disabled through a GPO setting (Computer Configuration > Administrative Templates > System > System Restore).

To enable this feature when installing or upgrading a single-session VDA:

- When using a VDA installer's graphical interface (such as using **Autostart** or the `XenDesktopVDASetup.exe` command without any restore or quiet options), select the **Enable automatic restore if update fails** check box on the **Summary** page.

If the install/upgrade completes successfully, the restore point is not used, but is retained.

- Run a VDA installer with either the `/enablerestore` or `/enablerestorecleanup` option.
  - If you use the `/enablerestorecleanup` option, and the install/upgrade completes successfully, the restore point is removed automatically.
  - If you use the `/enablerestore` option, and the install/upgrade completes successfully, the restore point is not used, but is retained.

## VDA installers

VDA installers can be downloaded directly from the Citrix Cloud console.

By default, files in the self-extracting installers are extracted to the `Temp` folder. The files extracted to the `Temp` folder are automatically deleted after the installation completes. Alternatively, you can use the `/extract` command with an absolute path.

Three standalone VDA installers are available for download.

**VDA Server Setup.exe** Installs a multi-session OS VDA.

**VDA Workstation Setup.exe** Installs a single-session OS VDA.

**VDA Workstation Core Setup.exe** Installs a single-session OS VDA that is optimized for Remote PC Access deployments or core VDI installations. Remote PC Access uses physical machines. Core VDI installations are VMs that are not being used as an image. This installer deploys only the core services necessary for VDA connections. Therefore, it supports only a subset of the options that are valid with the `VDA Workstation Setup` installer.

This installer for the current release does not install or contain the components used for:

- App-V.
- Profile Management. Excluding Citrix Profile Management from the installation affects Monitor displays.
- Machine Identity Service.
- Citrix Workspace app for Windows.
- Citrix Supportability Tools.
- Citrix Files for Windows.
- Citrix Files for Outlook.
- MCSIO write cache for storage optimization.

This installer does not install or contain a Citrix Workspace app for Windows.

This installer automatically installs the Browser Content Redirection MSI. Automatic installation applies to VDA release 2003 and later supported releases.

Using `VDAWorkstationCoreSetup.exe` is equivalent to using the `VDAWorkstationSetup.exe` installer to install a single-session OS VDA and either:

- In the graphical interface: Selecting the **Remote PC Access** option on the **Environment** page.
- In the command-line interface: Specifying the `/remotepc` option.
- In the command line interface: Specifying `/components vda` and `/exclude "Citrix Personalization for App-V - VDA""Personal vDisk""Machine Identity Service""Citrix Profile Management""Citrix Profile Management WMI Plugin""Citrix Supportability Tools""Citrix Files for Windows""Citrix Files for Outlook""Citrix MCS IODriver"`.

If you install a VDA with the `VDAWorkstationCoreSetup.exe` installer and later upgrade that VDA using the `VDAWorkstationSetup.exe` installer, you can optionally install the omitted components and features.

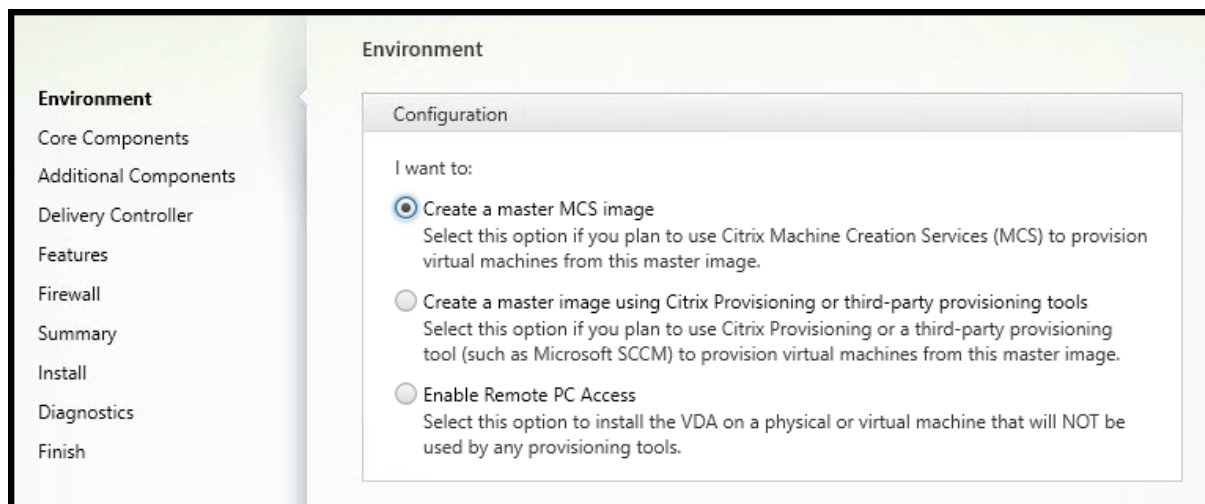
### Step 1. Download the product software and launch the wizard

1. On the machine where you're installing the VDA, sign in to [Citrix Cloud](#).
2. In the upper left menu, select the Citrix service in the **My Services** list.
3. On the right side, click **Downloads** and select **Download VDA**. You are redirected to the VDA download page. Find the VDA installer you want and then select **Download File**.
4. After the download completes, right-click the file and select **Run as administrator**. The installation wizard launches.

As an alternative to steps 1-3, you can download the VDA directly from the [Citrix download page](#).



## Step 2. Specify how the VDA will be used



On the **Environment** page, specify how you plan to use the VDA, indicating whether you'll use this machine as an image to provision machines. The option you choose affects which Citrix provisioning tools are installed automatically (if any), and the default values on the **Additional Components** page of the VDA installer.

Choose one of the following:

- **Create a master MCS image:** Select this option to install a VDA on a VM image, if you plan to use Machine Creation Services to provision VMs. This option installs the Machine Identity Service. This is the default option.

Command-line option: `/mastermcsimage` or `/masterimage`

- **Create a master image using Citrix Provisioning or third-party provisioning tools:** Select this option to install a VDA on a VM image, if you plan to use Citrix Provisioning or third-party provisioning tools (such as Microsoft System Center Configuration Manager). Use this option for previously provisioned VMs that were booted from a Citrix Provisioning read/write disk.

Command-line option: `/masterpvsimage`

- (Appears only on multi-session OS machines) **Enable brokered connections to a server:** Select this option to install a VDA on a physical or virtual machine that will not be used as an image.

Command-line option: `/remotepc`

- (Appears only on multi-session OS machines) **Enable Remote PC Access:** Select this option to install a VDA on a physical machine for use with Remote PC Access.

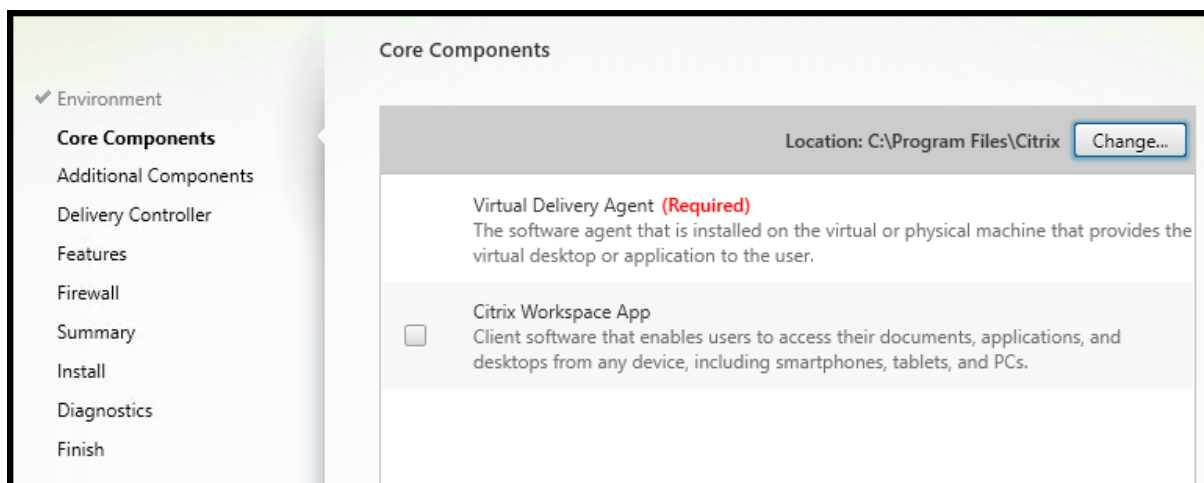
Command-line option: `/remotepc`

Select **Next**.

This page does not appear:

- When upgrading a VDA.
- When using the `VDAWorkstationCoreSetup.exe` installer.

### Step 3. Select the components to install and the installation location



On the **Core components** page:

- **Location:** By default, components are installed in `C:\Program Files\Citrix`. This default is fine for most deployments. If you specify a different location, that location must have execute permissions for the network service.

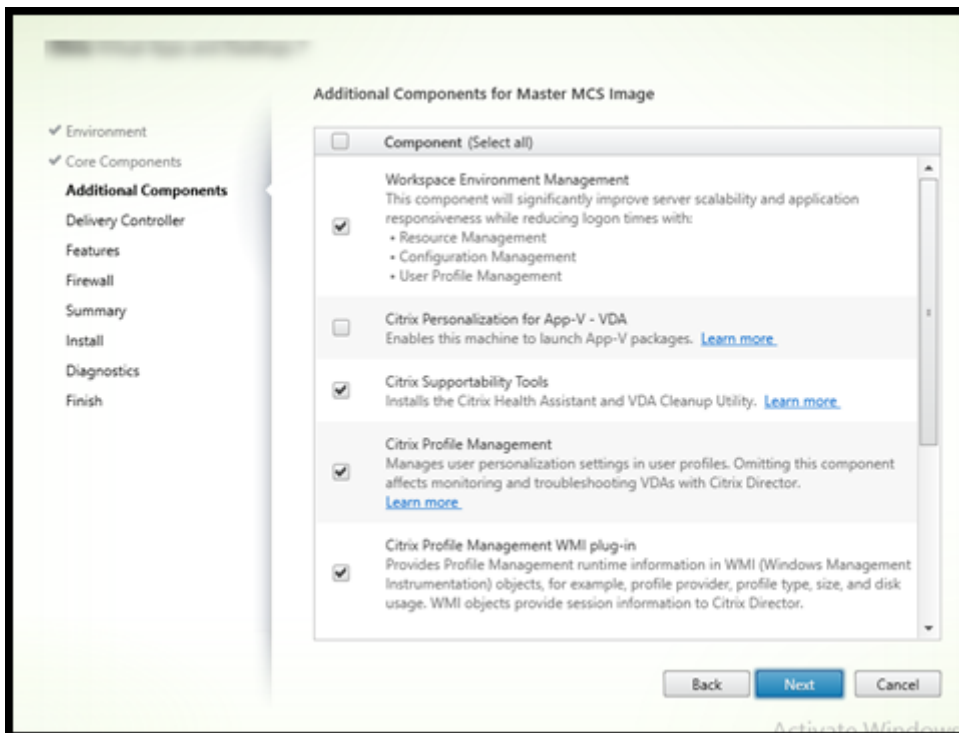
Command-line option: `/installdir`

- **Components:** By default, Citrix Workspace app for Windows is not installed with the VDA. If you are using the `VDAWorkstationCoreSetup.exe` installer, Citrix Workspace app for Windows is never installed, so this check box is not displayed.

Command-line option: `/components vda,plugin` to install the VDA and the Citrix Workspace app for Windows

Select **Next**.

### Step 4. Install additional components



The **Additional Components** page contains check boxes to enable or disable installation of other features and technologies with the VDA. In a command-line installation, you can use the `/exclude` or `/includeadditional` option to omit or include one or more available components.

The following table indicates the default setting of items on this page. The default setting depends on the option you selected on the **Environment** page.

Additional Components page	Environment page: "Master image with MCS" or "Master image with Citrix Provisioning ..." selected	Environment page: "Enable brokered connections to server" (for multi-session OS) or "Remote PC Access" (for single-session OS) selected
Workspace Environment Management	Not selected	Not selected
Citrix Personalization for App-V	Not selected	Not selected
User Personalization Layer	Not selected	Not shown because it's not valid for this use case
Citrix Supportability tools	Selected	Not selected
Citrix Profile Management	Selected	Not selected

Additional Components page	Environment page: “Master image with MCS” or “Master image with Citrix Provisioning ...” selected	Environment page: “Enable brokered connections to server” (for multi-session OS) or “Remote PC Access” (for single-session OS) selected
Citrix Profile Management WMI Plug-in	Selected	Not selected
Citrix Files for Windows	Not selected	Not selected
Citrix Files for Outlook	Not selected	Not selected
MCSIO write cache for storage optimization	Not selected	Not selected

This page does not appear when:

- Using the `VDAWorkstationCoreSetup.exe` installer. Also, the command-line options for the additional components are not valid with that installer.
- Upgrading a VDA and all the additional components are already installed. If some of the additional components are already installed, the page lists only the components that are not installed.

The components list can include:

- **Workspace Environment Management:** Install this component if your deployment uses Citrix Workspace Environment Management. For details, see [Workspace Environment Management](#).

Command-line options: `/includeadditional "Citrix WEM Agent"` to enable component installation, `/exclude "Citrix WEM Agent"` to prevent component installation

- **Citrix Personalization for App-V:** Install this component if you use applications from Microsoft App-V packages. For details, see [App-V](#).

Command-line option: `/includeadditional "Citrix Personalization for App-V - VDA"` to enable component installation, `/exclude "Citrix Personalization for App-V - VDA"` to prevent component installation

- **Citrix User Personalization Layer:** Installs the MSI for the user personalization layer. For details, see [User personalization layer](#).

This component appears only when installing a VDA on a single-session Windows 10 machine.

Command-line option: `/includeadditional "User Personalization Layer"` to enable component installation, `/exclude "User Personalization Layer"` to prevent component installation

- **Citrix Supportability Tools:** Installs the MSI that contains Citrix supportability tools.

Command-line option: `/includeadditional "Citrix Supportability Tools"` to enable component installation, `/exclude "Citrix Supportability Tools"` to prevent component installation

- **Citrix Profile Management:** This component manages user personalization settings in user profiles. For details, see [Profile Management](#).

Excluding Citrix Profile Management from the installation affects the monitoring and troubleshooting of VDAs in Citrix Cloud.

- On the **User details** and **EndPoint** pages of the **Monitor** tab, the **Personalization** panel and the **Logon Duration** panel fail.
- On the **Dashboard** and **Trends** pages, the **Average Logon Duration** panel display data only for machines that have Profile Management installed.

Even if you are using a third-party user profile management solution, Citrix recommends that you install and run the Citrix Profile Management Service. Enabling the Citrix Profile Management Service is not required.

Command-line option: `/includeadditional "Citrix Profile Management"` to enable component installation, `/exclude "Citrix Profile Management"` to prevent component installation

- **Citrix Profile Management WMI Plug-in:** This plug-in provides Profile Management runtime information in WMI (Windows Management Instrumentation) objects (for example, profile provider, profile type, size, and disk usage). WMI objects provide session information to Director.

Command-line option: `/includeadditional "Citrix Profile Management WMI Plug-in"` to enable component installation, `/exclude "Citrix Profile Management WMI Plug-in"` to prevent component installation

- **Citrix Files for Windows:** This component enables users to connect to their Citrix Files account. They can then interact with Citrix Files through a mapped drive in the Windows file system, without requiring a full sync of their content.

Command-line options: `/includeadditional "Citrix Files for Windows"` to enable component installation, `/exclude "Citrix Files for Windows"` to prevent component installation

- **Citrix Files for Outlook:** This component allows you to bypass file size restrictions and add security to your attachments or emails by sending them through Citrix Files. You can provide a secure file upload request directly in your email. For more information, see [Citrix Files for Outlook](#).

Command-line options: `/includeadditional "Citrix Files for Outlook"` to enable component installation, `/exclude "Citrix Files for Outlook"` to prevent component installation

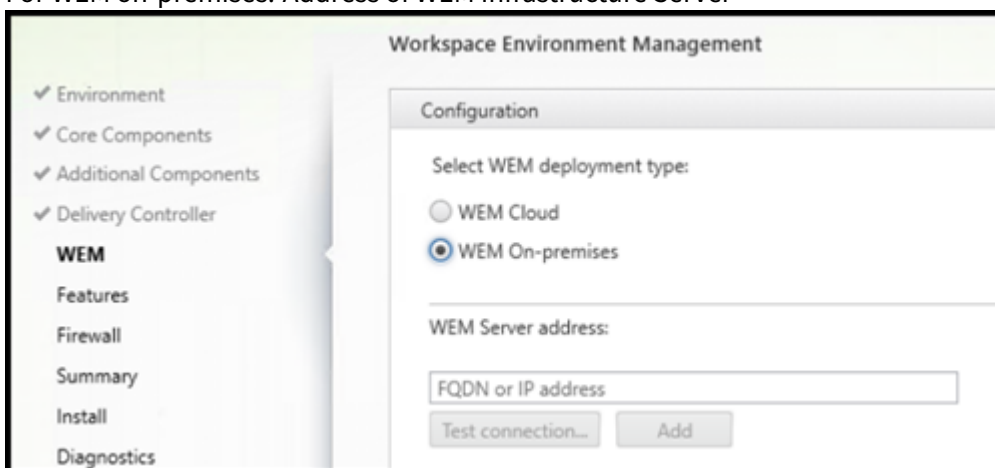
- **MCSIO write cache for storage optimization:** Installs the Citrix MCS IO driver. For more information, see [Storage shared by hypervisors](#) and [Configure cache for temporary data](#).

Command-line options: `/includeadditional "Citrix MCS IODriver"` to enable component installation, `/exclude "Citrix MCS IODriver"` to prevent component installation

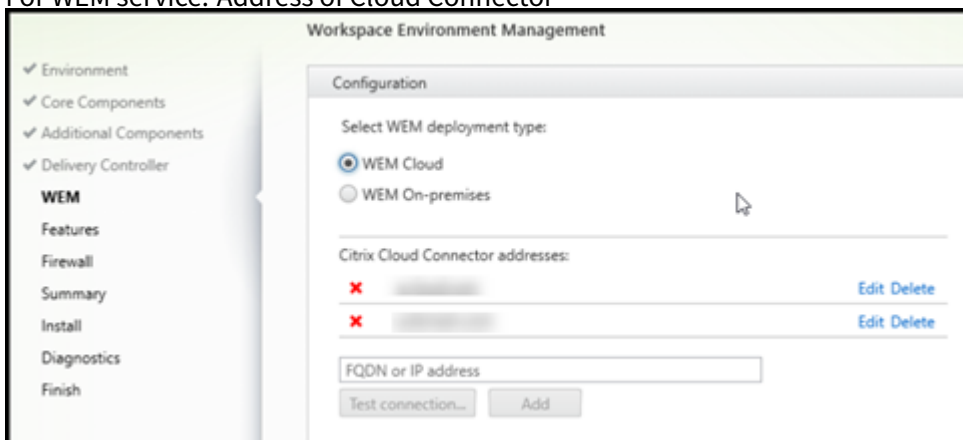
## Step 5. Workspace Environment Management

The **Workspace Environment Management (WEM)** page appears only when you enable the corresponding check box on the **Additional Components** page.

1. Select whether you have a WEM on-premises or WEM cloud (service) deployment.
2. Add an FQDN or IP address. Special characters are ignored.
  - For WEM on-premises: Address of WEM Infrastructure Server



- For WEM service: Address of Cloud Connector



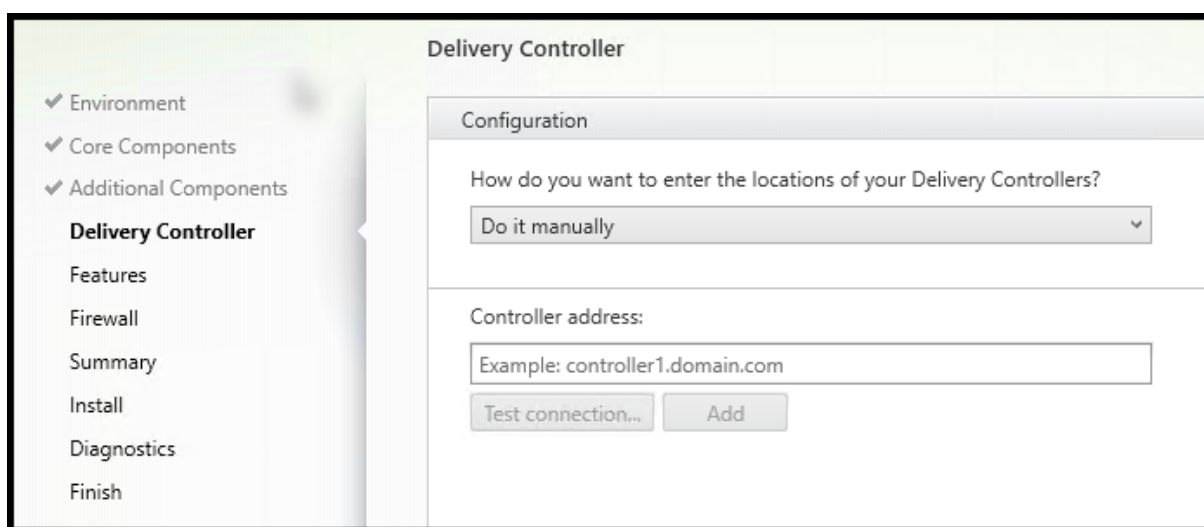
3. Select **Test Connection**. The port must be open in the firewall.

4. After a successful connection test, select **Add**.
5. For a WEM service deployment, repeat steps 2-4 for other Cloud Connectors.
6. Select **Next**.

Command-line options: `/wem_server`, `/wem_cloud_connectors`

More WEM agent configuration options are available in the [command-line interface](#).

## Step 6. Cloud Connector addresses



On the **Delivery Controller** page, select **Do it manually**. Enter the DNS name of an installed Cloud Connector and then select **Add**. If you've installed additional Cloud Connectors in the resource location, add their DNS names.


Select **Next**.

Considerations:

- The address can contain only alphanumeric characters.
- Successful VDA registration requires that the firewall ports used to communicate with the Cloud Connector are open. That action is enabled by default on the **Firewall** page of the wizard.

Command-line option: `/controllers`

## Step 7. Enable or disable features



On the **Features** page, use the check boxes to enable or disable features you want to use.

- **Use Windows Remote Assistance:** When this feature is enabled, Windows Remote Assistance is used with the user shadowing feature of the Director component in Citrix Cloud. Windows Remote Assistance opens the dynamic ports in the firewall. (Default = disabled)

Command-line option: `/enable_remote_assistance`

- **Use Real-Time Audio Transport for audio:** Enable this feature if voice-over-IP is widely used in your network. The feature reduces latency and improves audio resilience over lossy networks. It allows audio data to be transmitted using RTP over UDP transport. (Default = disabled)


Command-line option: `/enable_real_time_transport`

Select **Next**.

If this page contains a feature named **MCS I/O**, do not use it. The MCS IO feature is configured on the **Additional Components** page.



## Step 8. Firewall ports



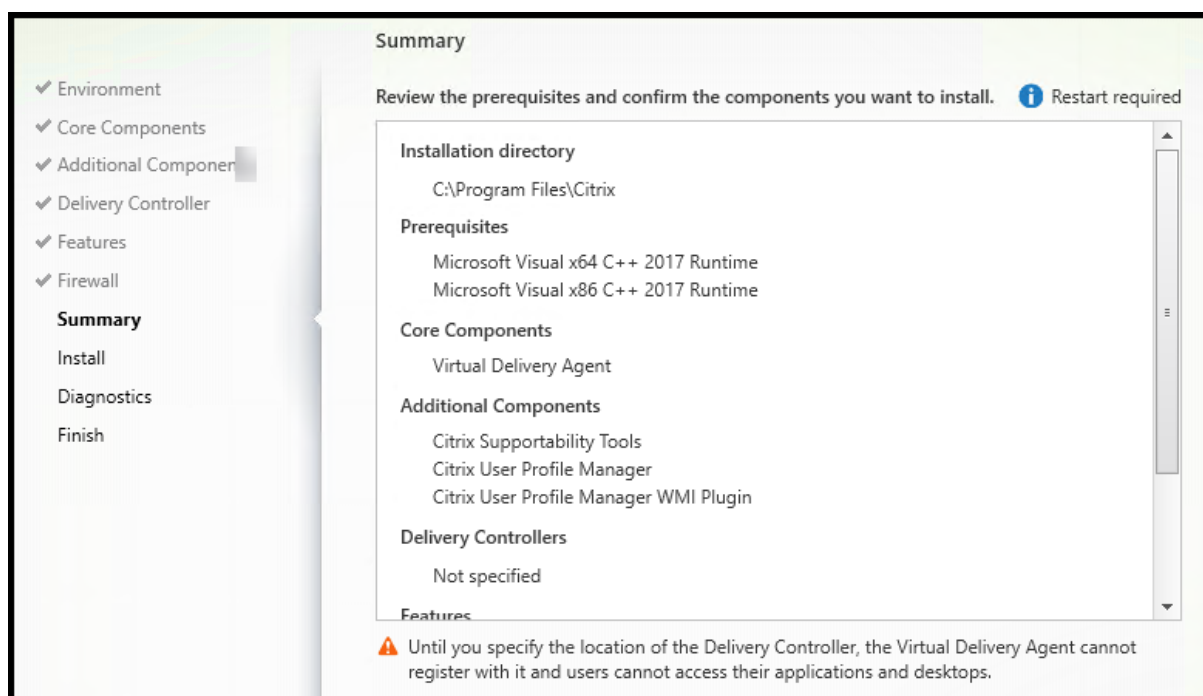
The **Firewall** page indicates which ports the VDA and Cloud Connectors use to communicate with each other. By default, these ports are opened automatically if the Windows Firewall Service is running, even if the firewall is not enabled. This default setting is fine for most deployments.

For port information, see [Network ports](#).

Select **Next**.

Command-line option: `/enable_hdx_ports`

## Step 9. Review prerequisites and confirm installation




The **Summary** page lists what will be installed. You can return to earlier wizard pages and change selections, if needed.

(Single-session VDAs only) Select the **Enable automatic restore if update fails** check box to enable the restore on failure feature. For details, see [Restore on install or upgrade failure](#).

When you're ready, select **Install**.

## Step 10. Diagnose



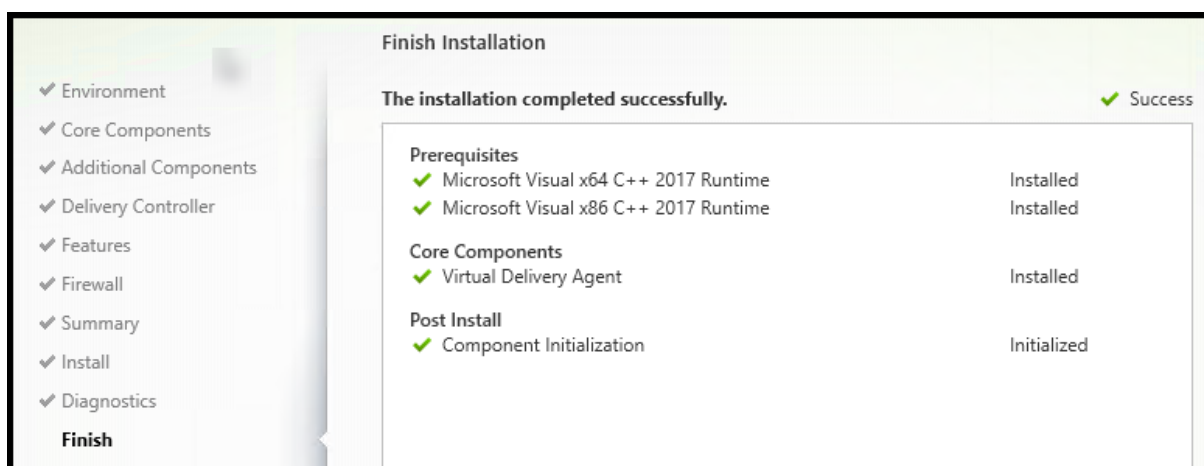
On the **Diagnostics** page, choose whether to participate in Citrix Call Home. If you choose to partici-

pate (the default), select **Connect**. When prompted, enter your Citrix account credentials.

After your credentials are validated (or if you choose not to participate), select **Next**.

For more information, see [Call Home](#).

## Step 11. Complete this installation



The **Finish** page contains green check marks for all prerequisites and components that installed and initialized successfully.

Select **Finish**. By default, the machine restarts automatically. Although you can disable this automatic restart, the VDA cannot be used until the machine restarts.

If you are installing a VDA on individual machines (rather than an image), repeat the steps above to install a VDA on other machines, as needed.

## Troubleshoot

In the **Manage > Full Configuration** display for a delivery group, the **Installed VDA version** entry in the details pane might not be the version installed on the machines. The machine's Windows Programs and Features display shows the actual VDA version.

## Citrix Optimizer

Citrix Optimizer is a tool for Windows OS that helps Citrix administrators optimize VDAs by removing and optimizing various components.

After installing a VDA and completing the final restart, download and install Citrix Optimizer. See [CTX224676](#). The CTX article contains the download package, plus instructions about installing and using Citrix Optimizer.

## Customize a VDA

Later, to customize (change information for) an installed VDA:

1. From the Windows feature for removing or changing programs, select **Citrix Virtual Delivery Agent** or **Citrix Remote PC Access/VDI Core Services VDA**. Then right-click and select **Change**.
2. Select **Customize Virtual Delivery Agent Settings**.

When the installer launches, change any available settings.

## Other information

- After you install a VDA, you can check the health and availability of the site and its components with a [Cloud Health check](#).

## Where to go next

[Create machine catalogs](#).

To review the entire configuration process, see [Plan and build a deployment](#).

## Install VDAs using the command line

July 6, 2021

### Introduction

This article applies to installing, upgrading, and customizing Virtual Delivery Agents (VDAs) on machines with Windows operating systems.

This article describes how to issue VDA installation commands. Before beginning an installation, review [Install VDAs](#) to learn about installation considerations, installers, and what you specify during installation.

### Install a VDA from the command line

Ensure that the machine is domain-joined before installing the VDA software.

To install a VDA (and see command execution progress and return values), you must have elevated administrative privileges or use **Run as administrator**.

1. On the machine where you're installing the VDA, sign to [Citrix Cloud](#).

2. In the upper left menu, select **My Services > Virtual Apps and Desktops**.
3. On the upper right side, click **Downloads** and select **Download VDA**. You are redirected to the [VDA download page](#). Find the VDA installer you want and click **Download File**.
4. After the download completes, run its name. Use the options described in this article.
  - For the multi-session OS Virtual Delivery Agent, run `VDAserverSetup.exe`
  - For the single-session OS Virtual Delivery Agent, run `VDAWorkstationSetup.exe`
  - For the single-session OS Core Services Virtual Delivery Agent, run `VDAWorkstationCoreSetup.exe`

To extract the files before installing them, use `/extract` with the absolute path, for example `.\VDAWorkstationCoreSetup.exe /extract %temp%\CitrixVDAInstallMedia`. (The directory must exist. Otherwise, the extract fails.) Then in a separate command, run the appropriate command, using the valid options listed in this article.

- For `VDAserverSetup_XXXX.exe`, run `<extract folder>\Extract\Image-Full\x64\XenDesktop Setup\XenDesktopVDASetup.exe`
- For `VDAWorkstationCoreSetup_XXXX.exe`, run `<extract folder>\Extract\Image-Full\x64\XenDesktop Setup\XenDesktopRemotePCSetup.exe`
- For `VDAWorkstationSetup_XXXX.exe`, run `<extract folder>\Extract\Image-Full\x64\XenDesktop Setup\XenDesktopVDASetup.exe`

## Command-line options to install a VDA

The following options are valid with one or more of the commands: `VDAserverSetup.exe`, `VDAWorkstationSetup.exe`, and `VDAWorkstationCoreSetup.exe`.

- **/components** *component[,component]*

Comma-separated list of components to install or remove. Valid values are:

- **VDA:** Virtual Delivery Agent
- **PLUGINS:** Citrix Workspace app for Windows

To install the VDA and Citrix Workspace app, specify `/components vda,plugins`.

If this option is omitted, only the VDA is installed (not the Citrix Workspace app).

This option is not valid when using the `VDAWorkstationCoreSetup.exe` installer. That installer cannot install Citrix Workspace app.

- **/controllers** “*controller [controller]...*”

Space-separated FQDNs of Citrix Cloud Connectors with which the VDA can communicate, enclosed in straight quotation marks. Do not specify both the `/site_guid` and `/controllers` options.

- **`/disableexperiencemetrics`**

Prevents the automatic upload of analytics collected during installation, upgrade, or removal to Citrix.

- **`/enable_hdx_ports`**

Opens ports in the Windows firewall required by the VDA and enabled features (except Windows Remote Assistance), if the Windows Firewall Service is detected, even if the firewall is not enabled. If you are using a different firewall or no firewall, you must configure the firewall manually. For port information, see [Network ports](#).

To open the UDP ports that HDX adaptive transport, specify the `/enable_hdx_udp_ports` option, in addition to the `/enable_hdx_ports` option.

- **`/enable_hdx_udp_ports`**

Opens UDP ports in the Windows firewall that HDX adaptive transport requires, if the Windows Firewall Service is detected, even if the firewall is not enabled. If you are using a different firewall or no firewall, you must configure the firewall manually. For port information, see [Network ports](#).

To open the ports that the VDA uses, specify the `/enable_hdx_ports` option, in addition to the `/enable_hdx_udp_ports` option.

- **`/enable_real_time_transport`**

Enables or disables use of UDP for audio packets (RealTime Audio Transport for audio). Enabling this feature can improve audio performance. Include the `/enable_hdx_ports` option if you want the UDP ports opened automatically when the Windows Firewall Service is detected.

- **`/enable_remote_assistance`**

Enables the shadowing feature in Windows Remote Assistance for use with the **Monitor** functions. If you specify this option, Windows Remote Assistance opens the dynamic ports in the firewall.

- **`/enablerestore` or `/enablerestorecleanup`**

(Valid only for single-session VDAs) Enables automatic return to the restore point, if the VDA install or upgrade fails.

If the install/upgrade completes successfully:

- `/enablerestorecleanup` instructs the installer to remove the restore point.

- `/enablerestore` instructs the installer to retain the restore point, even though it was not used.

For details, see [Restore on install or upgrade failure](#).

- **`/exclude`** *"component" [, "component"]*

Prevents installation of one or more comma-separated optional components, each enclosed in straight quotation marks. For example, installing or upgrading a VDA on an image not managed by MCS does not need the Machine Identity Service component. Valid values are:

- Machine Identity Service
- Citrix Profile Management
- Citrix Profile Management WMI Plug-in
- Citrix Universal Print Client
- Citrix Telemetry Service
- Citrix Personalization for App-V - VDA
- Citrix Supportability Tools
- Citrix Files for Windows
- Citrix Files for Outlook
- User personalization layer
- Citrix WEM Agent
- Citrix MCS IODriver

Excluding Citrix Profile Management from the installation (`/exclude "Citrix Profile Management"`) affects monitoring and troubleshooting of VDAs from the **Monitor** tab. On the **User details** and **EndPoint** pages, the Personalization panel and the Logon Duration panel fail. On the **Dashboard** and **Trends** pages, the Average Logon Duration panel display data only for machines that have Profile Management installed.

Even if you are using a third-party Profile Management solution, Citrix recommends that you install and run the Citrix Profile Management Service. Enabling the Citrix Profile Management Service is not required.

If you plan to use MCS to provision VMs, do not exclude the Machine Identity Service.

If you specify both `/exclude` and `/includeadditional` with the same component name, the component is not installed.

This option is not valid when using the `VDAWorkstationCoreSetup.exe` installer. That installer automatically excludes many of these items.

- **`/h` or `/help`**

Displays command help.

- **`/includeadditional`** *component" [, "component"] ...*

Includes installation of one or more comma-separated optional components, each enclosed in straight quotation marks. The component names are case-sensitive.

This option can be helpful when you are creating a Remote PC Access deployment, and want to install components that are not included by default. Valid values are:

- Citrix Profile Management
- Citrix Profile Management WMI Plug-in
- Citrix Universal Print Client
- Citrix Telemetry Service
- Citrix Personalization for App-V - VDA
- Citrix Supportability Tools
- Citrix Files for Windows
- Citrix Files for Outlook
- User personalization layer
- Citrix WEM Agent
- Citrix MCS IODriver

If you specify both `/exclude` and `/includeadditional` with the same component name, that component is not installed.

- **`/installdir`** *directory*

Existing empty directory where components will be installed. Default = `c:\Program Files\Citrix`.

- **`/install_mcsio_driver`**

Do not use. Instead, use `/includeadditional "Citrix MCS IODriver"` or `/exclude Citrix MCS IODriver"`

- **`/logpath`** *path*

Log file location. The specified folder must exist. The installer does not create it. Default = `"%TEMP%\Citrix\XenDesktop Installer"`

This option is not available in the graphical interface.

- **`/masterimage`**

Valid only when installing a VDA on a VM. Sets up the VDA as a master image. This option is equivalent to `/mastermcsimage`.

This option is not valid when using the `VDAWorkstationCoreSetup.exe` installer.

- **`/mastermcsimage`**

Specifies that this machine will be used as a master image with Machine Creation Services. This option is equivalent to `/masterimage`.



- **`/masterpvimage`**

Specifies that this machine will be used as a master image with either Citrix Provisioning or a third-party provisioning tool (such as Microsoft System Center Configuration Manager).

- **`/no_mediafoundation_ack`**

Acknowledges that Microsoft Media Foundation is not installed, and several HDX multimedia features will not be installed and will not work. If this option is omitted and Media Foundation is not installed, the VDA installation fails. Most supported Windows editions come with Media Foundation already installed, except N editions.

- **`/nodesktopexperience`**

Valid only when installing a multi-session OS VDA. Prevents enabling of the Enhanced Desktop Experience feature. This feature is also controlled with the Enhanced Desktop Experience Citrix policy setting.

- **`/noreboot`**

Prevents a restart after installation. The VDA cannot be used until after a restart.

- **`/noresume`**

By default, when a machine restart is needed during an installation, the installer resumes automatically after the restart completes. To override the default, specify `/noresume`. This can be helpful if you must remount the media or want to capture information during an automated installation.

- **`/portnumber` *port***

Valid only when the `/reconfig` option is specified. Port number to enable for communications between the VDA and the Controller. The previously configured port is disabled, unless it is port 80.

- **`/quiet` or `/passive`**

No user interface appears during the installation. The only evidence of the installation and configuration process is in Windows Task Manager. If this option is omitted, the graphical interface launches.

- **`/reconfigure`**

Customizes previously configured VDA settings when used with the `/portnumber`, `/controllers`, or `/enable_hdx_ports` options. If you specify this option without also specifying the `/quiet` option, the graphical interface for customizing the VDA launches.

- **`/remotepc`**

Valid only for Remote PC Access deployments (single-session OS) or brokered connections (multi-session OS). Excludes installation of the following components:

- Citrix Personalization for App-V
- Citrix Profile Management
- Citrix Profile Management WMI Plug-in
- Machine Identity Service
- Citrix Supportability Tools
- Citrix Files for Windows
- Citrix Files for Outlook
- User personalization layer
- Citrix WEM Agent

This option is not valid when using the `VDAWorkstationCoreSetup.exe` installer. That installer automatically excludes installation of these components.

- **`/remove_appdisk_ack`**

Authorizes the VDA installer to uninstall the AppDisks VDA plug-in if it's installed.

- **`/remove_pvd_ack`**

Authorizes the VDA installer to uninstall Personal vDisk if it's installed.

- **`/remove`**

Removes the components specified with the `/components` option.

- **`/removeall`**

Removes the VDA. It does not remove the Citrix Workspace app (if installed).

- **`/sendexperiencemetrics`**

Automatically sends analytics collected during the installation, upgrade, or removal to Citrix. If this option is omitted (or the `/disableexperiencemetrics` option is specified), analytics are collected locally, but not sent automatically.

- **`/servervdi`**

Installs a single-session OS VDA on a supported Windows server. Omit this option when installing a multi-session VDA on a Windows server. Before using this option, see [Server VDI](#).

- **`/site_guid`** *guid*

Globally Unique Identifier of the site Active Directory Organizational Unit (OU). This associates a virtual desktop with a site when you are using Active Directory for discovery (auto-update is the recommended and default discovery method). The site GUID is a site property displayed in **Manage > Full Configuration**. Do not specify both the `/site_guid` and `/controllers` options.

- **`/tempdir`** *directory*

Directory to hold temporary files during installation. Default = `c:\Windows\Temp`.

This option is not available in the graphical interface.

- **/virtualmachine**

Valid only when installing a VDA on a VM. Overrides detection by the installer of a physical machine, where BIOS information passed to VMs makes them appear as physical machines.

This option is not available in the graphical interface.

- **/wem\_add\_firewall\_rules**

Adds WEM agent firewall rules.

- **/wem\_agent\_cache\_location**

Alternative WEM agent cache location.

- **/wem\_agent\_port**

Port that WEM agent uses to communicate with WEM infrastructure server.

- **/wem\_cached\_data\_sync\_port**

Port that WEM agent uses to synchronize cache with WEM infrastructure server.

- **/wem\_cloud\_connectors**

Comma-separated list of Citrix Cloud Connectors in the resource location containing the WEM deployment.

- **/wem\_server**

FQDN or IP address of the WEM infrastructure server.

## Examples: Install a VDA

The following command installs a VDA on a multi-session OS. The VDA will be used as a master image.

```
VDAServerSetup.exe /quiet /controllers "Contr-East.domain.com"/enable_hdx_ports  
/masterimage
```

The following command installs a Core Services VDA on a single-session OS for use in a Remote PC Access or VDI deployment. Citrix Workspace app and other non-core services are not installed. The address of a Cloud Connector is specified, and ports in the Windows Firewall Service will be opened automatically. The administrator will handle restarts.

```
VDAServerSetup.exe /quiet /controllers "Contr-East.domain.com"/  
enable_hdx_ports /noreboot
```

## Customize a VDA using the command line

After you install a VDA, you can customize several settings. Run `XenDesktopVDASetup.exe`, using one or more of the following options.

- `/reconfigure` (required when customizing a VDA)
- `/h` or `/help`
- `/quiet`
- `/noreboot`
- `/controllers`
- `/portnumber port`
- `/enable_hdx_ports`

## Where to go next

- [Create machine catalogs](#)
- To review the entire configuration process, see [Plan and build a deployment](#).

## Create machine catalogs

September 9, 2021

### Note:

This article describes how to create catalogs using the Full Configuration interface. If you're using Quick Deploy to create Azure resources, follow the guidance in [Create catalogs using Quick Deploy](#).

Collections of physical or virtual machines are managed as a single entity called a machine catalog. All the machines in a catalog have the same type of operating system: multi-session OS or single-session OS. A catalog containing multi-session OS machines can contain either Windows or Linux machines, not both.

The **Manage > Full Configuration** interface guides you to create the first machine catalog. After you create the first catalog, you create the first delivery group. Later, you can change the catalog you created, and create more catalogs.

## Overview

When you create a catalog of VMs, you specify how to provision those VMs. You can use Machine Creation Services (MCS). Or, you can use your own tools to provide machines.

- If you use MCS to provision VMs, you provide an image (or snapshot) to create identical VMs in the catalog. Before you create the catalog, you first use hypervisor or cloud service tools to create and configure the image. This process includes installing a Virtual Delivery Agent (VDA) on the image. Then you create the machine catalog in the **Manage > Full Configuration** interface. You select that image (or a snapshot of an image), specify the number of VMs to create in the catalog, and configure additional information.
- If your machines are already available (so you do not need images), you must still create one or more machine catalogs for those machines.

When using MCS to create the first catalog, you specify a host connection that you created previously. Later (after you create your first catalog and delivery group), you can change information about that connection or create more connections.

If a Cloud Connector is not operating properly, MCS provisioning operations (such as catalog updates) take much longer than usual, and the management interface's performance degrades significantly.

### **Access images from Azure Shared Image Gallery**

When selecting an image to use for creating a machine catalog, you can select images you created in the Azure Shared Image Gallery. These images appear in the list of images in the Master Image screen of the Machine Catalog Setup wizard.

For these images to appear, you must:

1. Configure a Citrix Virtual Apps and Desktops site.
2. Connect to [the Azure Resource Manager](#).
3. In the Azure portal, create a resource group. For details, see [Create an Azure Shared Image Gallery using the portal](#).
4. In the resource group, create a Shared Image Gallery.
5. In the Shared Image Gallery, create an image definition.
6. In the image definition, create an image version.

### **RDS license check**

Creation of a machine catalog containing Windows multi-session OS machines includes an automatic check for valid Microsoft RDS licenses. The catalog is searched for a powered-on and registered machine to perform the check on.

- If a powered-on and registered machine cannot be found, a warning is displayed, explaining that the RDS licensing check cannot be performed.
- If a machine is found and an error is detected, **Manage > Full Configuration** displays a warning message for the catalog containing the detected issue. To remove an RDS license warning from

a catalog (so that it no longer appears in the display), select the catalog. Select **Remove RDS license warning**. When prompted, confirm the action.

### VDA registration

A VDA must be registered with a Cloud Connector to be considered when launching brokered sessions. Unregistered VDAs can result in underutilization of otherwise available resources. There are various reasons a VDA might not be registered, many of which you can troubleshoot. Troubleshooting information is provided in the catalog creation wizard, and after you add a catalog to a delivery group.

In the catalog creation wizard, after you add existing machines, the list of computer account names indicates whether each machine is suitable for adding to the catalog. Hover over the icon next to each machine to display an informative message about that machine.

If the message identifies a problematic machine, you can either remove that machine (using the **Remove** button), or add the machine. For example, if a message indicates that information cannot be obtained about a machine (perhaps because it had never registered), you might choose to add the machine anyway.

For more information about VDA registration troubleshooting, see [CTX136668](#).

### MCS catalog creation summary

Here's a brief overview of default MCS actions after you provide information in the catalog creation wizard.

- If you selected an image (rather than a snapshot), MCS creates a snapshot.
- MCS creates a full copy of the snapshot and places the copy on each storage location defined in the host connection.
- MCS adds the machines to Active Directory, which creates unique identities.
- MCS creates the number of VMs specified in the wizard, with two disks defined for each VM. In addition to the two disks per VM, a master is also stored in the same storage location. If you have multiple storage locations defined, each gets the following disk types:
  - The full copy of the snapshot (noted above), which is read-only and shared across the just-created VMs.
  - A unique 16 MB identity disk that gives each VM a unique identity. Each VM gets an identity disk.
  - A unique difference disk to store writes made to the VM. This disk is thin provisioned (if supported by the host storage) and increases to the maximum size of the master image, if necessary. Each VM gets a difference disk. The difference disk holds changes made during sessions. It is permanent for dedicated desktops. For pooled desktops, it is deleted and a new one created after each restart.

Alternatively, when creating VMs to deliver static desktops, you can specify (on the **Machines** page of the catalog creation wizard) thick (full copy) VM clones. Full clones do not require retention of the master image on every data store. Each VM has its own file.

### **MCS storage considerations**

There are many factors when deciding on storage solutions, configurations, and capacities for MCS. The following information provides proper considerations for storage capacity:

*Capacity considerations:*

- Disks

The Delta or Differencing (Diff) Disks consume the largest amount of space in most MCS deployments for each VM. Each VM created by MCS is given at minimum 2 disks upon creation.

- Disk0 = Diff Disk: contains the OS when copied from the Master Base Image.
- Disk1 = Identity Disk: 16 MB - contains Active Directory data for each VM.

As the product evolves, you might have to add more disks to satisfy certain use cases and feature consumption. For example:

- [MCS Storage Optimization](#) creates a write cache style disk for each VM.
- MCS added the ability to use [full clones](#) as opposed to the Delta disk scenario described in the previous section.

Hypervisor features might also enter into the equation. For example:

- [Citrix Hypervisor IntelliCache](#) creates a Read Disk on local storage for each Citrix Hypervisor. This option saves on IOPS against the image which might be held on the shared storage location.

- Hypervisor overhead

Different hypervisors utilize specific files that create overhead for VMs. Hypervisors also use storage for management and general logging operations. Calculate space to include overhead for:

- [Log files](#)
- Hypervisor specific files. For example:
  - \* VMware adds more files to the **VM storage** folder. See [VMware Best Practices](#).
  - \* Calculate your total virtual machine size requirements. Consider a virtual machine containing 20 GB for the virtual disk, 16 GB for the swap file, and 100 MB for log files consuming 36.1 GB total.
- [Snapshots for XenServer](#); [Snapshots for VMware](#).

- Process overhead

Creating a catalog, adding a machine, and updating a catalog have unique storage implications. For example:

- **Initial catalog creation** requires a copy of the base disk to be copied to each storage location.
  - \* It also requires you to create a **Preparation VM** temporarily.
- **Adding a machine** to a catalog does not require copying of the base disk to each storage location. Catalog creation varies based on the features selected.
- **Updating the catalog** to create an extra base disk on each storage location. Catalog updates also experience a temporary storage peak where each VM in the catalog has 2 Diff disks for a certain amount of time.

*More considerations:*

- **RAM sizing:** Affects the size of certain hypervisor files and disks, including I/O optimization disks, write cache, and snapshot files.
- **Thin / Thick provisioning:** NFS storage is preferred due to the thin provisioning capabilities.

### **Machine Creation Services (MCS) storage optimization**

The Machine Creation Services (MCS) storage optimization feature is also known as MCS I/O:

- The write cache container is *file-based*, the same functionality found in Citrix Provisioning. For example, the Citrix Provisioning write cache file name is `D:\vdiskdif.vhdx` and the MCS I/O write cache file name is `D:\mcsdif.vhdx`.
- Achieve diagnostic improvements by including support for a Windows crash dump file written to the write cache disk.
- MCS I/O retains the technology *cache in RAM with overflow to hard disk* to provide the most optimal multi-tier write cache solution. This functionality allows an administrator to balance between the cost in each tier, RAM and disk, and performance to meet the desired workload expectation.

Updating the write cache method from *disk-based* to *file-based* requires the following changes:

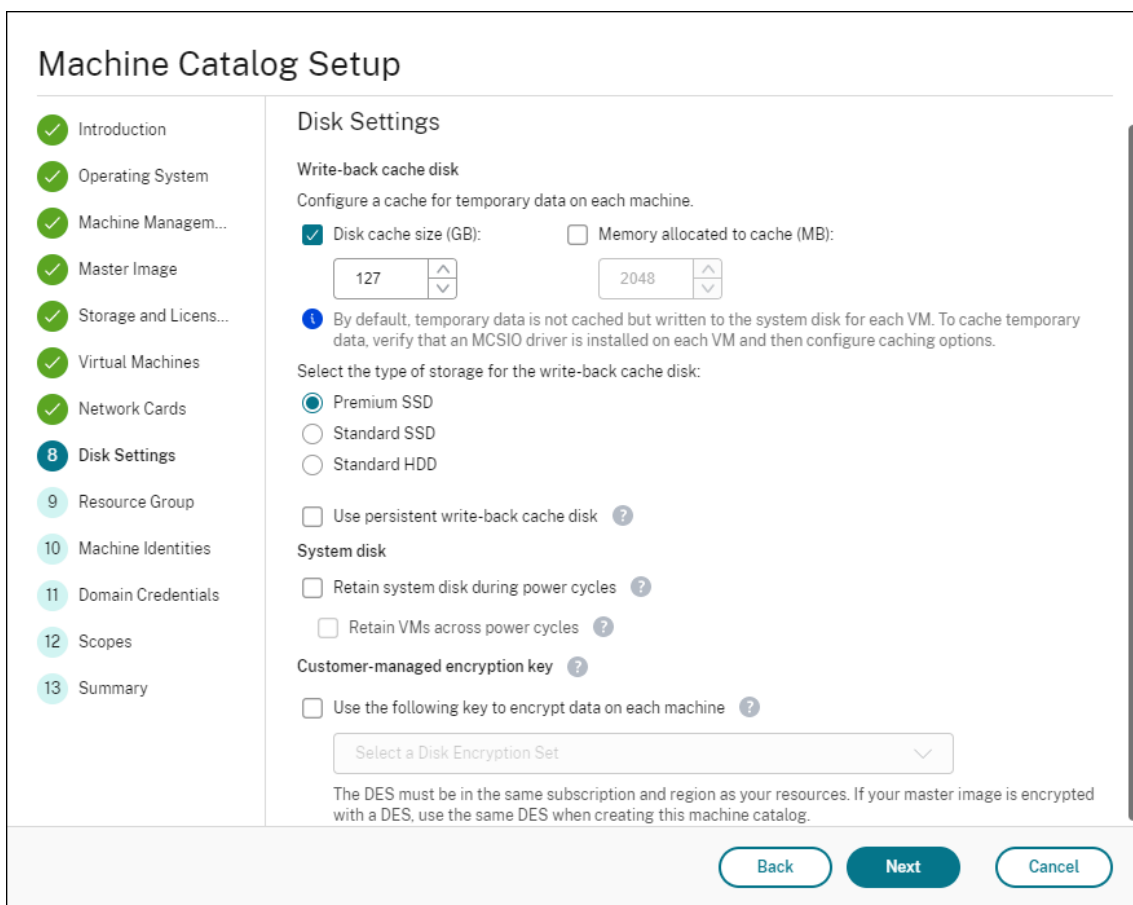
1. MCS I/O no longer supports RAM only cache. Specify a disk size during machine catalog creation.
2. The VM write cache disk is created and formatted automatically when booting a VM for the first time. Once the VM is up, the write cache file `mcsdif.vhdx` is written into the formatted volume `MCSWCDisk`.
3. The pagefile is redirected to this formatted volume, `MCSWCDisk`. As a result, this disk size considers the total amount of disk space. It includes the delta between the disk size and the generated workload plus the pagefile size. This is typically associated with VM RAM size.

### **Enabling MCS storage optimization updates**



When creating a machine catalog, the administrator can configure the RAM and disk size as follows:

- The machine catalog setup user interface of the web-based console:



- The machine catalog setup user interface of the legacy console:

**Machine Catalog Setup**

**Studio**

- ✓ Operating System
- ✓ Machine Management
- ✓ Desktop Experience
- ✓ Master Image
- Virtual Machines**
- Computer Accounts
- Summary

**Virtual Machines**

How many virtual machines do you want to create?

1 - +

Configure your machines.

Total memory (MB) on each machine: 4096 - +

Configure a cache for temporary data on each machine.

Memory allocated to cache (MB): 256 - +

Disk cache size (GB): 10 - +

**i** Caching should not be enabled if you intend to use this catalog to create AppDisks.

If you clear both check boxes, temporary data is not cached; it is written to the OS storage for each VM. (This is the provisioning action in releases earlier than 7.9.)

Back Next Cancel

To enable the MCS I/O storage optimization feature, upgrade the Delivery Controller and the VDA to the latest version of Citrix Virtual Apps and Desktops.

**Note:**

If you upgrade an existing deployment which has MCS I/O enabled, no additional configuration is required. The VDA and the Delivery Controller upgrade handle the MCS I/O upgrade.

With the MCS storage optimization feature enabled, you can do the following when creating a catalog:

- Configure the size of the disk and RAM used for caching temporary data.
- Select the storage type for the write-back cache disk. The following storage types are available to use for the write-back cache disk:
  - **Premium SSD.** Offers a high-performance, low-latency disk storage option suitable for VMs with I/O-intensive workloads.
  - **Standard SSD.** Offers a cost-effective storage option that is suitable for workloads that require consistent performance at lower IOPS levels. An Azure identity disk is always created using Standard SSD.
  - **Standard HDD.** Offers a reliable, low-cost disk storage option suitable for VMs that run latency-insensitive workloads.

MCS configures premium and standard disks to use Locally Redundant Storage (LRS). LRS makes

multiple synchronous copies of your disk data within a single data center. For details about Azure storage types and storage replication, see the following:

- <https://docs.microsoft.com/en-us/azure/storage/common/storage-introduction/>
  - <https://azure.microsoft.com/en-us/documentation/articles/storage-premium-storage/>
  - <https://azure.microsoft.com/en-us/documentation/articles/storage-redundancy/>
- Use **Use persistent write-back cache disk** to control whether the write-back cache disk persists for the provisioned VMs in Azure. Alternatively, you can use PowerShell instead. For details, see [Using PowerShell to create a catalog with persistent write-back cache disk](#).
    - **Use persistent write-back cache disk.** This option lets you control whether the write-back cache disk persists for the provisioned VMs. By default, persistent write-back cache disk is disabled, causing the disk to be deleted during power cycles and any data redirected to the disk to be lost. Enabling this option increases your storage costs. Enable **Disk cache size (GB)** to make this option available.
  - Use **Retain system disk during power cycles** to control whether to retain system disks for VDAs during power cycles.
    - **Retain system disk during power cycles.** By default, the system disk is deleted on shutdown and recreated on startup. This ensures that the disk is always in a clean state but results in longer VM restart times. If system writes are redirected to the RAM cache and overflow to the cache disk, the system disk remains unchanged. Enabling this option increases your storage costs but reduces VM restart times, retains your VM customization, and enables the VMs to be started through the Azure portal. Enable **Disk cache size (GB)** to make this option available.

### Using PowerShell to create a catalog with persistent write-back cache disk

To configure a catalog with persistent write-back cache disk, use the PowerShell parameter `New-ProvScheme CustomProperties`.

#### Tip:

The PowerShell parameter `New-ProvScheme CustomProperties` should only be used for cloud-based hosting connections. If you want to provision machines using a persistent write-back cache disk for an on-premises solution (for example, Citrix Hypervisor) PowerShell is not needed because the disk persists automatically.

This parameter supports an extra property, `PersistWBC`, used to determine how the write-back cache disk persists for MCS provisioned machines. The `PersistWBC` property is only used when the `UseWriteBackCache` parameter is specified, and when the `WriteBackCacheDiskSize` parameter is set to indicate that a disk is created.

**Note:**

This behavior only applies to Azure where the default MCSIO write-back cache disk is deleted and re-created when power cycling. You can choose to persist the disk to avoid the deletion and recreation of MCSIO write-back cache disk.

Examples of properties found in the `CustomProperties` parameter before supporting `PersistWBC` include:

```
1 <CustomProperties xmlns="http://schemas.citrix.com/2014/xd/
   machinecreation" xmlns:xsi="http://www.w3.org/2001/XMLSchema-
   instance">
2 <Property xsi:type="StringProperty" Name="UseManagedDisks" Value="true"
   />
3 <Property xsi:type="StringProperty" Name="StorageAccountType" Value="
   Premium_LRS" />
4 <Property xsi:type="StringProperty" Name="ResourceGroups" Value="
   benva1dev5RG3" />
5 </CustomProperties>
6 <!--NeedCopy-->
```

When using these properties, consider that they contain default values if the properties are omitted from the `CustomProperties` parameter. The `PersistWBC` property has two possible values: **true** or **false**.

Setting the `PersistWBC` property to **true** does not delete the write-back cache disk when the Citrix Virtual Apps and Desktops administrator shuts down the machine from the management interface.

Setting the `PersistWBC` property to **false** deletes the write-back cache disk when the Citrix Virtual Apps and Desktops administrator shuts down the machine from the management interface.

**Note:**

If the `PersistWBC` property is omitted, the property defaults to **false** and the write-back cache is deleted when the machine is shutdown from the management interface.

For example, using the `CustomProperties` parameter to set `PersistWBC` to true:

```
1 <CustomProperties xmlns="http://schemas.citrix.com/2014/xd/
   machinecreation" xmlns:xsi="http://www.w3.org/2001/XMLSchema-
   instance">
2 <Property xsi:type="StringProperty" Name="UseManagedDisks" Value="true"
   />
3 <Property xsi:type="StringProperty" Name="StorageAccountType" Value="
   Premium_LRS" />
4 <Property xsi:type="StringProperty" Name="ResourceGroups" Value="
   benva1dev5RG3" />
```

```
5 <Property xsi:type="StringProperty" Name="PersistWBC" Value="true" />
6 </CustomProperties>
7 <!--NeedCopy-->
```

**Important:**

The `PersistWBC` property can only be set using the `New-ProvScheme` PowerShell cmdlet. Attempting to alter the `CustomProperties` of a provisioning scheme after creation has no impact on the machine catalog and the persistence of the write-back cache disk when a machine is shut down.

For example, set `New-ProvScheme` to use the write-back cache while setting the `PersistWBC` property to `true`:

```
1 New-ProvScheme
2 -CleanOnBoot
3 -CustomProperties "<CustomProperties xmlns='http://schemas.citrix.com
  /2014/xd/machinecreation`" xmlns:xsi='http://www.w3.org/2001/
  XMLSchema-instance`"><Property xsi:type='`StringProperty`' Name='`
  UseManagedDisks`' Value='`true`' /><Property xsi:type='`
  StringProperty`' Name='`StorageAccountType`' Value='`Premium_LRS`'
  /><Property xsi:type='`StringProperty`' Name='`ResourceGroups`'
  Value='`benvaldev5RG3`' /><Property xsi:type='`StringProperty`' Name
  ='`PersistWBC`' Value='`true`' /></CustomProperties>"
4 -HostingUnitName "adSubnetScale1"
5 -IdentityPoolName "BV-WBC1-CAT1"
6 -MasterImageVM "XDHyp:\HostingUnits\adSubnetScale1\image.folder\
  GoldImages.resourcegroup\W10MCSI0-01
  _OsDisk_1_a940e6f5bab349019d57ccef65d2c7e3.manageddisk"
7 -NetworkMapping @{
8   "0"="XDHyp:\HostingUnits\adSubnetScale1\virtualprivatecloud.folder\
  CloudScale02.resourcegroup\adVNET.virtualprivatecloud\
  adSubnetScale1.network" }
9
10 -ProvisioningSchemeName "BV-WBC1-CAT1"
11 -ServiceOffering "XDHyp:\HostingUnits\adSubnetScale1\serviceoffering.
  folder\Standard_D2s_v3.serviceoffering"
12 -UseWriteBackCache
13 -WriteBackCacheDiskSize 127
14 -WriteBackCacheMemorySize 256
15 <!--NeedCopy-->
```

## Improve boot performance with MCSIO

You can improve boot performance for Azure managed disks when MCSIO is enabled. Use the PowerShell `PersistOsDisk` custom property in the `New-ProvScheme` command to configure this feature. Options associated with `New-ProvScheme` include:

```

1 <CustomProperties xmlns="http://schemas.citrix.com/2014/xd/
   machinecreation" xmlns:xsi="http://www.w3.org/2001/XMLSchema-
   instance">
2 <Property xsi:type="StringProperty" Name="UseManagedDisks" Value="true"
   />
3 <Property xsi:type="StringProperty" Name="StorageAccountType" Value="
   Premium_LRS" />
4 <Property xsi:type="StringProperty" Name="Resource<!--NeedCopy-->
5 ` ` ` ` ` ` <!--NeedCopy-->
6 ` ` ` ` ` ` Groups" Value="benvaldev5RG3" />
7 <Property xsi:type="StringProperty" Name="PersistOsDisk" Value="true"
   />
8 </CustomProperties>
9 <!--NeedCopy-->

```

To enable this feature, set the `PersistOsDisk` custom property to **true**. For example:

```

1 New-ProvScheme
2 -CleanOnBoot
3 -CustomProperties "<CustomProperties xmlns=`"http://schemas.citrix.com
   /2014/xd/machinecreation`" xmlns:xsi=`"http://www.w3.org/2001/
   XMLSchema-instance`"><Property xsi:type=`"StringProperty`" Name=`"
   UseManagedDisks`" Value=`"true`" /><Property xsi:type=`"
   StringProperty`" Name=`"StorageAccountType`" Value=`"Premium_LRS`"
   /><Property xsi:type=`"StringProperty`" Name=`"ResourceGroups`"
   Value=`"benvaldev5RG3`" /><Property xsi:type=`"StringProperty`" Name
   =`"PersistOsDisk`" Value=`"true`" /></CustomProperties>"
4 -HostingUnitName "adSubnetScale1"
5 -IdentityPoolName "BV-WBC1-CAT1"
6 -MasterImageVM "XDHyp:\HostingUnits\adSubnetScale1\image.folder\
   GoldImages.resourcegroup\W10MCSIO-01
   _OsDisk_1_a940e6f5bab349019d57ccef65d2c7e3.manageddisk"
7 -NetworkMapping @{
8   "0"="XDHyp:\HostingUnits\adSubnetScale1\virtualprivatecloud.folder\
   CloudScale02.resourcegroup\adVNET.virtualprivatecloud\
   adSubnetScale1.network" }
9
10 -ProvisioningSchemeName "BV-WBC1-CAT1"
11 -ServiceOffering "XDHyp:\HostingUnits\adSubnetScale1\serviceoffering.

```

```
    folder\Standard_D2s_v3.serviceoffering"  
12 -UseWriteBackCache  
13 -WriteBackCacheDiskSize 127  
14 -WriteBackCacheMemorySize 256  
15 <!--NeedCopy-->
```

### Azure Customer-managed encryption key

**Note:**

This feature is available only in the web-based console.

When creating a machine catalog, you can choose whether to encrypt data on the machines provisioned in the catalog. Server-side encryption with a customer-managed encryption key lets you manage encryption at a managed disk level and protect data on the machines in the catalog. A Disk Encryption Set (DES) represents a customer-managed key. To use this feature, you must first create your DES in Azure. A DES is in the following format:

- /subscriptions/12345678-1234-1234-1234-123456789012/resourceGroups/Sample-RG/providers/Microsoft.

Select a DES from the list. The DES you select must be in the same subscription and region as your resources. If your image is encrypted with a DES, use the same DES when creating the machine catalog. You cannot change the DES after you create the catalog.

If you create a catalog with an encryption key and later disable the corresponding DES in Azure, you can no longer power on the machines in the catalog or add machines to it.

### Azure dedicated hosts

You can use MCS to provision virtual machines on Azure dedicated hosts. Before provisioning Virtual machines on Azure dedicated hosts:

- Create a host group.
- Create hosts in that host group.
- Ensure there is sufficient host capacity reserved for creating catalogs and virtual machines.

You can create a catalog of machines with host tenancy defined through the following PowerShell script:

```
1 New-ProvScheme <otherParameters> -CustomProperties '<CustomProperties  
    xmlns="http://schemas.citrix.com/2014/xd/machinecreation" xmlns:xsi  
    ="http://www.w3.org/2001/XMLSchema-instance">  
2 <Property xsi:type="StringProperty" Name="HostGroupId" Value="  
    myResourceGroup/myHostGroup" />  
3 ...other Custom Properties...  
4 </CustomProperties>
```

When using MCS to provision virtual machines on Azure dedicated hosts, consider:

- A *Dedicated host* is a catalog property and cannot be changed once the catalog is created. Dedicated tenancy is currently not supported on Azure.
- A pre-configured Azure host group, in the region of the hosting unit, is required when using the `HostGroupId` parameter.
- Azure auto-placement is required. This functionality makes a request to onboard the subscription associated with the host group. For more information, see [VM Scale Set on Azure Dedicated Hosts - Public Preview](#). If auto-placement is not enabled, MCS will throw an error during catalog creation.

### **AWS dedicated host tenancy support**

You can use MCS to provision AWS dedicated hosts. An administrator can create a catalog of machines with host tenancy defined through PowerShell.

An Amazon [EC2] dedicated host is a physical server with [EC2] instance capacity that is fully dedicated, allowing you to use existing per-socket, or per-VM software licenses.

Dedicated hosts have preset utilization based on instance type. For example, a single allocated dedicated host of C4 Large instance types is limited to running 16 instances. See the [AWS site](#) for more information.

The requirements for provisioning to AWS hosts include:

- An imported BYOL (bring your own license) image (AMI). With dedicated hosts, use and manage your existing licenses.
- An allocation of dedicated hosts with sufficient utilization to satisfy provisioning requests.
- enable **auto-placement**.

To provision to a dedicated host in AWS using PowerShell, use the **New-ProvScheme** cmdlet with the parameter `TenancyType` set to *Host*.

Refer to the [Citrix Developer Documentation](#) for more information.

### **AWS instance property capturing**

When you create a catalog to provision machines using Machine Creation Services (MCS) in AWS, you select an AMI to represent the master/golden image of that catalog. From that AMI, MCS uses a snapshot of the disk. In previous releases, if you wanted roles and/or tags on your machines you would use the AWS console to set them individually. This functionality is enabled by default.



**Tip:**

To use AWS instance property capturing, you must have a VM associated with the AMI.

To improve this process, MCS reads properties from the instance from which the AMI was taken and applies the Identity Access Management (IAM) role and tags of the machine to the machines provisioned for a given catalog. When using this optional feature, the catalog creation process finds the selected AMI source instance, reading a limited set of properties. These properties are then stored in an AWS Launch Template, which is used to provision machines for that catalog. Any machine in the catalog inherits the captured instance properties.

Captured properties include:

- IAM roles – applied to provisioned instances
- Tags - applied to provisioned instances, their disk, and NICs. These tags are applied to transient Citrix resources, including: S3 bucket and objects, volume and worker resources, and AMIs, snapshots, and launch templates.

**Tip:**

The tagging of transient Citrix resources is optional and is configurable using the custom property `AwsOperationalResourcesTagging`.

### Capturing the AWS instance property

You can use this feature by specifying a custom property, `AwsCaptureInstanceProperties`, when creating a provisioning scheme for an AWS hosting connection:

```
New-ProvScheme -CustomProperties "AwsCaptureInstanceProperties, true" ...<
standard provscheme parameters
```

To use this feature, you must specify a broader set of permissions for the AWS service key. These permissions include:

- ec2:AssociateIamInstanceProfile
- ec2:CreateLaunchTemplate
- ec2>DeleteLaunchTemplate
- ec2>DeleteTags
- ec2:DisassociateIamInstanceProfile
- ec2:DescribeIamInstanceProfileAssociations
- ec2:DescribeLaunchTemplates
- ec2:DescribeLaunchTemplateVersions
- ec2:DescribeSnapshots
- ec2:DescribeTags
- iam:PassRole

- s3:PutBucketTagging
- s3:PutObjectTagging

Refer to the [Citrix Developer Documentation](#) for more information.

### **AWS operational resource tagging**

An Amazon Machine Image (AMI) represents a type of virtual appliance used to create a virtual machine within the Amazon Cloud environment, commonly referred to as EC2. You use an AMI to deploy services that use the EC2 environment. When you create a catalog to provision machines using MCS for AWS, you select the AMI to act as the golden image for that catalog.

#### **Important:**

Creating catalogs by capturing an instance property and launch template is required for using operational resource tagging. For details, see the preceding section [AWS instance property capturing](#).

To create an AWS catalog, you must first create an AMI for the instance you want to be the golden image. MCS reads the tags from that instance and incorporates them into the launch template. The launch template tags are then applied to all Citrix resources created in your AWS environment, including:

- Virtual Machines
- VM disks
- VM network interfaces
- S3 buckets
- S3 objects
- Launch templates
- AMIs

### **Tagging an operational resource**

To use PowerShell to tag resources:

1. Open a PowerShell window from the DDC host.
2. Run the command `asnp citrix` to load Citrix-specific PowerShell modules.

To tag a resource for a provisioned VM, use the new custom property `AwsOperationalResourcesTagging`. The syntax for this property is:

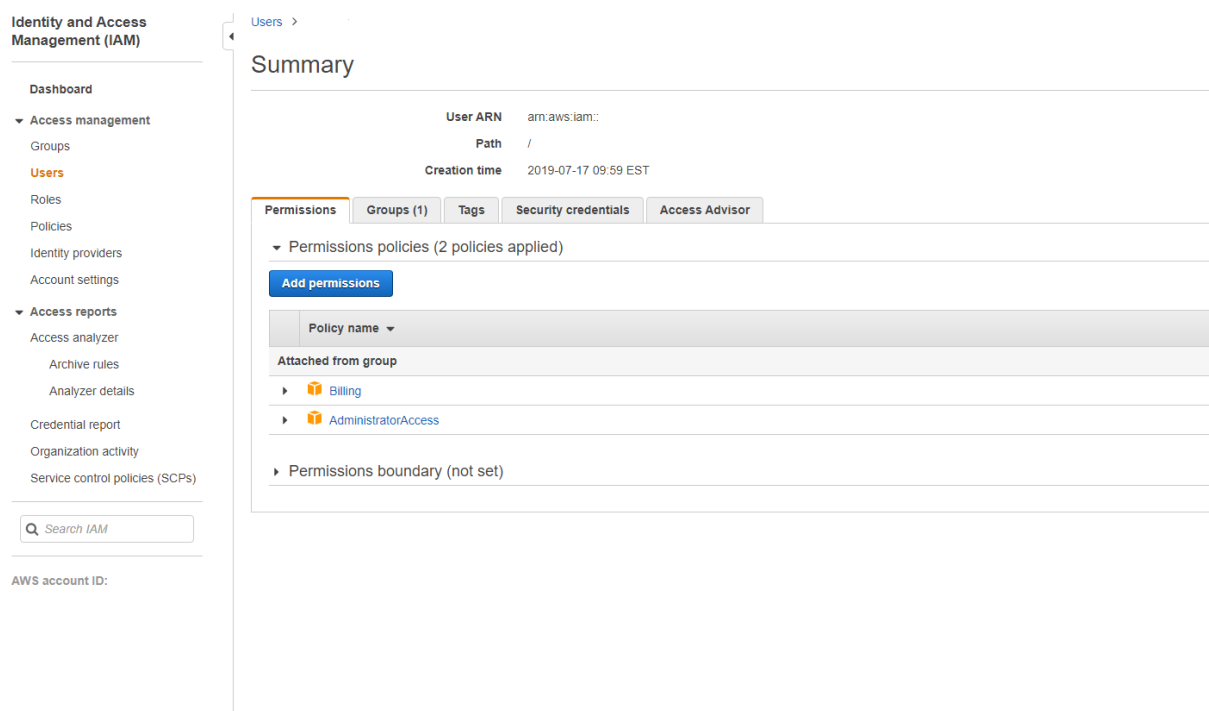
```
New-ProvScheme -CustomProperties "AwsCaptureInstanceProperties,true;  
AwsOperationalResourcesTagging,true" ...<standard provscheme parameters>
```

To use the `AwsOperationalResourcesTagging` custom property, ensure that the following new permissions exist for the AWS service key:

- ec2:CreateTags
- ec2>DeleteTags
- ec2:DescribeTags
- s3:PutBucketTagging
- s3:PutObjectTagging

Set these permissions in the **IAM** section of the AWS Management Console:

1. In the **Summary** panel, select the **Permissions** tab.
2. Select **Add permissions**.



In the **Add Permissions** screen, grant permissions:

Add permissions to

Grant permissions

Use IAM policies to grant permissions. You can assign an existing policy or create a new one.



Create policy

Filter policies	Search	Type	Used as
<input type="checkbox"/>	▶ AdministratorAccess	Job function	Permissions policy (8)
<input type="checkbox"/>	▶ AlexaForBusinessDeviceSetup	AWS managed	None
<input type="checkbox"/>	▶ AlexaForBusinessFullAccess	AWS managed	None
<input type="checkbox"/>	▶ AlexaForBusinessGatewayExecution	AWS managed	None
<input type="checkbox"/>	▶ AlexaForBusinessPolyDelegatedAccessPolicy	AWS managed	None
<input type="checkbox"/>	▶ AlexaForBusinessReadOnlyAccess	AWS managed	None
<input type="checkbox"/>	▶ AmazonAPIGatewayAdministrator	AWS managed	None
<input type="checkbox"/>	▶ AmazonAPIGatewayInvokeFullAccess	AWS managed	None

Use the following as an example in the **JSON** tab:

Create policy 1 2

A policy defines the AWS permissions that you can assign to a user, group, or role. You can create and edit a policy in the visual editor and using JSON. [Learn more](#)

Visual editor **JSON** Import managed policy

```

1 {
2   "Version": "2012-10-17",
3   "Statement": [
4     {
5       "Sid": "VisualEditor0",
6       "Effect": "Allow",
7       "Action": [
8         "ec2:CreateTags",
9         "ec2:DeleteTags",
10        "ec2:DescribeTags",
11        "ec2:PutObjectTagging",
12        "ec2:PutBucketTagging"
13      ],
14      "Resource": "*"
15    },
16    {
17      "Sid": "VisualEditor1",
18      "Effect": "Allow",
19      "Action": "iam:PassRole",
20      "Resource": "arn:aws:iam:*:role/*"
21    }
22  ]
23 }

```

Character count: 304 of 6,144. Cancel **Review policy**

**Tip:**

The noted JSON example may not include all the permissions for your environment. See [How to Define Identity Access Management Permissions Running CVAD on Amazon Web Services](#) for more information.

## Prepare a master image on the hypervisor or cloud service

The master image contains the operating system, non-virtualized applications, VDA, and other software.

Good to know:

- A master image might also be known as a clone image, golden image, base VM, or base image. Host vendors and cloud service providers may use different terms.
- Ensure that the hypervisor or cloud service has enough processors, memory, and storage to accommodate the number of machines created.
- Configure the correct amount of hard disk space needed for desktops and applications. That value cannot be changed later or in the machine catalog.
- Remote PC Access machine catalogs do not use master images.

- Microsoft KMS activation considerations when using MCS: If your deployment includes 7.x VDAs with a XenServer 6.1 or 6.2, vSphere, or Microsoft System Center Virtual Machine Manager host, you do not need to manually rearm Microsoft Windows or Microsoft Office.

Install and configure the following software on the master image:

- Integration tools for your hypervisor (such as Citrix VM Tools, Hyper-V Integration Services, or VMware tools). If you omit this step, applications and desktops might not function correctly.
- A VDA. Citrix recommends installing the latest version to allow access to the newest features. Failure to install a VDA on the master image causes the catalog creation to fail.
- Third-party tools as needed, such as antivirus software or electronic software distribution agents. Configure services with settings that are appropriate for users and the machine type (such as updating features).
- Third-party applications that you are not virtualizing. Citrix recommends virtualizing applications. Virtualizing reduces costs by eliminating having to update the master image after adding or reconfiguring an application. Also, fewer installed applications reduce the size of the master image hard disks, which saves storage costs.
- App-V clients with the recommended settings, if you plan to publish App-V applications. The App-V client is available from Microsoft.
- When using MCS, if you localize Microsoft Windows, install the locales and language packs. During provisioning, when a snapshot is created, the provisioned VMs use the installed locales and language packs.

**Important:**

If you are using MCS, do not run Sysprep on master images.

To prepare a master image:

1. Using your hypervisor's management tool, create a master image and then install the operating system, plus all service packs and updates. Specify the number of vCPUs. You can also specify the vCPU value if you create the machine catalog using PowerShell. You cannot specify the number of vCPUs when creating a catalog from **Manage > Full Configuration**. Configure the amount of hard disk space needed for desktops and applications. That value cannot be changed later or in the catalog.
2. Ensure that the hard disk is attached at device location 0. Most standard master image templates configure this location by default, but some custom templates might not.
3. Install and configure the software listed above on the master image.
4. If you are not using MCS, join the master image to the domain where applications and desktops are members. Ensure that the master image is available on the host where the machines are created. If you are using MCS, joining the master image to a domain is not required. The provisioned machines are joined to the domain specified in the catalog creation wizard.
5. Citrix recommends that you create and name a snapshot of your master image so that it can be

identified later. If you specify a master image rather than a snapshot when creating a catalog, the management interface creates a snapshot, but you cannot name it.

## Start creating the catalog

Before creating a catalog:

- Review this section to learn about the choices you make and information you supply.
- Ensure that you have created a connection to the hypervisor, cloud service, other resource that hosts your machines.
- If you have created a master image to provision machines, ensure that you have installed a VDA on that image.

To start the catalog creation wizard:

1. Sign in to [Citrix Cloud](#). In the upper left menu, select **My Services > Virtual Apps and Desktops**.
2. Select **Manage**.
3. If this is the first catalog being created, you are guided to the correct selection (such as “Set up the machines and create machine catalogs to run apps and desktops.”). The catalog creation wizard opens and walks you through the items described below.

If you already created a catalog and want to create another, from **Manage > Full Configuration**, select **Machine Catalogs** in the left pane. Then select **Create Machine Catalog**.

The wizard walks you through the pages described below. The pages you see may differ, depending on the selections you make, and the connection (to a host) you use. [Hosts / virtualization resources](#) lists information sources for the supported host types.

## Operating system

Each catalog contains machines of only one type:

- **Multi-session OS:** A multi-session OS catalog provides hosted shared desktops. The machines can be running supported versions of the Windows or Linux operating systems, but the catalog cannot contain both.
- **Single-session OS:** A single-session OS catalog provides VDI desktops that you can assign to various different users.
- **Remote PC Access:** A Remote PC Access catalog provides users with remote access to their physical office desktop machines. Remote PC Access does not require a VPN to provide security.

## Machine management

This page does not appear when you are creating a Remote PC Access catalog.

The **Machine Management** page indicates how machines are managed and which tool you use to deploy machines.

Choose if machines in the catalog will be power managed through the Full Configuration interface.

- Machines are power managed through the Full Configuration interface or provisioned through a cloud environment, for example, VMs or blade PCs. This option is available only if you already configured a [connection](#) to a hypervisor or cloud service.
- Machines are not power managed through the Full Configuration interface, for example, physical machines.

If you indicated that machines are power managed through the Full Configuration interface or provisioned through a cloud environment, choose which tool to use to create VMs.

- **Citrix Machine Creation Services (MCS):** Uses a master image to create and manage virtual machines. Machine catalogs in cloud environments use MCS. MCS is not available for physical machines.
- **Other:** A tool that manages machines already in the data center. Citrix recommends that you use Microsoft System Center Configuration Manager or another third-party application to ensure that the machines in the catalog are consistent.

## Desktop types (desktop experience)

This page appears only when you are creating a catalog containing single-session OS machines.

The **Desktop Experience** page determines what occurs each time a user logs on. Select one of:

- Users connect to a new (random) desktop each time they log on.
- Users connect to the same (static) desktop each time they log on.

## Master image

This page appears only when you are using MCS to create VMs.

Select the connection to the host hypervisor or cloud service, and then select the snapshot or VM created earlier.

### Note:

- When you are using MCS, do not run Sysprep on master images.
- If you specify a master image rather than a snapshot, the management interface creates a snapshot, but you cannot name it.

Do not change the default minimum VDA version selection. To enable use of the latest product features, ensure that the master image has the latest VDA version installed.

An error message appears if you select a snapshot or VM that is not compatible with the machine management technology you selected earlier in the wizard.

## Cloud platform and service environments

When you are using a cloud service or platform to host VMs, the catalog creation wizard may contain extra pages specific to that host. For example, when using an Azure Resource Manager master image, the catalog creation wizard contains a **Storage and License Types** page.

For host-specific information, follow the appropriate link listed in Start creating the catalog.

## Machines

This page does not appear when you are creating Remote PC Access catalogs.

The title of this page depends on what you selected on the **Machine Management** page: **Machines**, **Virtual Machines**, or **VMs and users**.

- **When using MCS to create machines:**

- Specify how many virtual machines to create.
- Choose the amount of memory (in MB) each VM has.
- **Important:** Each created VM has a hard disk. Its size is set in the master image; you cannot change the hard disk size in the catalog.
- If you indicated on the **Desktop Experience** page that user changes to static desktops should be saved on a separate Personal vDisk, specify the virtual disk size in GB and the drive letter.
- If your deployment uses more than one zone (resource location), you can select a zone for the catalog.
- If you are creating static desktop VMs, select a virtual machine copy mode. See Virtual machine copy mode.
- If you are creating random desktop VMs that do not use personal vDisks, you can configure a cache to be used for temporary data on each machine. See Configure cache for temporary data.

- **When using other tools to provide machines:**

Add (or import a list of) Active Directory machine account names. You can change the Active Directory account name for a VM after you add/import it. If you specified static machines on the **Desktop Experience** wizard page, you can optionally specify the Active Directory user name for each VM you add.

After you add or import names, you can use the **Remove** button to delete names from the list, while you are still on this wizard page.



- **When using other tools (but not MCS):**

An icon and tooltip for each machine added (or imported) help identify machines that might not be eligible to add to the catalog, or be unable to register with a Cloud Connector.

### Virtual machine copy mode

The copy mode you specify on the **Machines** page determines whether MCS creates thin (fast copy) or thick (full copy) clones from the master image. (Default = thin clones)

- Use fast copy clones for more efficient storage use and faster machine creation.
- Use full copy clones for better data recovery and migration support, with potentially reduced IOPS after the machines are created.

### Configure cache for temporary data

Caching temporary data locally on the VM is optional. You can enable use of the temporary data cache on the machine when you use MCS to manage pooled (not dedicated) machines in a catalog. If the catalog uses a connection that specifies storage for temporary data, you can enable and configure the temporary data cache information when you create the catalog.

To enable the caching of temporary data, the VDA on each machine in the catalog must be minimum version 7.9. This feature is referred to as **MCSIO**.


#### **Important:**

This feature requires a current MCSIO driver. Installing this driver is an option when you install or upgrade a VDA. By default, that driver is not installed.

You specify whether temporary data uses shared or local storage when you create the connection that the catalog uses. For details, see [Connections and resources](#). Enabling and configuring the temporary cache in the catalog includes two check boxes and values: **Memory allocated to cache (MB)** and **Disk cache size (GB)**. By default, these check boxes are cleared. When you enable one or both check boxes, the default values differ according to the connection type. Generally, the default values are sufficient for most cases; however, take into account the space needed for:

- Temporary data files created by Windows itself, including the Windows page file.
- User profile data.
- ShareFile data that is synced to users' sessions.
- Data that may be created or copied by a session user or any applications users may install inside the session.

If you enable the **Disk cache size** check box, temporary data is initially written to the memory cache. When the memory cache reaches its configured limit (the **Memory allocated to cache** value), the oldest data is moved to the temporary data cache disk.



The memory cache is part of the total amount of memory on each machine; therefore, if you enable the **Memory allocated to cache** check box, consider increasing the total amount of memory on each machine.

If you clear the **Memory allocated to cache** check box and leave the **Disk cache size** check box enabled, temporary data is written directly to the cache disk, using a minimal amount of memory cache.

Changing the **Disk cache size** from its default value can affect performance. The size must match user requirements and the load placed on the machine.

**Important:**

If the disk cache runs out of space, the user's session becomes unusable.

If you clear the **Disk cache size** check box, no cache disk is created. In this case, specify a **Memory allocated to cache** value that is large enough to hold all of the temporary data. This is feasible only if large amounts of RAM are available for allocation to each VM.

If you clear both check boxes, temporary data is not cached. It is written to the difference disk (located in the OS storage) for each VM. (This is the provisioning action in releases earlier than 7.9.)

Do not enable caching if you intend to use this catalog to create AppDisks.

You cannot change the cache values in a machine catalog after it is created.

### Using CSV files to bulk add machines

If you use the **Full Configuration** management interface, you can bulk add machines by using CSV files. The feature is available to all catalogs except catalogs created through MCS.

A general workflow to use CSV files to bulk add machines is as follows:

1. On the **Machines** page, select **Add CSV File**. The **Add Machines in Bulk** window appears.
2. Select **Download CSV Template**.
3. Fill out the template file.
4. Drag or browse to the file to upload it.
5. Select **Validate** to perform validation checks on your import.
6. Select **Import** to complete.

For information about CSV file considerations, see [Considerations when using CSV files to add machines](#).

## NIC (NICs)

This page does not appear when you are creating Remote PC Access catalogs.

If you plan to use multiple NICs, associate a virtual network with each card. For example, you can assign one card to access a specific secure network, and another card to access a more commonly used network. You can also add or remove NICs from this page.

## Machine accounts

This page appears only when creating Remote PC Access catalogs.

Specify the Active Directory machine accounts or Organizational Units (OUs) to add that correspond to users or user groups. Do not use a forward slash (/) in an OU name.

You can choose a previously configured power management connection or elect not to use power management. If you want to use power management but a suitable connection hasn't been configured yet, you can create that connection later and then edit the machine catalog to update the power management settings.

If you use the *web-based console*, you can also bulk add machines by using CSV files. A general workflow to do that is as follows:

1. On the **Machine Accounts** page, select **Add CSV File**. The **Add Machines in Bulk** window appears.
2. Select **Download CSV Template**.
3. Fill out the template file.
4. Drag or browse to the file to upload it.
5. Select **Validate** to perform validation checks on your import.

6. Select **Import** to complete.

For information about CSV file considerations, see [Considerations when using CSV files to add machines](#).

## Machine identities

This page appears only when using MCS to create VMs.

Each machine in the catalog must have a unique identity. This page lets you configure identities for machines in the catalog. The machines are joined to the identity after they are provisioned. You cannot change the identity type after you create the catalog.

A general workflow to configure settings on this page is as follows:

1. Select an identity from the list.
2. Indicate whether to create accounts or use existing ones, and the location (domain) for those accounts.

You can select one of the following options:

- **On-premises Active Directory.** Machines owned by an organization and signed into with an Active Directory account that belongs to that organization. They exist on-premises.
- **Non-domain-joined.** Machines not joined to any domain.

### Important:

- If you select **On-premises Active Directory** as the identity type, each machine in the catalog must have a corresponding Active Directory computer account.
- The **Non-domain-joined** identity type requires version 1811 or later of the VDA as the minimum functional level for the catalog. To make it available, update the minimum functional level.

If you create accounts, you must have permission to create computer accounts in the OU where the machines reside.

- Each machine in the catalog must have a unique computer name. Specify the account naming scheme for the machines you want to create. Use hash marks to indicate where sequential numbers or letters appear. Do not use a forward slash (/) in an OU name. A name cannot begin with numbers or blank spaces. For example, a naming scheme of PC-Sales-## (with 0-9 selected) results in computer accounts named PC-Sales-01, PC-Sales-02, PC-Sales-03, and so on.

If you use existing accounts, browse to the accounts or click **Import** and specify a .csv file containing account names. The imported file content must use the format:

- [ADComputerAccount] ADcomputeraccountname.domain

Ensure that there are enough accounts for all the machines you are adding. The Full Configuration interface manages those accounts, so either allow that interface to reset the passwords for all the accounts or specify the account password, which must be the same for all accounts.

For catalogs containing physical or existing machines, select or import existing accounts and assign each machine to both an Active Directory computer account and to a user account.

## Domain credentials

Select **Enter credentials** and enter user credentials with sufficient permissions to create machine accounts in Active Directory.

### Tip:

The account you used to log into the Full Configuration interface is the same account you use to interact with Active Directory. When provisioning a new catalog, adding machines to a catalog, or removing a catalog and the machine names from Active Directory, a dialog appears requesting your Active Directory domain administrator credentials.

## Summary, name, and description

On the **Summary** page, review the settings you specified. Enter a name and description for the catalog. This information appears in the Full Configuration management interface.

When you're done, select **Finish** to start the catalog creation.

## More information

- [Citrix Virtual Apps and Desktops Image Management](#)
- [Connections and resources](#)
- [Manage machine catalogs](#)

## Where to go next

If this is the first catalog created, you are guided to [create a delivery group](#).

To review the entire configuration process, see [Plan and build a deployment](#).

## Manage machine catalogs

August 11, 2021

**Note:**

This article describes how to manage catalogs using the Full Configuration interface. If you created the catalog using the Quick Deploy interface, and want to continue using that interface to manage the catalog follow the guidance in [Manage catalogs in Quick Deploy](#).

## Introduction

You can add or remove machines from a machine catalog, and rename, change the description, or manage a catalog's Active Directory computer accounts.

Maintaining catalogs can also include making sure that each machine has the latest OS updates, antivirus software updates, operating system upgrades, or configuration changes.

- Catalogs containing pooled random machines created using Machine Creation Services (MCS) maintain machines by updating the image used in the catalog and then updating the machines. This method enables you to efficiently update large numbers of user machines.
- For catalogs containing static, permanently assigned machines, and for Remote PC Access Machine catalogs, you manage updates to users' machines outside of the Full Configuration management interface. Perform this task either individually or collectively using third-party software distribution tools.

For information about creating and managing connections to host hypervisors and cloud services, see [Connections and resources](#).

**Note:**

MCS does not support Windows 10 IoT Core and Windows 10 IoT Enterprise. Refer to the [Microsoft site](#) for more information.

## About persistent instances

When updating an MCS catalog created using persistent, or dedicated instances, any new machines created for the catalog use the updated image. Pre-existing instances continue to use the original instance. The process of updating an image is done the same way for any other type of catalog. Consider the following:

- With persistent disk catalogs, the pre-existing machines are not updated to the new image, but any new machines added to the catalog use the new image.
- For non-persistent disk catalogs, the machine image is updated the next time the machine is reset.
- With persistent machine catalogs, updating the image also updates the catalog instances that use it.

- For catalogs that do not persist, if you want different images for different machines, the images must reside in separate catalogs.

## Add machines to a catalog

Before you start:

- Make sure the virtualization host (hypervisor or cloud service provider) has sufficient processors, memory, and storage to accommodate the additional machines.
- Make sure that you have enough unused Active Directory computer accounts. If you are using existing accounts, the number of machines you can add is limited by the number of accounts available.
- If you use the Full Configuration management interface to create Active Directory computer accounts for the additional machines, you must have appropriate domain administrator permission.

To add machines to a catalog:

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select a machine catalog and then select **Add machines** in the action bar.
3. Select the number of virtual machines to add.
4. If there are insufficient existing Active Directory accounts for the number of VMs you are adding, select the domain and location where the accounts are created. Specify an account naming scheme, using hash marks to indicate where sequential numbers or letters appear. Do not use a forward slash (/) in an OU name. A name cannot begin with a number. For example, a naming scheme of PC-Sales-## (with 0-9 selected) results in computer accounts named PC-Sales-01, PC-Sales-02, PC-Sales-03, and so on.
5. If you use existing Active Directory accounts, either browse to the accounts or select **Import** and specify a .csv file containing account names. Make sure that there are enough accounts for all the machines you are adding. The Full Configuration interface manages these accounts. Either allow that interface to reset the passwords for all the accounts, or specify the account password, which must be the same for all accounts.

The machines are created as a background process, and can take much time when creating many machines. Machine creation continues even if you close the Full Configuration management interface.

## Use CSV files to bulk add machines to a catalog

### Important:

This feature is available only in the **Full Configuration** management interface.

You can bulk add machines by using CSV files. The feature is available to all catalogs except catalogs created through MCS.

To bulk add machines to a catalog, complete the following steps:

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select a machine catalog and then select **Add Machines** in the action bar. The **Add Machines** window appears.
3. Select **Add CSV File**. The **Add Machines in Bulk** window appears.
4. Select **Download CSV Template**.
5. Fill out the template file.
6. Drag or browse to the file to upload it.
7. Select **Validate** to perform validation checks on your import.
8. Select **Import** to complete the process.

### Considerations when using CSV files to add machines

When editing the CSV template file, keep the following in mind:

- The feature gives you more flexibility to bulk add machines through a CSV file. In the file, you can add only machines (for use with user auto-assignments) or add machines along with user assignments. Type your data in the following format:
  - For machine account and user name (samName) pairs:
    - \* Domain\ComputerName1,Domain\Username1
    - \* Domain\ComputerName2,Domain\Username1;Domain\Username2
  - For machine accounts only:
    - \* Domain\ComputerName1
    - \* Domain\ComputerName2
  - For VM and user name pairs:
    - \* XDHyp:\Connections\ConnectioName\Region\vm.folder\VMName1.vm,Domain\ComputerName1
    - \* XDHyp:\Connections\ConnectioName\Region\vm.folder\VMName2.vm,Domain\ComputerName2
  - For VMs only:
    - \* XDHyp:\Connections\ConnectioName\Region\vm.folder\VMName1.vm,Domain\ComputerName1
    - \* XDHyp:\Connections\ConnectioName\Region\vm.folder\VMName2.vm,Domain\ComputerName2
- The maximum number of machines that a file can contain is 1,000. To import more than 1,000 machines, spread them across different files and then import those files one by one. We recommend that you import no more than 1,000 machines. Otherwise, catalog creation can take a long time to complete.



## Delete machines from a catalog

After you delete a machine from a machine catalog, users can no longer access it, so before deleting a machine, ensure that:

- User data is backed up or no longer required.
- All users are logged off. Turning on maintenance mode stops new connections from being made to a machine.
- Machines are powered off.

To delete machines from a catalog:

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select a catalog and then select **View Machines** in the action bar.
3. Select one or more machines and then select **Delete** in the action bar.

Choose whether to delete the machines being removed. If you choose to delete the machines, indicate whether the Active Directory accounts for those machines are to be retained, disabled, or deleted.

When you delete an Azure Resource Manager machine catalog, the associated machines and resource groups are deleted from Azure, even if you indicate that they are to be retained.

## Change a catalog description or change Remote PC Access settings

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select a catalog and then select **Edit Machine Catalog** in the action bar.
3. (Remote PC Access catalogs only) On the **Power Management** page, you can change the power management settings and select a power management connection. On the **Organizational Units** page, add or remove Active Directory OUs.
4. On the **Description** page, change the catalog description.

## Rename a catalog

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select a catalog and then select **Rename Machine Catalog** in the action bar.
3. Enter the new name.

## Move a catalog to a different zone

If your deployment has more than one zone, you can move a catalog from one zone to another.

Keep in mind that moving a catalog to a different zone than the hypervisor or cloud service containing the VMs in that catalog can affect performance.

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.

2. Select a catalog and then select **Move** in the action bar.
3. Select the zone where you want to move the catalog.

## Delete a catalog

Before deleting a catalog, ensure that:

- All users are logged off and that no disconnected sessions are running.
- Maintenance mode is turned on for all machines in the catalog so that new connections cannot be made.
- All machines in the catalog are powered off.
- The catalog is not associated a delivery group. In other words, the delivery group does not contain machines from the catalog.

To delete a catalog:

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select a catalog and then select **Delete Machine Catalog** in the action bar.
3. Indicate whether the machines in the catalog are to be deleted. If you choose to delete the machines, indicate whether the Active Directory computer accounts for those machines are to be retained, disabled, or deleted.

## Manage Active Directory computer accounts in a catalog

To manage Active Directory accounts in a machine catalog, you can:

- Free unused machine accounts by removing Active Directory computer accounts from single-session and multi-session catalogs. Those accounts can then be used for other machines.
- Add accounts so that when more machines are added to the catalog, the computer accounts are already in place. Do not use a forward slash (/) in an OU name.

To manage Active Directory accounts:

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select a catalog and then select **Manage AD accounts** in the action bar.
3. Choose whether to add or delete computer accounts. If you add accounts, specify what to do with the account passwords: either reset them all or enter a password that applies to all accounts.

You might reset passwords if you do not know the current account passwords; you must have permission to perform a password reset. If you enter a password, the password is changed on the accounts as they are imported. If you delete an account, choose whether the account in Active Directory are to be kept, disabled, or deleted.

You can also indicate whether Active Directory accounts are to be retained, disabled, or deleted when you remove machines from a catalog or delete a catalog.

## Update a catalog

Citrix recommends that you save copies or snapshots of images before you update the machines in the catalog. The database keeps a historical record of the images used with each machine catalog. Roll back, or revert, machines in a catalog to use the previous version of the image. Perform this task if users encounter problems with updates you deployed to their desktops, minimizing user downtime. Do not delete, move, or rename images. Otherwise, you cannot revert a catalog to use them.

After a machine is updated, it restarts automatically.

## Update or create an image

Before you update the Machine Catalog, either update an existing image or create a one on your host hypervisor.

1. On your hypervisor or cloud service provider, take a snapshot of the current VM and give the snapshot a meaningful name. This snapshot can be used to revert (roll back) machines in the catalog, if needed.
2. If necessary, power on the image, and log on.
3. Install updates or make any required changes to the image.
4. If the image uses a Personal vDisk, update the inventory.
5. Power off the VM.
6. Take a snapshot of the VM, and give the snapshot a meaningful name that is recognized when the catalog is updated. Although the management interface can create a snapshot, Citrix recommends that you create a snapshot using the hypervisor management console, and then select that snapshot in the Full Configuration management interface. This enables you to provide a meaningful name and description rather than an automatically generated name. For GPU images, you can change the image only through the XenServer XenCenter console.

## Update the catalog

To prepare and roll out the update to all machines in a catalog:

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select a catalog and then select **Update Machines** in the action bar.
3. On the **Master Image** page, select the host and the image you want to roll out.
4. On the **Rollout Strategy** page, choose when the machines in the Machine Catalog are updated with the new image: on the next shutdown or immediately.

5. Verify the information on the **Summary** page and then select **Finish**. Each machine restarts automatically after it is updated.

If you are updating a catalog using the PowerShell SDK, you can specify a hypervisor template (`VMTemplates`), as an alternative to an image or a snapshot of an image.

#### **Rollout strategy:**

Updating the image on next shutdown will immediately affect any machines not currently in use, that is, machines that do not have an active user session. A system that is in use receives the update when the current active session ends. Consider the following:

- New sessions cannot be launched until the update has completed on applicable machines.
- For single-session machines, machines are immediately updated when the machine is not in use, or when users are not logged in.
- For a multi-session OS with child machines, reboots do not occur automatically. They must be manually shut down and restarted.

#### **Tip:**

Limit the number of machines being rebooted by using the advanced settings for a host connection. Use these settings to modify the actions taken for a given catalog; advanced settings vary depending on the hypervisor.

If you choose to update the image immediately, configure a distribution time and notifications.

- **Distribution time:** You can choose to update all machines at the same time, or specify the total length of time it is expected to take to begin updating all machines in the catalog. An internal algorithm determines when each machine is updated and restarted during that interval.
- **Notification:** In the left notification list, choose whether to display a notification message on the machines before an update begins. By default, no message is displayed.

If you choose to display a message 15 minutes before the update begins, you can choose (in the right list) to repeat the message every five minutes after the initial message. By default, the message is not repeated.

Unless you choose to update all machines at the same time, the notification message displays on each machine at the appropriate time before the update begins, calculated by an internal algorithm.

#### **Roll back an update**

After you roll out an updated/new image, you can roll it back. This might be necessary if issues occur with the newly updated machines. When you roll back, machines in the catalog are rolled back to the last working image. Any new features that require the newer image will no longer be available. As with the rollout, rolling back a machine includes a restart.

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select the catalog and then select **Rollback machine update** in the action bar.
3. Specify when to apply the earlier image to machines, as described for the rollout operation.

The rollback is applied only to machines that need to be reverted. For machines that have not been updated with the new/updated image (for example, machines with users who have not logged off), users do not receive notification messages and are not forced to log off.

## Upgrade a catalog or revert an upgrade

Upgrade the machine catalog after you upgrade the VDAs on the machines to a newer version. Citrix recommends upgrading all VDAs to the latest version to enable access to all the newest features.

Before upgrading a catalog:

- Start the upgraded machines so that they register with the Controller. This lets the management interface determine that the machines in the catalog need upgrading.

To upgrade a catalog:

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select the catalog. The **Details** tab in the lower pane displays version information.
3. Select **Upgrade Catalog**. If the management interface detects that the catalog needs upgrading, it displays a message. Follow the prompts. If one or more machines cannot be upgraded, a message explains why. Citrix recommends you resolve machine issues before upgrading the catalog to ensure that all machines function properly.

After the catalog upgrade completes, you can revert the machines to their previous VDA versions by selecting the catalog and then selecting **Undo** in the action bar.

## Clone a catalog

Before cloning a catalog, be aware of the following considerations:

- You cannot change settings associated with [operating system](#) and [machine management](#). The cloned catalog inherits those settings from the original.
- Cloning a catalog can take some time to complete. Select **Hide progress** to run the cloning in the background if needed.
- The cloned catalog inherits the name of the original and has a suffix [Copy](#). You can change the name. See [Rename a catalog](#).
- After cloning completes, be sure to assign the cloned catalog to a delivery group.

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select a catalog and then select **Clone** in the action bar.

3. In the **Clone Selected Machine Catalog** window, view the settings for the cloned catalog and configure settings as applicable. Select **Next** to proceed to the next page.
4. On the **Summary** page, view a summary of the settings and select **Finish** to start cloning.
5. Select **Hide progress** to run the cloning in the background if needed.

## Manage tags

You can use the Full Configuration management interface to apply or remove a tag to or from a catalog. This feature is available only in the web-based console.

1. From **Manage > Full Configuration**, select **Machine Catalogs** in the left pane.
2. Select a machine catalog and then select **More**.
3. Select **Manage Tags** from the menu. The **Manage Tags** window appears.
4. If a tag exists, select the check box next to its name. Otherwise, select **Create** and then specify a name for the tag. After the tag is created, you automatically return to the **Manage Tags** window, with the check box next to its name selected.
5. In the **Manage Tags** window, select **Save** to apply your changes and to exit the window.

## Configure support for non-domain joined catalogs

Using the Citrix Virtual Apps and Desktops service, you can create catalogs based on workgroup, or, non-domain joined machines. Creating non-domain joined machines depends on how the account identity pool is created. The account identity pool is the mechanism used by MCS to create and track machine names during catalog provisioning.

For example, in past releases all Active Directory fields were supplied in a single instance:

```
1 New-AcctIdentityPool AllowUnicode -Domain "awsdevexample.local" -
  IdentityPoolName "DedicatedHostCatalog" -NamingScheme "MH-DHost##" -
  NamingSchemeType "Numeric" *-OU "CN=Computers,DC= awsdevexample,DC=
  local"* -Scope @() -ZoneUid "81291221-d2f2-49d2-ab12-bae5bbd0df05"
2 <!--NeedCopy-->
```

MCS uses a new PoSH parameter, `WorkgroupMachine`, to create a workgroup catalog. Using the same example, noted above, this parameter removes the requirement to specify all the AD-specific parameters, including domain administrator credentials:

```
1 New-AcctIdentityPool AllowUnicode -WorkgroupMachine -IdentityPoolName "
  DedicatedHostCatalog" -NamingScheme "MH-DHost##" -NamingSchemeType "
  Numeric" -Scope @() -ZoneUid "81291221-d2f2-49d2-ab12-bae5bbd0df05"
2 <!--NeedCopy-->
```

**Note:**

When using the `WorkgroupMachine` parameter, consider that non-domain joined machine catalogs are only supported through PowerShell for all catalog lifecycle events including provisioning, adding/removing machines from the catalog, updating, and power management.

## Troubleshoot

- For machines with `Power State Unknown` status, see [CTX131267](#) for guidance.
- If a Cloud Connector is not operating properly, MCS provisioning operations (such as catalog updates) take much longer than usual and the management console performance degrades significantly.

## Quick Deploy

August 24, 2021

### Introduction

In the Citrix Virtual Apps and Desktops service, the **Manage > Quick Deploy** interface offers fast deployment of apps and desktops when you're using Microsoft Azure to host your desktops and apps. This interface offers basic configuration, without advanced features.

Use Quick Deploy to:

- Provision virtual machines and catalogs that deliver desktops and apps hosted in Microsoft Azure.
- Create Remote PC Access catalogs for existing machines.

With Quick Deploy, you can use a [Citrix Managed Azure](#) subscription, or your own Azure subscription.

As an alternative to Quick Deploy, the **Full Configuration** interface offers advanced configuration features. For information about the service's **Manage** tab options, see [Management interfaces](#).

### Differences between management interfaces

The following table compares the Full Configuration and Quick Deploy interfaces.

Feature	Quick Deploy	Full Configuration
Deploy using Azure	Yes	Yes *
Deploy using other cloud services	No	Yes
Deploy using on-premises hypervisors	No	Yes
Citrix prepared images available	Yes	No
Simplified user experience	Yes	No

\* When using a Citrix Managed Azure subscription, you must use Quick Deploy when creating an image or catalog.

If you are familiar with using Full Configuration to create and manage catalogs, Quick Deploy has the following differences.

- Different terminology.
  - In Quick Deploy, you create a catalog.
  - In Full Configuration, you create a machine catalog. In practice, it is often referred to as simply catalog.
- Resource location and Cloud Connectors.
  - Quick Deploy automatically creates a resource location containing two Cloud Connectors when you create your first catalog.
  - In Full Configuration, creating a resource location and adding Cloud Connectors are separate steps that you must complete in Citrix Cloud before creating a catalog.
- Images used to create catalogs.
  - Quick Deploy offers several Citrix prepared images of Windows and Linux machines. You can use these images to create catalogs. You can also use these images to create images and then customize the new images to suit your unique deployment needs. This feature is known as the image builder. You can also import images from your own Azure subscription.
  - In Full Configuration, you customize images from the supported host you're using. Citrix prepared images are not available.
- Catalog displays:
  - Catalogs created in Quick Deploy are visible in the Quick Deploy and Full Configuration displays.
  - Catalogs created in Full Configuration are not visible in the Quick Deploy display.
- Delivery groups and application groups.



- You do not create delivery groups in Quick Deploy. In Quick Deploy, you specify the machines, applications, desktops, and users (subscribers) in the catalog. Citrix automatically creates a delivery group for each Quick Deploy catalog, using the same name as the catalog. That action occurs behind the scenes. You don't need to do anything to create the delivery group. The delivery group appears only in the Full Configuration interface, not in Quick Deploy.
- In Full Configuration, you create a delivery group and indicate which machines it contains. Optionally, you also specify applications, desktops, and users. (Or, you can specify users and applications in the Citrix Cloud library.) You can also create application groups.
- Layout and user interface.
  - The Quick Deploy interface has a different layout and style from Full Configuration. Quick Deploy contains more on-screen guidance.

The interfaces are not mutually exclusive. You can use Quick Deploy to create some catalogs, and then use Full Configuration to create other catalogs.

**Important:**

You can use Full Configuration to edit a catalog created with Quick Deploy. However, if you use Full Configuration to change that catalog or its corresponding delivery group, you can no longer edit that catalog in the Quick Deploy interface.

Similarly, if you create a catalog in Full Configuration, you cannot edit that catalog in Quick Deploy. That catalog does not appear in the Quick Deploy interface.

After a catalog created in Quick Deploy is edited in Full Configuration, it still appears in the Quick Deploy display, but cannot be managed or deleted in that interface. You must use Full Configuration to make further changes to that catalog.

The Quick Deploy interface is easy to use, especially if you don't need the advanced features of the Full Configuration interface.

Although the names are similar, Quick Deploy is not the same as the Quick Create way to create catalogs in Quick Deploy.

### **Replacement of earlier Azure Quick Deploy interface**

Quick Deploy replaces an earlier interface named Azure Quick Deploy. The Quick Deploy display includes all the catalogs you created using Azure Quick Deploy.

If you started creating a catalog in Azure Quick Deploy, but did not finish it, that catalog appears in the Quick Deploy catalog list. However, the only available action in Quick Deploy is to delete it.

## Requirements

- Quick Deploy supports only Azure workloads. It is not available with any other cloud host types, services, or hypervisors.
- Quick Deploy is available only in the Citrix Virtual Apps and Desktops service Standard for Azure, Premium, and Advanced editions, and Workspace Premium Plus.
- You must have a Citrix Cloud account and a subscription to the Citrix Virtual Apps and Desktops service.
- If you ordered the [Citrix Managed Azure Consumption Fund](#), you can use a Citrix Managed Azure subscription when you create catalogs and images.

If you did not order the Consumption Fund (or prefer to use your own Azure subscription), you must have an Azure subscription.

- You must have appropriate permission in the Virtual Apps and Desktops service to see the **Manage** tab. For details, see [Delegated administration](#).

### Important:

To ensure that you get important information about Citrix Cloud and the Citrix services you subscribe to, make sure you can receive all email notifications. For example, Citrix sends monthly informational notification emails detailing your Azure consumption (usage).

In the upper right corner of the Citrix Cloud console, expand the menu to the right of the customer name and OrgID fields. Select **Account Settings**. On the **My Profile** tab, select all entries in the **Email Notifications** section.

## Citrix Gateway consideration

If you use your own Citrix Gateway, it must have access to the VNet specified in the catalog creation wizard. A VPN can provide that access.

The Citrix Gateway Service works automatically with Quick Deploy catalogs.

## What's next

Follow the Quick Deploy setup guidance in [Get started](#).

After setting up your deployment using Quick Deploy, you can continue using that interface for the following management tasks.

- [Manage the catalog](#). Catalog management includes adding or deleting machines, managing apps, and managing power management schedules.

- [Manage images](#). Image management includes preparing or importing images, updating catalogs with a new image, renaming or deleting images, and installing or upgrading VDAs on an image.
- [Add or remove users in a catalog](#).
- [Manage resource locations](#).

## Get started with Quick Deploy

July 21, 2021

This article summarizes the setup tasks for delivering desktops and apps using the service's Quick Deploy interface. We recommend that you review each procedure before actually doing it, so you know what to expect.

To use Quick Deploy to set up a Remote PC Access deployment, see [Remote PC Access](#).

### Setup task summary

The following sections of this article guide you through setup tasks:

1. Review and complete necessary tasks in system requirements and preparation.
2. Set up a quick proof of concept deployment or a production deployment.
3. Provide the workspace URL to your users.

### System requirements and preparation

- [Sign up for Citrix Cloud and the Citrix Virtual Apps and Desktops service](#).

Also, if you plan to use [Citrix Managed Azure](#), make sure to order the Citrix Azure Consumption Fund (in addition to the service), either through Citrix or Azure Marketplace.

- **Windows licensing:** Ensure that you are properly licensed for Remote Desktop Services to run either Windows Server workloads or Azure Virtual Desktop Licensing for Windows 10. For more information, see [Configure a Microsoft RDS license server](#).
- If you plan to use a Citrix Managed Azure subscription, and want to join VDAs to a domain using Active Directory Group Policy, you must be an administrator with permission to perform that action in Active Directory. For details, see [Customer responsibility](#).
- Configuring connections to your corporate on-premises network has extra requirements.
  - Any connection (Azure VNet peering or SD-WAN): [Requirements for all connections](#).
  - Azure VNet peering connections: [VNet peering requirements and preparation](#).

- SD-WAN connections: [SD-WAN connection requirements and preparation](#).
- If you plan to use your own Azure images when creating a catalog, those [images must meet certain requirements](#).
- Internet connectivity requirements: [System and connectivity requirements](#).
- Resource limits in a service deployment: [Limits](#).

### Supported operating systems

When using Quick Deploy with a Citrix Managed Azure subscription:

- Windows 7 (VDA must be 7.15 LTSR with latest Cumulative Update)
- Windows 10 single-session
- Windows 10 multi-session
- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2019
- Red Hat Enterprise Linux and Ubuntu

When using Quick Deploy with a customer-managed Azure subscription:

- Windows 7 (VDA must be 7.15 LTSR with latest Cumulative Update)
- Windows 10 Enterprise single-session
- Windows 10 Enterprise Virtual Desktop multi-session
- Windows Server 2008 R2
- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2019
- Red Hat Enterprise Linux and Ubuntu

### Set up a quick proof of concept deployment

This procedure requires a Citrix Managed Azure subscription.

1. [Create a catalog using quick create](#).
2. [Add your users to the Managed Azure AD](#).
3. [Add your users to the catalog](#).
4. Notify your users of the Workspace URL.

### Set up a production deployment

1. If you're using your own Active Directory or Azure Active Directory to authenticate users, [connect and set that method in Citrix Cloud](#).

2. If you're using domain-joined machines, [verify that you have valid DNS server entries](#).
3. If you're using your own Azure subscription (instead of a Citrix Managed Azure subscription), [add your Azure subscription](#).
4. [Create or import an image](#). Although you can use one of the Citrix prepared images as-is in a catalog, they're intended primarily for proof of concept deployments.
5. If you're using a Citrix Managed Azure subscription, and want your users to be able to access items in your network (such as file servers), set up an [Azure VNet peering](#) or [Citrix SD-WAN](#) connection.
6. [Create a catalog using custom create](#).
7. If you're creating a catalog of multi-session machines, [add apps to the catalog](#), if needed.
8. If you're using the Citrix Managed Azure AD to authenticate your users, [add users to the directory](#).
9. [Add users to the catalog](#).
10. Notify your users of the Workspace URL.

After you set up the deployment, use the **Quick Deploy > Monitor** dashboard to see [desktop usage](#), [sessions](#), and [machines](#).

## Workspace URL

After you create catalogs and assign users, notify users where to find their desktops and apps: the Workspace URL. The Workspace URL is the same for all catalogs and users.

The Workspace URL is available in two locations:

- From **Manage > Quick Deploy** in the service, view the URL by expanding **User Access & Authentication** on the right.
- From the Citrix Cloud console, select **Workspace Configuration** from the upper left menu. The **Access tab** contains the Workspace URL.

For information about customizing the Workspace URL, see [Customize the Workspace URL](#).

After users navigate to the Workspace URL and authenticate, they can start their desktops and apps.

## Get help

- Review the [Troubleshoot](#) article.
- If you still have problems with the service, open a ticket by following the instructions in [How to Get Help and Support](#).

## Create catalogs using Quick Deploy

July 22, 2021

Use the procedures in this article to create a catalog of Microsoft Azure machines using the Quick Deploy management interface.

Review the entire procedure before creating a catalog, so you know what to expect.

To create a catalog using the Full Configuration interface, see [Create machine catalogs](#).

### Machine types

A Quick Deploy catalog can contain one of the following types of machines:

- **Static:** The catalog contains single-session static machines (also known as personal, dedicated, or persistent desktops). Static means that when a user starts a desktop, that desktop “belongs” to that user. Any changes that that user makes to the desktop are retained at logoff. Later, when that user returns to Citrix Workspace and starts a desktop, it is the same desktop.
- **Random:** The catalog contains single-session random machines (also known as non-persistent desktops). Random means that when a user starts a desktop, any changes that that user makes to that desktop are discarded after logoff. Later, when that user returns to Citrix Workspace and starts a desktop, it might or might not be the same desktop.
- **Multi-session:** The catalog contains machines with apps and desktops. More than one user can access each of those machines simultaneously. Users can launch a desktop or apps from their workspace. App sessions can be shared. Session sharing is not permitted between an app and a desktop.
  - When you create a multi-session catalog, you select the work load: light (such as data entry), medium (such as office apps), heavy (such as engineering), or custom. Each option represents a specific number of machines and sessions per machine, which yields the total number of sessions that the catalog supports.
  - If you select the custom work load, you then select from available combinations of CPUs, RAM, and storage. Type the number of machines and sessions per machine, which yields the total number of sessions that the catalog supports.

When deploying desktops, the static and random machine types are sometimes called “desktop types”.

### Ways to create a catalog using Quick Deploy

There are several ways to create and configure a catalog:

- **Quick create** is the fastest way to get started. You provide minimal information, and the service takes care of the rest. A quick create catalog is great for a test environment or proof of concept.
- **Custom create** allows more configuration choices than quick create. It's more suited to a production environment than a quick create catalog.
- **Remote PC Access** catalogs contain existing machines (usually physical) that users access remotely. For details and instructions about these catalogs, see [Remote PC Access](#).

Here's a comparison of quick create and custom create:

Quick create	Custom create
Less information to provide.	More information to provide.
Fewer choices for some features.	More choices for some features.
Citrix-managed Azure Active Directory user authentication.	Choice of: Citrix-managed Azure Active Directory, or your Active Directory/Azure Active Directory.
No connection to your on-premises network.	Choice of: No connection to your on-premises network, Azure VNet peering, and SD-WAN.
Uses a Citrix prepared Windows 10 image. That image contains a current desktop VDA.	Choice of: Citrix prepared images, your images that you import from Azure, or images you've built in the service from a Citrix prepared or imported image.
Each desktop has Azure standard disk (HDD) storage.	Several storage options are available.
Static desktops only.	Static, random, or multi-session desktops.
A power management schedule cannot be configured during creation. The machine hosting the desktop powers off when the session ends. (You can change this setting later.)	A power management schedule can be configured during creation. (A Quick Deploy power management schedule differs from a power management schedule you can create using the Full Configuration management interface.)
Must use a <a href="#">Citrix Managed Azure</a> subscription.	Can use the Citrix Managed Azure subscription or your own Azure subscription.

For procedure details, see:

- Create a Quick Deploy catalog using quick create
- Create a Quick Deploy catalog using custom create

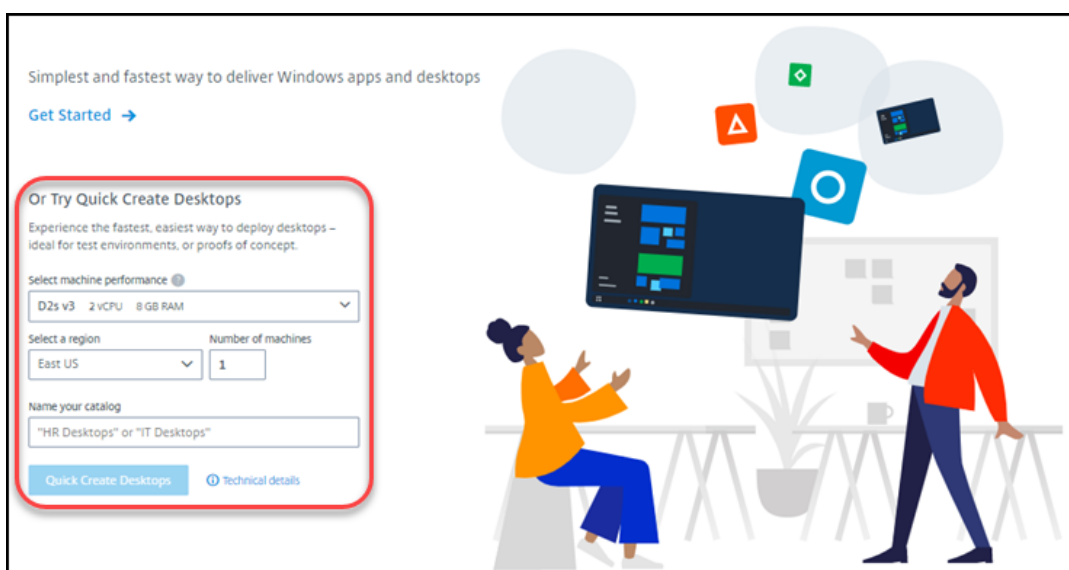
**Important:**

When you create a catalog (or an image) using a Citrix Managed Azure subscription for the first time, you are asked to acknowledge and consent to your responsibility for charges incurred. Reminders of that consent can also appear when creating more catalogs or images using the Citrix Managed Azure subscription.

**Create a Quick Deploy catalog using quick create**

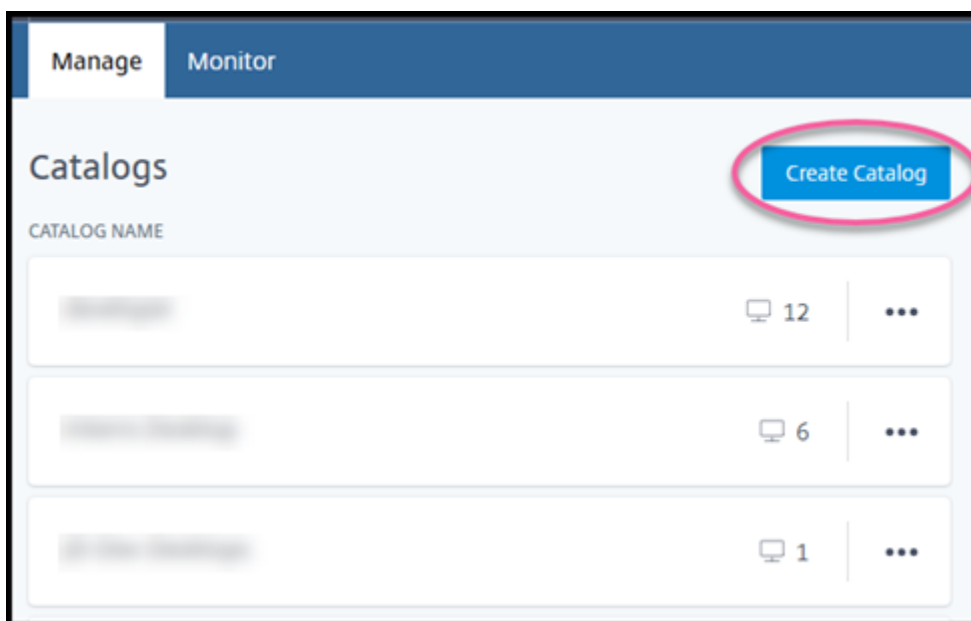
The quick create method uses a Citrix Managed Azure subscription and a Citrix prepared Windows 10 image to create a catalog containing static machines. Power management settings use the Cost Saver preset values. There is no connection to your corporate network. Users must be added using Citrix Managed Azure AD.

1. Sign in to [Citrix Cloud](#).
2. In the upper left menu, select **My Services > Virtual Apps and Desktops**.
3. Select **Manage > Quick Deploy**.
4. If a catalog has not yet been created, you're taken to the **Welcome** page. Choose one of:
  - Configure the catalog on this page. Continue with steps 6 through 10.




- Select **Get Started**. You're taken to the **Manage > Quick Deploy** dashboard. Select **Create Catalog**.
5. If a catalog has already been created (and you're creating another one), you're taken to the **Manage > Quick Deploy** dashboard. Select **Create Catalog**.





6. Select **Quick Create** at the top of the page, if it is not already selected.



- **Machine performance:** Select the machine type. Each choice has a unique combination of CPUs, RAM, and storage. Higher-performance machines have higher monthly costs.
- **Region:** Select a region where you want the machines created. You might select a region that's close to your users.

- **Name:** Type a name for the catalog. This field is required, and there is no default value.
  - **Number of machines:** Type the number of machines you want.
7. When you're done, select **Create Catalog**. (If you're creating the first catalog from the **Welcome** page, select **Quick Create Desktops**.)
  8. If this is the first catalog you're creating using a Citrix Managed Azure subscription, when prompted, acknowledge your responsibility for related charges.

While the catalog is being created, the catalog's name is added to the list of catalogs, indicating its progress through creation.

The service also automatically creates a resource location and adds two Citrix Cloud Connectors.

What to do next:

- You can [add users to the Managed Azure AD directory](#) while the catalog is being created.
- After the catalog is created, [add users to the catalog](#).

### Create a Quick Deploy catalog using custom create

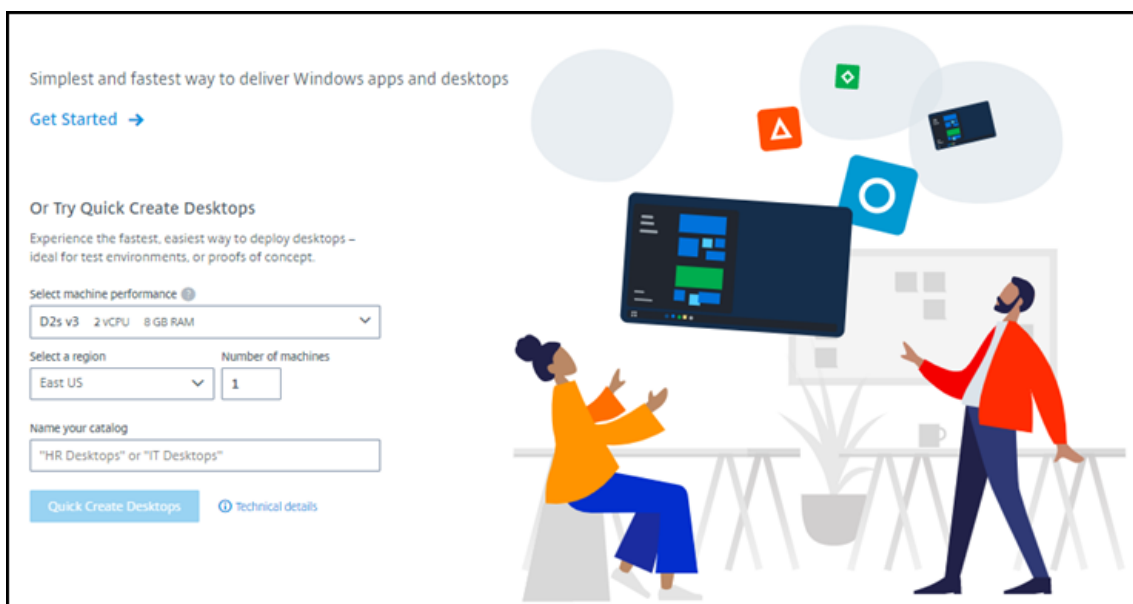
If you are using a Citrix Managed Azure subscription, and plan to use a connection to your on-premises network resources, [create that network connection](#) before creating the catalog. To allow your users access to your on-premises or other network resources, you also need Active Directory information for that location.

If you do not have a Citrix Managed Azure subscription, you can:

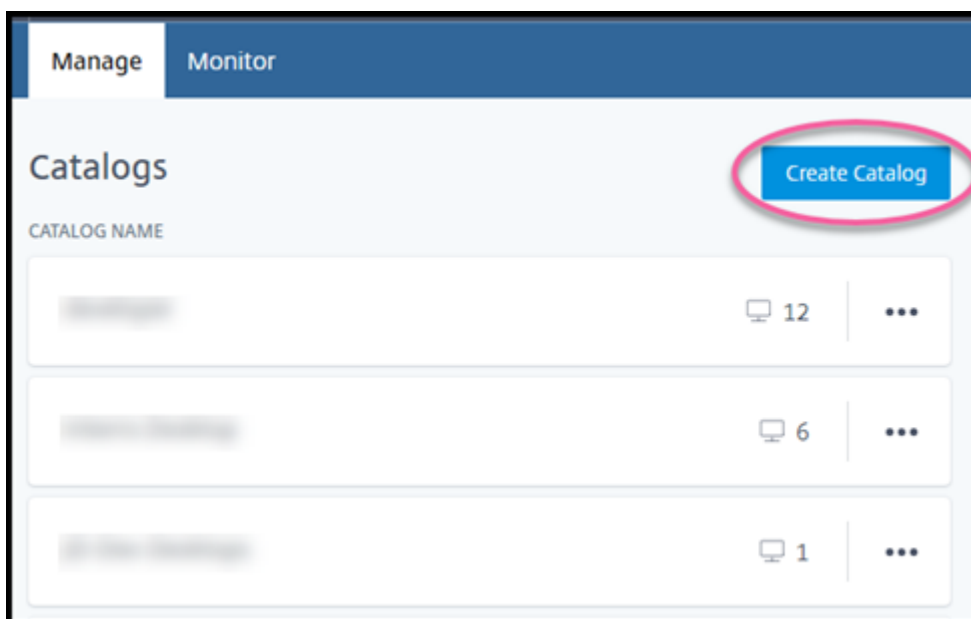
- [Order the Azure Consumption Fund](#) through Azure Marketplace, which provides you with a Citrix Managed Azure subscription.
- [Import \(add\) one or more of your own Azure subscriptions](#) to the service before creating a catalog.

To create a catalog:

1. Sign in to [Citrix Cloud](#).
2. In the upper left menu, select **My Services > Virtual Apps and Desktops**.
3. Select **Manage > Quick Deploy**.
4. If a catalog has not yet been created, you're taken to the **Welcome** page. Select **Get Started**. At the end of the introduction page, you're taken to the **Manage > Quick Deploy** dashboard. Select **Create Catalog**.



If a catalog has already been created, you're taken to the **Manage > Quick Deploy** dashboard. Select **Create Catalog**.



5. Select **Custom Create** at the top of the page, if it's not already selected.

The screenshot shows the 'Custom Create' configuration page for Citrix Virtual Apps and Desktops. The page is divided into several sections:

- Machine type:** Radio buttons for Multi-session (selected), Static (personal desktops), and Random (pooled desktops).
- Subscription:** A dropdown menu showing 'Citrix Managed'.
- Select a master Image:** A dropdown menu showing 'Win 2016 Server + VDA 2009'.
- Network connection:** A dropdown menu showing 'No connectivity to corporate network'.
- Region:** A dropdown menu showing 'East US'.
- Qualify for Linux compute rates?** Radio buttons for Yes (selected) and No.
- Select a machine:** A section containing:
  - Storage type:** A dropdown menu showing 'Standard disks (HDD)'.
  - Work Load:** A dropdown menu showing 'Light 16 sessions (D2s v2, 2 vCPU, 8 GB RAM)'.
  - Table:**

Machines	Sessions per machine	Total sessions
1	16	16

6. Complete the following fields. (Some fields are valid only for certain machine types. The field order might differ.)

- **Machine type.** Select a machine type. For details, see Machine types.
- **Subscription.** Select an [Azure subscription](#).
- **Master image:** Select an operating system [image](#) to be used for the catalog's machines.
- **Network connection:** Select the [network connection](#) to use for accessing resources in your network.

If you selected a Citrix Managed Azure subscription, the choices are:

- **No Connectivity:** Users cannot access locations and resources on your on-premises corporate network.
- **Connections:** Select a previously created connection, such as a VNet peering or SD-WAN connection.

If you selected a customer-managed Azure subscription, select the appropriate resource group, virtual network, and subnet.

- **Region:** (Available only if you selected **No Connectivity** in **Network connection**.) Select a region where you want the desktops created. You might select a region that's close to your users.

If you selected a connection in **Network connection**, the catalog uses that network's region.

- **Qualify for Linux compute rates?** (Available only if you selected a Windows image.) You can save money when you use your eligible license or Azure Hybrid Benefit.

**Windows Virtual Desktop benefit:** Eligible Windows 10 or Windows 7 per user licenses for:

- Microsoft 365 E3/ES
- Microsoft 365 A3/AS/Student Use Benefits
- Microsoft 365 F3
- Microsoft 365 Business Premium
- Windows 10 Enterprise E3/E5
- Windows 10 Education A3/A5
- Windows 10 VDA per user

Per user or per device license of RDS CAL with Software Assurance for Windows Server workloads.

**Azure Hybrid benefit:** Windows Server licenses with active Software Assurance or the equivalent qualifying subscription licenses. See <https://azure.microsoft.com/en-us/pricing/hybrid-benefit/>.

- **Machine:**
  - **Storage type.** HDD or SSD.
  - **Machine performance** (for **Static** or **Random** machine type), or **Workload** (for multi-session machine type). Choices include only options that match the generation type (gen1 or gen2) of the image you selected.

If you select the custom work load, type the number of machines and sessions per machine in the **Machine Performance** field.
  - **Machines.** How many machines you want in this catalog.
- **Machine naming scheme:** See Machine naming scheme.
- **Name:** Type a name for the catalog. This name appears on the **Manage** dashboard.
- **Power schedule:** By default, the **I'll configure this later** check box is selected. For details, see [Power management schedules](#). (This power management schedule differs from power management features available in the service's Full Configuration management interface.)

- **Join the local Active Directory domain:** (Available only if you selected an Azure VNet peering connection in **Network connection**.) Select **Yes** or **No**. If you select **Yes**, enter the:
  - FQDN of the domain (for example, Contoso.com).
  - Organization Unit: To use the default OU (Computers), leave this field empty.
  - Service account name: Must be a domain or enterprise administrator in the format name@domain or domain\name.
  - Password for the service account name.
- **Advanced settings:** See Resource location settings when creating a catalog.

7. When you're done, select **Create Catalog**.

8. If this is the first catalog you're creating using a Citrix Managed Azure subscription, when prompted, acknowledge your responsibility for related charges.

The **Manage > Quick Deploy** dashboard indicates when your catalog is created. The service also automatically creates a resource location and adds two Citrix Cloud Connectors.

What to do next:

- If you haven't done it already, [configure the authentication method](#) for your users to authenticate to Citrix Workspace.
- After the catalog is created, [add users to the catalog](#).
- If you created a multi-session catalog, [add applications](#) (before or after adding users).

## Resource location settings when creating a catalog

When creating a catalog, you can optionally configure several resource location settings.

When you select **Advanced settings** in the catalog creation dialog, the service retrieves resource location information.

- If you already have a resource location for the domain and network connection selected for the catalog, you can save it for use by the catalog you're creating.

If that resource location has only one Cloud Connector, another one is installed automatically. You can optionally specify advanced settings for the Cloud Connector you're adding.

- If you don't have a resource location set up for the domain and network connection selected for the catalog, you're prompted to configure one.

Configure advanced settings:

- (Required only when the resource location is already set up.) A name for the resource location.
- External connectivity type: through the Citrix Gateway service, or from within your corporate network.

- Cloud Connector settings:
  - (Available only when using a customer-managed Azure subscription) Machine performance. This selection is used for the Cloud Connectors in the resource location.
  - (Available only when using a customer-managed Azure subscription) Azure resource group. This selection is used for the Cloud Connectors in the resource location. The default is the resource group last used by the resource location (if applicable).
  - Organizational Unit (OU). The default is the OU last used by the resource location (if applicable).

When you're done with the advanced settings, select **Save** to return to the catalog creation dialog.

After you create a catalog, several resource location actions are available. For details, see [Resource location actions](#).

## Machine naming scheme

To specify a machine naming scheme when creating a catalog, select **Specify machine naming scheme**. Use from 1-4 wildcards (hash marks) to indicate where sequential numbers or letters appear in the name. Rules:

- The naming scheme must contain at least one wildcard, but not more than four wildcards. All the wildcards must be together.
- The entire name, including wildcards, must be between 2 and 15 characters.
- A name cannot include blanks (spaces), slashes, backslashes, colons, asterisks, angle brackets, pipes, commas, tildes, exclamation points, at signs, dollar signs, percent signs, carets, parentheses, braces, or underscores.
- A name cannot begin with a period.
- A name cannot contain only numbers.
- Do not use the following letters at the end of a name: **-GATEWAY**, **-GW**, and **-TAC**.

Indicate whether the sequential values are numbers (0-9) or letters (A-Z).

For example, a naming scheme of **PC-Sales-####** (with **0-9** selected) results in computer accounts named **PC-Sales-01**, **PC-Sales-02**, **PC-Sales-03**, and so on.

Leave enough room for growth.

- For example, a naming scheme with 2 wildcards and 13 other characters (for example, **MachineSales-####**) uses the maximum number of characters (15).
- Once the catalog contains 99 machines, the next machine creation fails. The service tries to create a machine with three digits (100), but that would create a name with 16 characters. The maximum is 15.

- So, in this example, a shorter name (for example, `PC-Sales-####`) allows scaling beyond 99 machines.

If you do not specify a machine naming scheme, the service uses the default naming scheme `DAS %%%%-**-#####`.

- %%%%= five random alphanumeric characters matching the resource location prefix
- \*\* = two random alphanumeric characters for the catalog
- ##### = three digits.

## Related information

- [Remote PC Access catalogs](#)
- [Create a catalog in a network that uses a proxy server](#)
- [Display catalog information](#)
- [Manage catalogs in Quick Deploy](#)

## Manage catalogs in Quick Deploy

July 28, 2021

This article describes the catalog management tasks you can use to manage catalogs that were created in Quick Deploy.

Remember: If you used Quick Deploy to create a catalog, and then use the Full Configuration interface to perform any management tasks on that catalog, you can no longer use the Quick Deploy interface for that catalog.

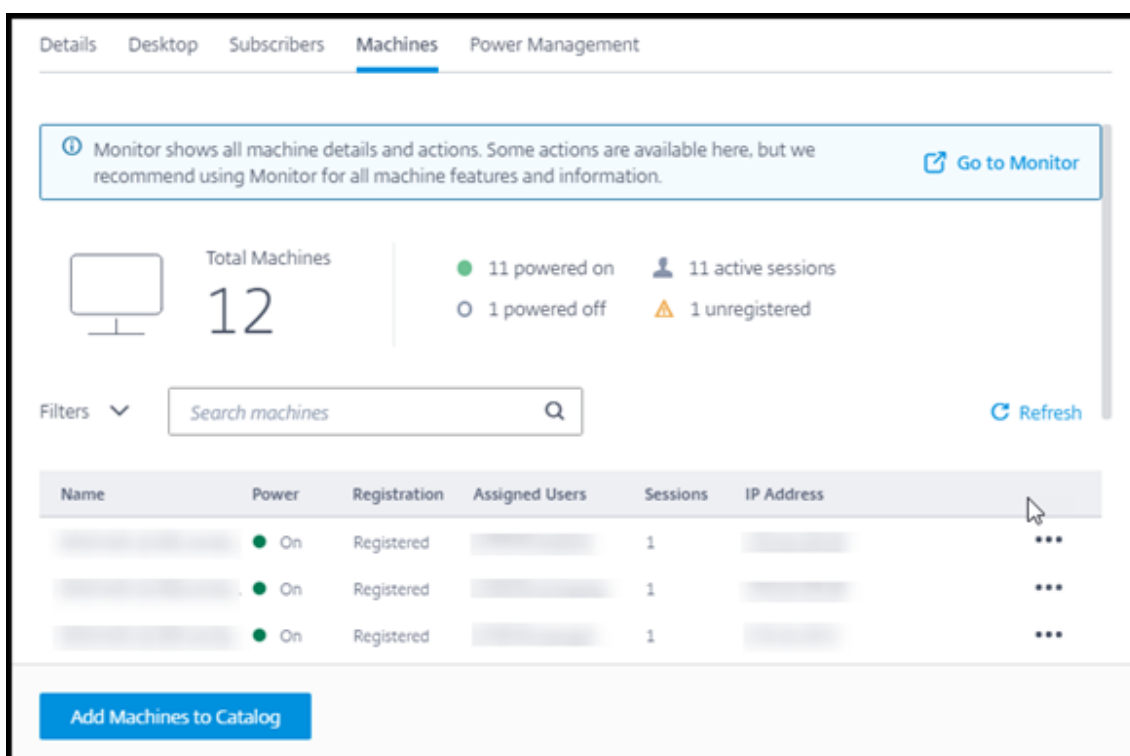
(For information about managing catalogs in the Full Configuration management interface, see [Manage machine catalogs](#).)

### Add machines to a catalog

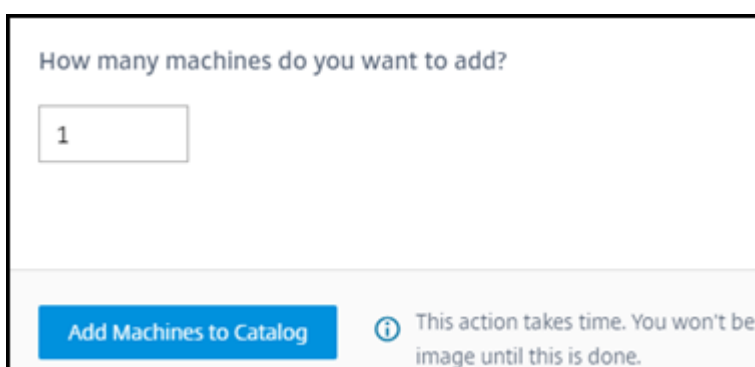
While machines are being added to a Quick Deploy catalog, you cannot make any other changes to that catalog.

1. From **Manage > Quick Deploy**, click anywhere in the catalog's entry.
2. On the **Machines** tab, select **Add Machines to Catalog**.





3. Enter the number of machines you want to add to the catalog.



4. (Valid only if the catalog is domain-joined.) Type the user name and password for the service account.
5. Select **Add Machines to Catalog**.

You cannot reduce the machine count for a catalog. However, you can use power management schedule settings to control how many machines are powered on.

### Change the number of sessions per machine

Changing the number of sessions per multi-session machine can affect users' experience. Increasing this value can reduce the compute resources allocated to concurrent sessions.

Recommendation: Observe your usage data to determine the appropriate balance between user experience and cost.

1. From **Manage > Quick Deploy**, select a catalog containing multi-session machines.
2. On the **Details** tab, select **Edit** next to **Sessions per Machine**.
3. Enter a new number of sessions per machine.
4. Select **Update Number of Sessions**.
5. Confirm your request.

This change does not affect current sessions. When you change the maximum number of sessions to a value that is lower than a machine's currently active sessions, the new value is implemented through the normal attrition of active sessions.

If a failure occurs before the update process begins, the catalog's **Details** display retains the correct number of sessions. If a failure occurs during the update process, the display indicates the number of sessions you wanted.

## Manage machines in a catalog

### Note:

Many of the actions that are available from **Manage > Quick Deploy** are also available from the **Monitor** tab in Quick Deploy.

To select actions from **Manage > Quick Deploy**:

1. From **Manage > Quick Deploy**, click anywhere in a catalog's entry.
2. On the **Machines** tab, find the machine you want to manage. In the ellipsis menu for that machine, select the desired action:
  - **Restart:** Restarts the selected machine.
  - **Start:** Starts the selected machine. This action is available only if the machine is powered off.
  - **Shutdown:** Shut down the selected machine. This action is available only if the machine is powered on.
  - **Turn maintenance mode on/off:** Turns maintenance mode on (if it is off) or off (if it is on) for the selected machine. By default, maintenance mode is turned off for a machine.

Turning on maintenance mode prevents new connections from being made to that machine. Users can connect to existing sessions on that machine, but they cannot start new sessions on that machine.

You might place a machine in maintenance mode before applying patches, or for troubleshooting.

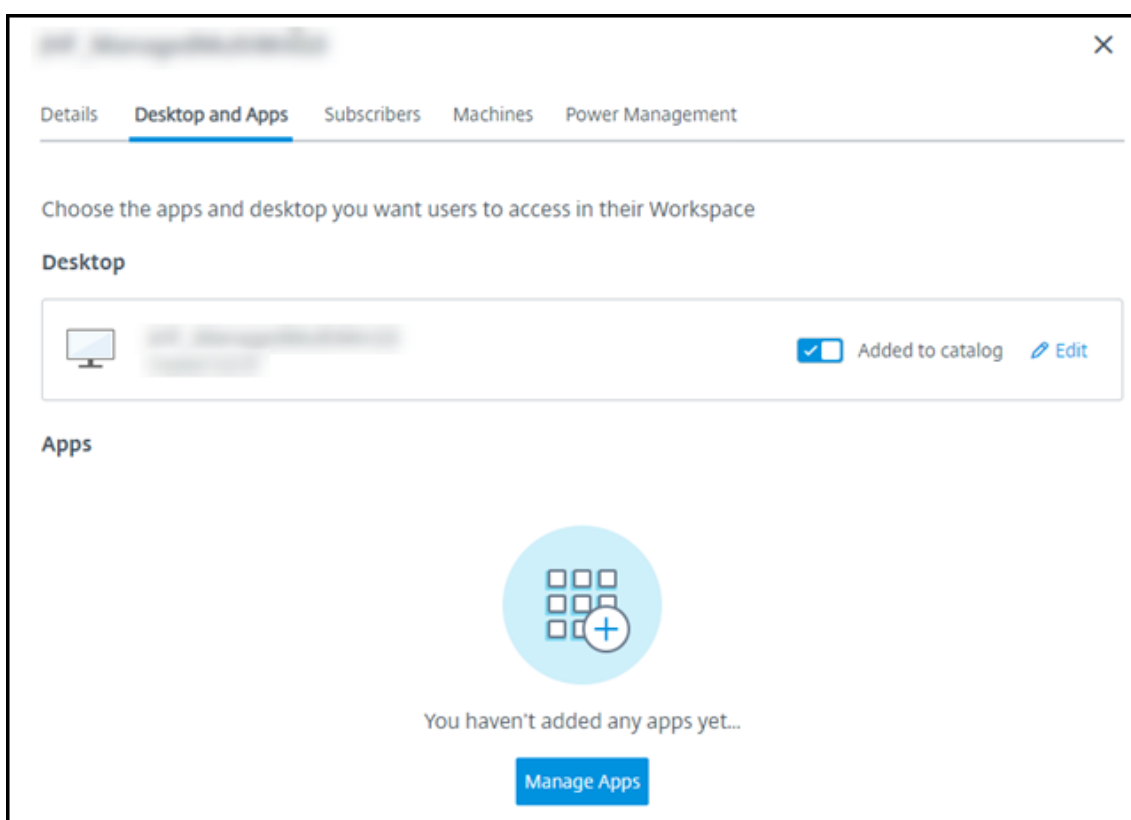
- **Delete:** Deletes the selected machine. This action is available only when the machine's session count is zero. Confirm the deletion.

When a machine is deleted, all data on the machine is removed.

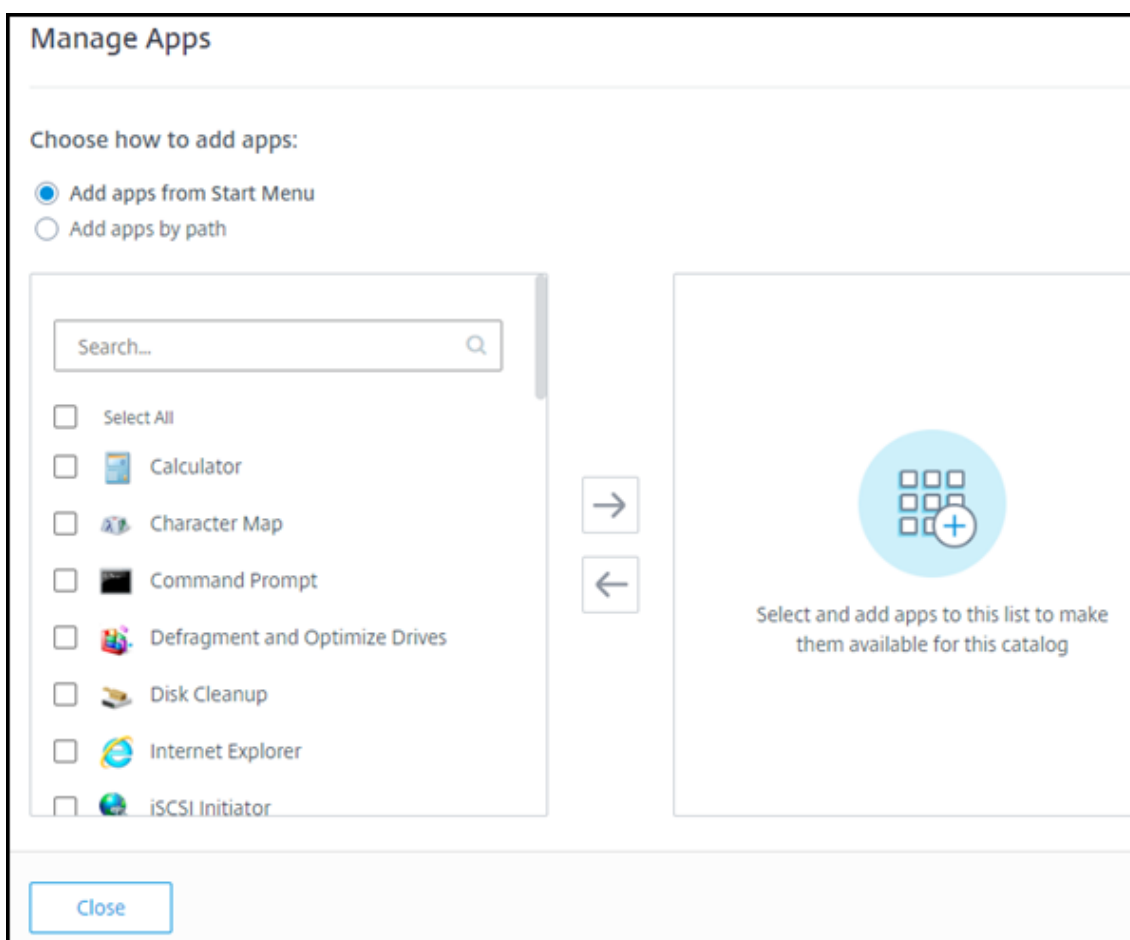
- **Force restart:** Forces a restart of the selected machine. Select this action only if a **Restart** action for the machine fails.

## Add apps to a catalog

1. From **Manage > Quick Deploy**, click anywhere in the catalog's entry.
2. On the **Desktop and Apps** tab, select **Manage Apps**.



3. Select how you are adding apps: from the **Start** menu of machines in the catalog, or from a different path on the machines.
4. To add apps from the **Start** menu:



- Select available apps in the left column. (Use **Search** to tailor the apps list.) Select the right arrow between the columns. The selected apps move to the right column.
- Similarly, to remove apps, select them in the right column. Select the left arrow between columns.
- If the **Start** menu has more than one version of the same app, with the same name, you can add only one. To add another version of that app, edit that version to change its name. Then you can add that version of the app.

5. To add apps by path:

- Enter the name for the app. This is the name users see in Citrix Workspace.
- The icon shown is the icon users see in Citrix Workspace. To select another icon, select **Change icon** and navigate to the icon you want to display.
- (Optional) Enter a description of the application.
- Enter the path to the app. This field is required. Optionally, add command line parameters and the working directory. For details about command line parameters, see Pass parameters to published applications.


6. When you're finished, select **Close**.

On Windows Server 2019 VDAs, some application icons might not appear correctly during configuration and in the users' workspace. As a workaround, after the app is published, edit the app and use the **Change icon** feature to assign a different icon that displays correctly.

### Edit an app in a catalog

1. From **Manage > Quick Deploy**, click anywhere in the catalog's entry.

2. On the **Desktop and Apps** tab, click anywhere on the row containing the app you want to edit.
3. Select the pencil icon.



4. Type changes in any of the following fields:
  - **Name:** The name users see in Citrix Workspace.
  - **Description**
  - **Path:** The path to the executable.
  - **Command line parameters:** For details, see Pass parameters to published applications.
  - **Working directory**
5. To change the icon users see in their Citrix Workspace, select **Change icon** and navigate to the icon you want to display.
6. When you're done, select **Save**.

### Pass parameters to published applications

When you associate a published application with file types, the percent and star symbols (enclosed in double quotation marks) are appended to the end of the command line. These symbols act as a placeholder for parameters passed to user devices.

- If a published application does not launch when expected, verify that its command line contains the correct symbols. By default, parameters supplied by user devices are validated when the symbols are appended.

For published applications that use customized parameters supplied by the user device, the symbols are appended to the command line to bypass command-line validation. If you do not see these symbols in a command line for the application, add them manually.

- If the path to the executable file includes directory names with spaces (such as “C:\Program Files” ), enclose the command line for the application in double quotation marks to indicate that the space belongs in the command line. Add double quotation marks around the path, and another set of double quotation marks around the percent and star symbols. Add a space between the closing quotation mark for the path and the opening quotation mark for the percent and star symbols.

For example, the command line for the published application Windows Media Player is: “C:\Program Files\Windows Media Player\mplayer1.exe” “%\*”

## Remove apps from a catalog

Removing an app from a catalog does not remove it from the machines. It just prevents it from appearing in Citrix Workspace.

1. From **Manage > Quick Deploy**, click anywhere in the catalog’s entry.
2. On the **Desktop and Apps** tab, select the trash icon next to the apps you want to remove.

## Delete a catalog

When you delete a catalog, all the machines in the catalog are permanently destroyed. Deleting a catalog cannot be reversed.

1. From **Manage > Quick Deploy**, click anywhere in the catalog’s entry.
2. On the **Details** tab, select **Delete Catalog**.
3. Confirm the deletion.

To help identify residual Active Directory machine accounts that you must delete, you can download a list of machine and Cloud Connector names.

## Manage power management schedules

A power management schedule affects all machines in a catalog. A schedule provides:

- Optimal user experience: Machines are available for users when they’re needed.
- Security: Desktop sessions that remain idle for a specified interval are disconnected, requiring users to launch a new session in their workspace.
- Cost management and power savings: Machines with desktops that remain idle are powered-off. Machines are powered on to meet scheduled and actual demand.

You can configure a power schedule when you create a custom catalog or do it later. If no schedule is selected or configured, a machine powers off when a session ends.

You cannot select or configure a power schedule when creating a catalog with quick create. By default, quick create catalogs use the Cost Saver preset schedule. You can select or configure a different schedule later for that catalog.

Schedule management includes:

- Knowing what information a schedule contains
- Creating a schedule

### **Information in a schedule**

The following diagram shows the schedule settings for a catalog containing multi-session machines. Settings for a catalog containing single-session (random or static) machines differ slightly.



Details Desktop and Apps Subscribers Machines **Power Management**

Presets  
Cost Saver ▾

General

Disconnect desktop sessions when idle  
After 15 Minutes ▾

Log Off Disconnected Sessions  
After 15 Minutes ▾

Power Off Delay  
After 30 Minutes ▾

Work hours ⓘ

Time Zone  
(UTC-05:00) Eastern Time (US & Canada) ▾

Power on machines  
SUN MON TUE WED THU FRI SAT

Start End  
▾ ▾ ▾ ▾

Capacity buffer  
10 %

Minimum running machines  
1

After-hours ⓘ

Capacity buffer  
10 %

Minimum running machines  
1

Save Changes

A power management schedule contains the following information.

### Preset schedules

The service offers several preset schedules. You can also configure and save custom schedules. Although you can delete custom presets, you cannot delete Citrix-provided presets.

### Time zone

Used with the power-on machines setting to establish work hours and after hours, based on the selected time zone.

This setting is valid for all machine types.

### **Power on machines: Work hours and after hours**

The days of the week and start-stop hours of the day that form your work hours. This generally indicates the intervals when you want machines powered on. Any time outside of those intervals is considered after-hours. Several schedule settings allow you to enter separate values for work hours and after-hours. Other settings apply all the time.

This setting is valid for all machine types.

### **Disconnect desktop sessions when idle**

How long a desktop can remain idle (not used) before the session is disconnected. After a session is disconnected, the user must go to Workspace and start a desktop again. This is a security setting.

This setting is valid for all machine types. One setting applies all the time.

### **Power off idle desktops**

How long a machine can remain disconnected before it is powered off. After a machine is powered off, the user must go to Workspace and start a desktop again. This is a power-saving setting.

For example, let's say you want desktops to disconnect after they have been idle for 10 minutes. Then, power off the machines if they remain disconnected for another 15 minutes.

If Tom stops using his desktop and walks away for a one-hour meeting, the desktop will be disconnected after 10 minutes. After another 15 minutes, the machine will be powered off (25 minutes total).

From a user standpoint, the two idle settings (disconnect and power-off) have the same effect. If Tom stays away from his desktop for 12 minutes or an hour, he must start a desktop again from Workspace. The difference in the two timers affects the state of the virtual machine providing the desktop.

This setting is valid for single-session (static or random) machines. You can enter values for work hours and after-hours.

### **Log off disconnected sessions**

How long a machine can remain disconnected before the session is closed.

This setting is valid for multi-session machines. One setting applies all the time.

### **Power-off delay**

The minimum amount of time a machine must be powered-on before it is eligible for power-off (along with other criteria). This setting keeps machines from “flip-flopping” on and off during volatile session demands.

This setting is valid for multi-session machines, and applies all the time.

### **Minimum running machines**

How many machines must remain powered-on, regardless of how long they are idle or disconnected.

This setting is valid for random and multi-session machines. You can enter values for work hours and after-hours.

### **Capacity buffer**

A capacity buffer helps accommodate sudden spikes in demand, by keeping a buffer of machines powered-on. The buffer is specified, as a percentage of current session demand. For example, if there are 100 active sessions and the capacity buffer is 10%, the service provides capacity for 110 sessions. A spike in demand might occur during work hours or adding new machines to the catalog.

A lower value decreases the cost. A higher value helps ensure an optimized user experience. When launching sessions, users do not have to wait for extra machines to power on.

When there are more than enough machines to support the number of powered-on machines needed in the catalog (including the capacity buffer), extra machines are powered off. Power-off might occur because of off-peak time, session logoffs, or fewer machines in the catalog. The decision to power off a machine must meet the following criteria:

- The machine is powered on and not in maintenance mode.
- The machine is registered as available or waiting to register after power-on.
- The machine has no active sessions. Any remaining sessions have ended. (The machine was idle for the idle timeout period.)
- The machine has been powered on for at least “X” minutes, where “X” is the power-off delay specified for the catalog.

In a static catalog, after all machines in the catalog are assigned, the capacity buffer does not play a role in powering machines on or off.

This setting is valid for all machine types. You can enter values for work hours and after-hours.

### Create a power management schedule

1. From **Manage > Quick Deploy**, click anywhere in the catalog's entry.
2. On the **Power Management** tab, determine whether any of the preset schedules (in the menu at the top) meet your needs. Select a preset to see the values it uses. If you want to use a preset, leave it selected.
3. If you change the values in any fields (such as days, times, or intervals), the preset selection changes to **Custom** automatically. An asterisk indicates that custom settings have not been saved.
4. Set the values you want for the custom schedule.
5. Select **Custom** at the top, and then save the current settings as a new preset. Enter a name for the new preset and select the check mark.
6. When you're done, select **Save Changes**.

Later, you can edit or delete a custom preset by using the pencil or trash icons in the **Presets** menu. You cannot edit or delete common presets.

### Related information

- [Update a catalog with a new image](#)
- [Add and remove users in a catalog](#)

## Azure subscriptions in Quick Deploy

August 10, 2021

### Introduction

When you create a catalog or build an image in Quick Deploy, you choose among the available Azure subscriptions. Quick Deploy supports both Citrix Managed Azure subscriptions and your own, customer-managed Azure subscriptions.

- To use your own Azure subscription, you first import (add) one or more of those subscriptions to the service. That action enables the Citrix service to access your Azure subscriptions.
- Using a Citrix Managed Azure subscription requires no subscription configuration. However, a Citrix Managed Azure subscription is available only when you [order the Citrix Azure Consumption Fund](#), in addition to the Citrix Virtual Apps and Desktops service.

Some service features differ, depending on whether the catalog uses a Citrix Managed Azure subscription or in your own Azure subscription.

Citrix Managed Azure subscription	Your own Azure subscription
Supports domain-joined or non-domain-joined machines.	Supports only domain-joined machines.
Supports quick create and custom create catalogs.	Supports only custom create catalogs.
Always available when creating catalogs and images.	Must add the Azure subscription to the service before creating a catalog.
For user authentication, supports Citrix Managed Azure Active Directory or your own Active Directory.	Can connect your own Active Directory and Azure Active Directory.
Network connection options include <b>No connectivity</b> .	Network connection options include only your own virtual networks.
When using Azure VNet peering to connect to your resources, you must create a VNet peer connection in the service.	Select an existing virtual network.
When importing an image from Azure, you specify the image's URI.	When importing an image, you can select a VHD or browse storage in the Azure subscription.
Can create a bastion machine in customer's Azure subscription to troubleshoot machines.	No need to create a bastion machine because you can already access the machines in your subscription.

## View Azure subscriptions

To view Azure subscription details, from **Manage > Quick Deploy**, expand **Cloud Subscriptions** on the right. Then select a subscription entry.

- The **Details** page includes the number of machines, plus the numbers and names of catalogs and images using the subscription.
- The **Resource Locations** page lists the resource locations where the subscription is used.

## Add customer-managed Azure subscriptions

To use a customer-managed Azure subscription, you must add it to the Citrix service before creating a catalog or building an image that uses that subscription. You have two options when adding your Azure subscriptions:

- **If you are a Global Administrator for the directory and have contributor privileges for the subscription:** Simply authenticate to your Azure account.
- **If you are not a Global Administrator and have contributor privileges on the subscription:** Before adding the subscription to the service, create an Azure app in your Azure AD and then add that app as a contributor of the subscription. When you add that subscription to the service, you provide relevant app information.

### **Add customer-managed Azure subscriptions if you're a Global Administrator**

This task requires Global Administrator privileges for the directory, and contributor privileges for the subscription.

1. From **Manage > Quick Deploy**, expand **Cloud Subscriptions** on the right.
2. Select **Add Azure subscription**.
3. On the **Add Subscriptions** page, select **Add your Azure subscription**.
4. Select the button that allows the service to access your Azure subscriptions on your behalf.
5. Select **Authenticate Azure Account**. You're taken to the Azure sign-in page.
6. Enter your Azure credentials.
7. You're returned automatically to the service. The **Add Subscription** page lists the discovered Azure subscriptions. Use the search box to filter the list, if needed. Select one or more subscriptions. When you're done, select **Add Subscriptions**.
8. Confirm that you want to add the selected subscriptions.

The Azure subscriptions you selected are listed when you expand **Subscriptions**. The added subscriptions are available for selection when you create a catalog or image.

### **Add customer-managed Azure subscriptions if you're not a Global Administrator**

Adding an Azure subscription when you're not a global administrator is a two-part process:

- Before you add a subscription to the service, create an app in Azure AD and then add that app as a contributor of the subscription.
- Add the subscription to the service, using information about the app you created in Azure.

### **Create an app in Azure AD and add it as a contributor**

1. Register a new application in Azure AD:
  - a) From a browser, navigate to <https://portal.azure.com>.
  - b) In the upper left menu, select **Azure Active Directory**.
  - c) In the **Manage** list, select **App registrations**.
  - d) Select **+ New registration**.

- e) On the **Register an application** page, provide the following information:
    - **Name:** Enter the connection name
    - **Application type:** Select **Web app / API**
    - **Redirect URI:** leave blank
  - f) Select **Create**.
2. Create the application's secret access key and add the role assignment:
    - a) From the previous procedure, select **App Registration** to view details.
    - b) Make a note of the **Application ID** and **Directory ID**. You'll use this later when adding your subscription to the service.
    - c) Under **Manage**, select **Certificates & secrets**.
    - d) On the **Client secrets** page, select **+ New client secret**.
    - e) On the **Add a client secret** page, provide a description and select an expiration interval. Then select **Add**.
    - f) Make a note of the client secret value. You'll use this later when adding your subscription to the service.
    - g) Select the Azure subscription you want to link (add) to the service, and then select **Access control (IAM)**.
    - h) In the **Add a role assignment** box, select **Add**.
      - i) On the **Add role assignment** tab, select the following:
        - **Role:** Contributor
        - **Assign access to:** Azure AD user, group, or service principal
        - **Select:** The name of the Azure app you created earlier.
      - j) Select **Save**.

### Add your subscription to the service

You need the application ID, directory ID, and client secret value from the app you created in Azure AD.

1. From **Manage > Quick Deploy**, expand **Cloud Subscriptions** on the right.
2. Select **Add Azure subscription**.
3. On the **Add Subscriptions** page, select **Add your Azure subscriptions**.
4. Select **I have an Azure App with contributor role to the subscription**.
5. Enter the tenant ID (directory ID), client ID (application ID), and client secret for the app you created in Azure.

6. Select **Select your subscription** and then select the subscription you want.

Later, from the subscription's **Details** page in the service dashboard, you can update the client secret or replace the Azure app from the ellipsis menu.

If the service can't access an Azure subscription after it's added, several catalog power management and individual machine actions aren't allowed. A message provides an option to add the subscription again. If the subscription was originally added using an Azure app, you can replace the Azure app.

### Add Citrix Managed Azure subscriptions

A Citrix Managed Azure subscription supports a certain number of machines. (In this context, *machines* refers to VMs that have a Citrix VDA installed. These machines deliver apps and desktops to users. It does not include other machines in a resource location, such as Cloud Connectors.)

If your Citrix Managed Azure subscription is likely to reach its limit soon, and you have enough Citrix licenses, you can request another Citrix Managed Azure subscription. The dashboard contains a notification when you're close to the limit.

You can't create a catalog (or add machines to a catalog) if the total number of machines for all catalogs that use that Citrix Managed Azure subscription would exceed the limit.

For example, assume a hypothetical limit of 1,000 machines per Citrix Managed Azure subscription.

- Let's say you have two catalogs (**Cat1** and **Cat2**) that use the same Citrix Managed Azure subscription. **Cat1** currently contains 500 machines, and **Cat2** has 250.
- As you plan for future capacity needs, you add 200 machines to **Cat2**. The Citrix Managed Azure subscription now supports 950 machines (500 in **Cat 1** and 450 in **Cat 2**). The dashboard indicates that the subscription is near its limit.
- When you need 75 more machines, you can't use that subscription to create a catalog with 75 machines (or add 75 machines to an existing catalog). That would exceed the subscription limit. Instead, you request another Citrix Managed Azure subscription. Then, you can create a catalog using that subscription.

When you have more than one Citrix Managed Azure subscription:

- Nothing is shared between those subscriptions.
- Each subscription has a unique name.
- You can choose among the Citrix Managed Azure subscriptions (and any customer-managed Azure subscriptions that you've added) when:
  - Creating a catalog.
  - Building or importing an image.
  - Creating a VNet peering or SD-WAN connection.



Requirement:

- You must have enough Citrix licenses to warrant adding another Citrix Managed Azure subscription. Using the previous hypothetical example, if you have 2,000 Citrix licenses in anticipation of deploying at least 1,500 machines through Citrix Managed subscriptions, you can add another Citrix Managed Azure subscription.

To add a Citrix Managed Azure subscription:

1. Contact your Citrix representative to request another Citrix Managed Azure subscription. You are notified when you can proceed.
2. From **Manage > Quick Deploy**, expand **Cloud Subscriptions** on the right.
3. Select **Add Azure subscription**.
4. On the **Add Subscriptions** page, select **Add a Citrix Managed Azure subscription**.
5. On the **Add a Citrix Managed Subscription** page, select **Add Subscription** at the bottom of the page.

If you're notified that an error occurred during creation of a Citrix Managed Azure subscription, contact Citrix Support.

## Remove Azure subscriptions

Before you can remove an Azure subscription, you must delete all catalogs and images that use it.

If you have one or more Citrix Managed Azure subscriptions, you cannot remove all of them. At least one must remain.

1. From **Manage > Quick Deploy**, expand **Cloud Subscriptions** on the right.
2. Select the subscription entry.
3. On the **Details** tab, select **Remove Subscription**.
4. Select **Authenticate Azure Account**. You're taken to the Azure sign-in page.
5. Enter your Azure credentials.
6. You're returned automatically to the service. Confirm the deletion and then select **Yes, Delete Subscription**.

## Images in Quick Deploy

August 11, 2021

When you create a catalog to deliver desktops or apps, an image is used (with other settings) as a template for creating the machines.

Quick Deploy provides a set of prepared images that you can choose from to build and customize an image in Quick Deploy. You can also import (add) images from your own Azure subscription.

## Citrix prepared images

Quick Deploy provides several Citrix prepared images:

- Windows 10 Enterprise (single-session)
- Windows 10 Enterprise Virtual Desktop (multi-session)
- Windows 10 Enterprise Virtual Desktop (multi-session) with Office 365 ProPlus
- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2019
- Linux Ubuntu (single-session and multi-session)

The Citrix prepared images have a current Citrix Virtual Delivery Agent (VDA) and troubleshooting tools installed. The VDA is the communication mechanism between your users' machines and the Citrix Cloud infrastructure that manages the service. Images provided by Citrix have a **CITRIX** notation.

Citrix prepared images are not available in the service's Full Configuration interface.

You can also import and use your own image from Azure.

## Ways to use images in Quick Deploy

You can:

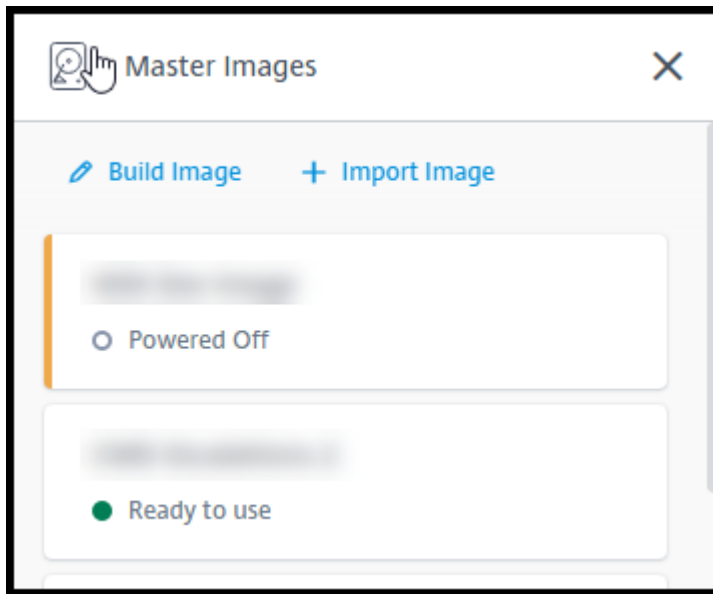
- **Use a Citrix prepared image when creating a catalog.** This choice is recommended only for proof of concept deployments.
- **Use a Citrix prepared image to create another image.** After the new image created, you customize it by adding applications and other software that your users need. Then, you can use that customized image when creating a catalog.
- **Import an image from Azure.** After you import an image from Azure, you can then use that image when creating a catalog.

Or, you can use that image to create a new image, and then customize it by adding apps. Then, you can use that customized image when creating a catalog.


When you create a catalog, the service verifies that the image uses a valid operating system, and has a Citrix VDA and troubleshooting tools installed (along with other checks).

## Display image information

1. From **Manage > Quick Deploy**, expand **Master Images** on the right. The display lists the Citrix prepared images, and any images you imported.



2. Select an image to display its details.



From the details card, you can:

- Change (edit) the image's name.
- Add and edit notes (available only for images you prepared or imported, not Citrix prepared images).
- Delete the image.

### Prepare a new image

Preparing a new image includes creating the image and then customizing it. When you create an image, a new VM is created to load the new image.

Requirements:

- Know the performance characteristics that the machines need. For example, running CAD apps might require different CPU, RAM, and storage than other office apps.

- If you plan to use a connection to your on-premises resources, set up that connection before creating the image and the catalog. For details, see [Network connections](#).

When using a Citrix prepared Ubuntu image to build a new image, a root password is created for the new image. You can change that root password, but only during the image creation and customization process. (You cannot change the root password after the image is used in a catalog.)

- When the image is created, the administrator account that you specified (**Login details for image building machine**) is added to the `sudoers` group.
- After you RDP to the machine containing the new image, launch the terminal application and type `sudo passwd root`. When prompted, provide the password you specified when creating the image. After verification, you're prompted to enter a new password for the root user.

To create an image:

1. From **Manage > Quick Deploy**, expand **Master Images** on the right.
2. Select **Build Image**.

The screenshot shows a configuration form for building a new master image. The form is titled "Name the new master image" and contains several sections:

- Name the new master image:** A text input field.
- Select a master image as base:** A dropdown menu with "Win 10 EVD (Multi-session) 1909 + Office 365 ProPlus + VE" selected.
- Subscription:** A dropdown menu with "Citrix Managed" selected.
- Network connection:** A dropdown menu with "No connectivity to corporate network" selected.
- Region:** A dropdown menu with "East US" selected.
- Set log-on credentials for the image machine:** A section with three text input fields: "Username", "Password", and "Confirm password".
- Performance (the machine that runs the image):** A dropdown menu with "D2s v3 2 vCPU 8 GB RAM" selected.
- Restricted IP access:** A section with a "+ Add IP addresses" link.
- Add Notes:** A text input field.

3. Enter values in the following fields:

- **Name:** Enter a name for the new image.
- **Master image:** Select an existing image. This is the base image that is used to create the new image.
- **Subscription:** Select an Azure subscription.
- **Network connection:**
  - If using a Citrix Managed Azure subscription, select **No connectivity** or a previously created connection.
  - If using your own customer-managed Azure subscription, select your resource group, virtual network, and subnet. Then add domain details: FQDN, OU, service account name, and credentials.
- **Region:** (Available only for **No connectivity**.) Select a region where you want the machine containing the image to be created.
- **Logon credentials for image machine:** You'll use these credentials later when you connect (RDP) to the machine containing the new image, so that you can install apps and other software.
- **Machine performance:** This is CPU, RAM, and storage information for the machine that runs the image. Select a machine performance that meets your apps' requirements.
- **Restricted IP access:** If you want to restrict access to specific addresses, select **Add IP addresses** and then enter one or more addresses. After adding the addresses, select **Done** to return to the **Build image** card.
- **Notes:** Optionally add up to 1024 characters of notes. After the image is created, you can update the notes from the image's details display.
- **Local domain join:** Indicate whether you want to join the local Active Directory domain.
  - If you select **Yes**, enter the FQDN, OU, service account name, and credentials.
  - If you select **No**, enter the credentials for the host machine.

4. When you're done, select **Build Image**.

An image can take up to 30 minutes to build. From **Manage > Quick Deploy**, expand **Master Images** on the right to see the current state (such as *Building image* or *Ready to customize*).

What to do next: Connect to a new image and customize it.

### **Connect to a new image and customize it**

After a new image is created, its name is added to the images list, with a status of *Ready to customize* (or similar wording). To customize that image, you first download an RDP file. When you use that file to connect to the image, you can then add applications and other software to the image.

1. From **Manage > Quick Deploy**, expand **Master Images** on the right. Select the image you want to connect to.
2. Select **Download RDP file**. An RDP client downloads.  
The image machine might power off if you do not RDP to it shortly after it's created. This saves costs. When that happens, select **Power On**.
3. Start the downloaded RDP client. It automatically attempts to connect to the address of the machine containing the new image. When prompted, enter the credentials you specified when creating the image.
4. After you connect to the machine, add or remove apps, install updates, and finish any other customization work.  
Do **NOT** Sysprep the image.
5. When you're done customizing the new image, return to the **Master Images** box and select **Finish build**. The new image automatically undergoes validation testing.

Later, when you create a catalog, the new image is included in the list of images you can select.

From **Manage > Quick Deploy**, the image display on the right indicates how many catalogs and machines use each image.

**Note:**

After you finalize an image, you cannot edit it. You must create a new image (optionally using the previous image as a starting point), and then update the new image.

## Import an image from Azure

When you import an image from Azure that has a Citrix VDA and applications your users need, you can use it to create a catalog or replace the image in an existing catalog.

### Imported image requirements

**Note:**

This service does not support importing disks that are associated with Azure generation 2 VMs.


Citrix runs validation tests on the imported image. Ensure that the following requirements are met when you prepare the image that you'll import into the service.

- **Supported OS:** The image must be a [supported OS](#). To check a Windows OS version, run `Get-WmiObject Win32_OperatingSystem`.
- **Supported generation:** Only generation 1 VMs are supported.
- **Not generalized:** The image must not be generalized.

- **No configured Delivery Controllers:** Ensure that no Citrix Delivery Controllers are configured in the image. Ensure that the following registry keys are cleared.
  - HKLM:\SOFTWARE\Citrix\VirtualDesktopAgent\ListOfDDCs
  - HKLM:\SOFTWARE\Policies\Citrix\VirtualDesktopAgent\ListOfDDCs
  - HKLM:\SOFTWARE\Citrix\VirtualDesktopAgent\FarmGUID
  - HKLM:\SOFTWARE\Policies\Citrix\VirtualDesktopAgent\FarmGUID
- **Personality.ini file:** The `personality.ini` file must exist on the system drive.
- **Valid VDA:** The image must have a Citrix VDA newer than 7.11 installed.
  - Windows: To check, use `Get-Childitem HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\Citrix Virtual Desktop Agent`. For installation guidance, see [Install a Windows VDA on an image](#).
  - Red Hat Enterprise Linux and Ubuntu: For installation guidance, see the [product documentation](#).
- **Azure Virtual Machine Agent:** Before importing an image, make sure that the Azure Virtual Machine Agent is installed on the image. For more information, see the Microsoft article [Azure Virtual Machine Agent overview](#).

### Import the image using Quick Deploy

1. From **Manage > Quick Deploy**, expand **Master Images** on the right.



2. Select **Import Image**.



The screenshot shows a web form titled "Choose how to import your image". It has two radio buttons at the top: "Browse storage account" (selected) and "Use Azure public URL". Below are several dropdown menus: "Subscription", "Choose resource group", "Storage account", and "Choose master image". There are two radio buttons for "Master image type": "Windows" (selected) and "Linux". A text input field is labeled "Name the new master image" with a placeholder "Eg. 'Windows 10 + My Apps'". At the bottom is a text area labeled "Add Notes" with a placeholder "Enter notes here (up to 1024 characters). You can see and change them in the image's details."

3. Choose how to import the image.

- For managed disks, use the export feature to generate a SAS URL. Set the expiration time to 7200 seconds or more.
- For VHDs in a storage account, choose one of the following:
  - Generate a SAS URL for the VHD file.
  - Update the access level of a block storage container to blob or container. Then, get the file's URL.

4. If you selected **Browse storage account**:

- a) Sequentially select a subscription > resource group > storage account > image.
- b) Name the image.

5. If you selected **Azure public URL**:

- a) Enter the Azure-generated URL for the VHD. For guidance, select the link to the Microsoft document [Download a Windows VHD from Azure](#).
- b) Select a subscription. (A Linux image can be imported only if you select a customer-managed subscription.)
- c) Name the image.

- When you're done, select **Import Image**.

## Update a Quick Deploy catalog with a new image

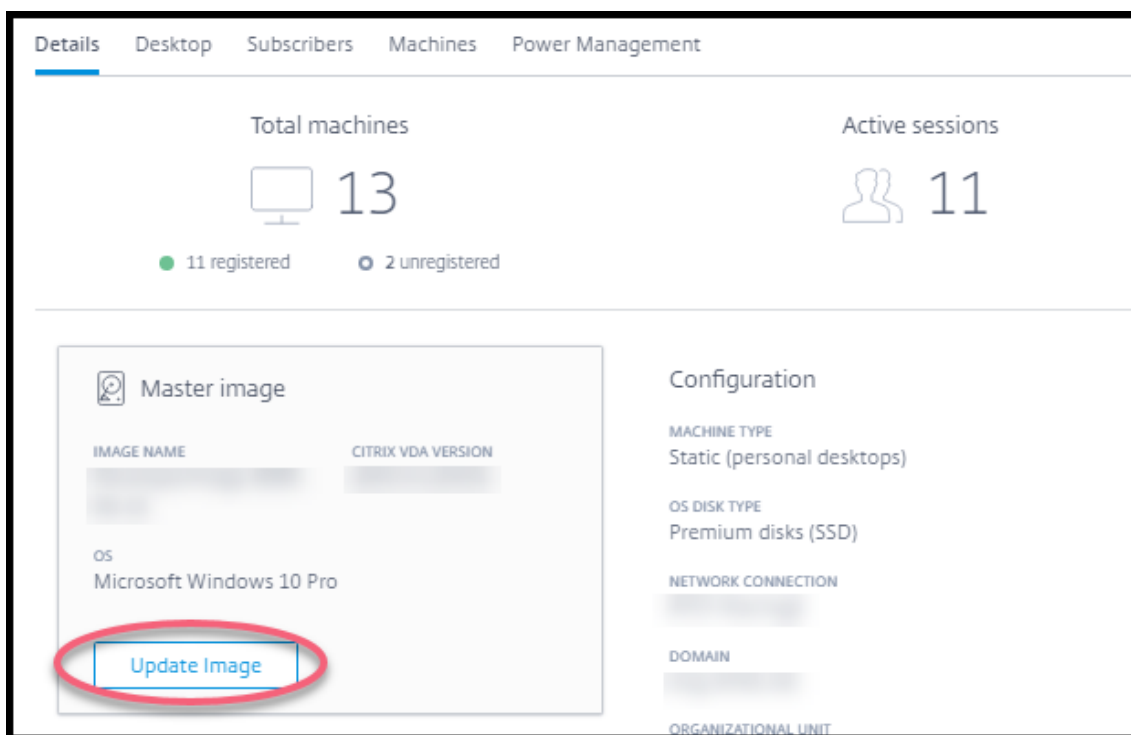
The catalog type determines which machines are updated when you update the catalog.

- For a random catalog, all the machines currently in the catalog are updated with the latest image. If you add more desktops to that catalog, they are based on the latest image.
- For a static catalog, the machines currently in the catalog are not updated with the latest image. Machines currently in the catalog continue to use the image they were created from. However, if you add more machines to that catalog, they are based on the latest image.

You can update a catalog containing machines with gen1 images with a gen2 image, if the catalog's machines support gen2. Similarly, you can update a catalog containing gen2 machines with a gen1 image, if the catalog's machines support gen1.

To update a catalog with a new image:

- From **Manage > Quick Deploy**, click anywhere in the catalog's entry.
- On the **Details** tab, select **Update Image**.



- Select an image.
- For random or multi-session catalogs: Select a logoff interval. After the service completes the initial image processing, subscribers receive a warning to save their work and log off from their

desktops. The logoff interval indicates how long subscribers have after receiving the message until the session ends automatically.

5. Select **Update Image**.

### Delete an image from Quick Deploy

1. From **Manage > Quick Deploy**, expand **Master Images** on the right.
2. Select the image you want to delete.
3. Select **Delete Image** at the bottom of the card. Confirm the deletion.

### Install a Windows VDA on an image

Use the following procedure when preparing a Windows image that you plan to import into the Citrix service.

For Linux VDA installation guidance, see the [Linux VDA product documentation](#).

1. In your Azure environment, connect to the image VM (if you're not already connected).
2. You can download a VDA by using the **Downloads** link on the Citrix Cloud navigation bar. Or, use a browser to navigate to the Citrix Virtual Apps and Desktops service [download](#) page.  
  
Download a VDA onto the VM. There are separate VDA download packages for a desktop (single-session) OS and a server (multi-session) OS.
3. Launch the VDA installer by double-clicking the downloaded file. The installation wizard launches.
4. On the **Environment** page, select the option to create an image using MCS, and then select **Next**.
5. On the **Core Components** page, select **Next**.
6. On the **Delivery Controller** page, select **Let Machine Creation Services do it automatically** and then select **Next**. Confirm your selection, if prompted.
7. Leave the default settings on the **Additional Components, Features**, and **Firewall** pages, unless Citrix instructs you otherwise. Select **Next** on each page.
8. On the **Summary** page, select **Install**. Prerequisites begin to install. When prompted to restart, agree.
9. The VDA installation resumes automatically. Prerequisite installation completes and then the components and features are installed. On the **Call Home** page, leave the default setting (unless Citrix instructs you otherwise). After you connect, select **Next**.
10. Select **Finish**. The machine restarts automatically.

11. To ensure that the configuration is correct, launch one or more of the applications you installed on the VM.
12. Shut down the VM. Do not Sysprep the image.

For more information about installing VDAs, see [Install VDAs](#).

## Network connections in Quick Deploy

July 21, 2021

### Introduction

This article provides details about how to create network connections to your corporate resources when using a Citrix Managed Azure subscription.

When using your own customer-managed Azure subscription, there is no need to create a network connection.

When creating a Quick Deploy catalog, you indicate if and how users access locations and resources on their corporate on-premises network from their Citrix desktops and apps. When using a connection, you must create the connection before creating the catalog.

When using a Citrix Managed Azure subscription, the choices are:

- No connectivity
- Azure VNet peering
- SD-WAN


You cannot change a catalog's connection type after the catalog is created.

### Requirements for all network connections

- When creating a connection, you must have valid [DNS server entries](#).
- When using Secure DNS or a third-party DNS provider, you must add the address range that is allocated for use by the service to the DNS provider's IP addresses on the allow list. That address range is specified when you create a connection.
- All service resources that use the connection (domain-joined machines) must be able to reach your Network Time Protocol (NTP) server, to ensure time synchronization.

## No connectivity

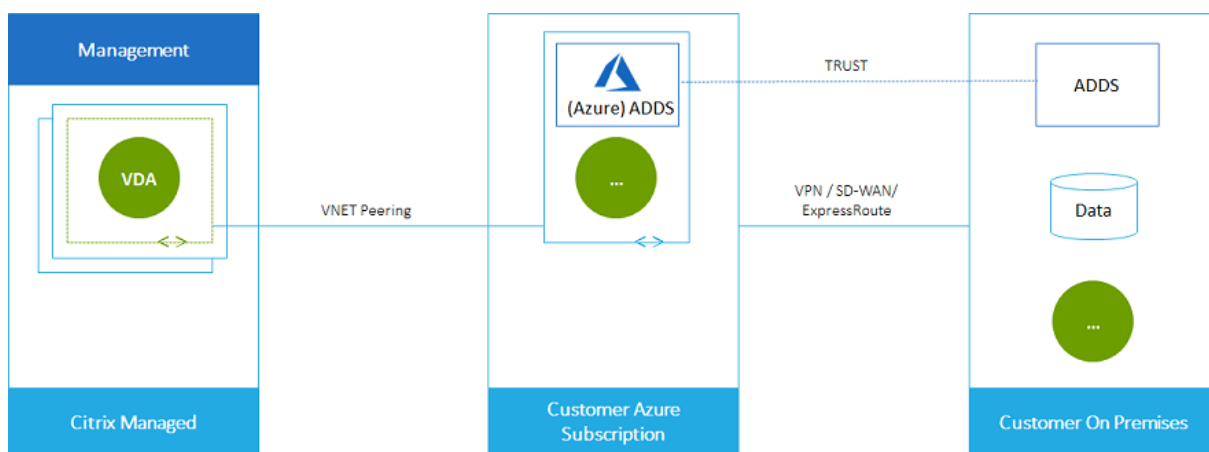
When a catalog is configured with **No connectivity**, users cannot access resources on their on-premises or other networks. This is the only choice when creating a catalog using quick create.




## About Azure VNet peering connections

Virtual network peering seamlessly connects two Azure virtual networks (VNETs): yours and the Citrix service VNet. Peering also helps enable users to access files and other items from your on-premises networks.

As shown in the following graphic, you create a connection using Azure VNet peering from the Citrix Managed Azure subscription to the VNet in your company's Azure subscription.



Here's another illustration of VNet peering.



Users can access their network resources (such as file servers) by joining the local domain when you create a catalog. (That is, you join the AD domain where file shares and other needed resources reside.) Your Azure subscription connects to those resources (in the graphics, using a VPN or Azure ExpressRoute). When creating the catalog, you provide the domain, OU, and account credentials.

**Important:**

- Learn about Azure VNet peering before using it in this service.
- Create a VNet peering connection before creating a catalog that uses it.

### Azure VNet peering custom routes

Custom, or user-defined, routes override Azure's default system routes for directing traffic between virtual machines in a VNet peering, on-premises networks, and the Internet. You might use custom routes if there are networks that Citrix service resources are expected to access but aren't directly connected through VNet peering. For example, you might create a custom route that forces traffic through a network appliance to the Internet or to an on-premises network subnet.

To use custom routes:

- You must have an existing Azure virtual network gateway or a network appliance such as Citrix SD-WAN in your Citrix service environment.
- When you add custom routes, you must update your company's route tables with the Citrix service's destination VNet information to ensure end-to-end connectivity.
- Custom routes are displayed in the Citrix service in the order in which they are entered. This display order does not affect the order in which Azure selects routes.

Before using custom routes, review the Microsoft article [Virtual network traffic routing](#) to learn about using custom routes, next hop types, and how Azure selects routes for outbound traffic.

You can add custom routes when you create an Azure VNet peering connection or to existing ones in your Citrix service environment. When you're ready to use custom routes with your VNet peering, refer to the following sections in this article:

- For custom routes with new Azure VNet peerings: [Create an Azure VNet peering connection](#)
- For custom routes with existing Azure VNet peerings: [Manage custom routes for existing Azure VNet peer connections](#)

### Azure VNet peering requirements and preparation

- Credentials for an Azure subscription owner. This must be an Azure Active Directory account. This service does not support other account types, such as live.com or external Azure AD accounts (in a different tenant).
- An Azure subscription, resource group, and virtual network (VNet).
- Set up the Azure network routes so that VDAs in the Citrix Managed Azure subscription can communicate with your network locations.
- Open Azure network security groups from your VNet to the specified IP range.
- **Active Directory:** For domain-joined scenarios, we recommend that you have some form of Active Directory services running in the peered VNet. This takes advantage of the low latency characteristics of the Azure VNet peering technology.

For example, the configuration might include Azure Active Directory Domain Services (AADDs), a domain controller VM in the VNet, or Azure AD Connect to your on-premises Active Directory.

After you enable AADDs, you cannot move your managed domain to a different VNet without deleting the managed domain. So, it's important to select the correct VNet to enable your managed domain. Before proceeding, review the Microsoft article [Networking considerations for Azure AD Domain Services](#).

- **VNet IP range:** When creating the connection, you must provide an available CIDR address space (IP address and network prefix) that is unique among the network resources and the Azure VNets being connected. This is the IP range assigned to the VMs within the Citrix service's peered VNet.

Ensure that you specify an IP range that does not overlap any addresses that you use in your Azure and on-premises networks.

- For example if your Azure VNet has an address space of 10.0.0.0 /16, create the VNet peering connection in the Citrix service as something such as 192.168.0.0 /24.
- In this example, creating a peering connection with a 10.0.0.0 /24 IP range would be considered an overlapping address range.

If addresses overlap, the VNet peering connection might not be created successfully. It also does not work correctly for site administration tasks.

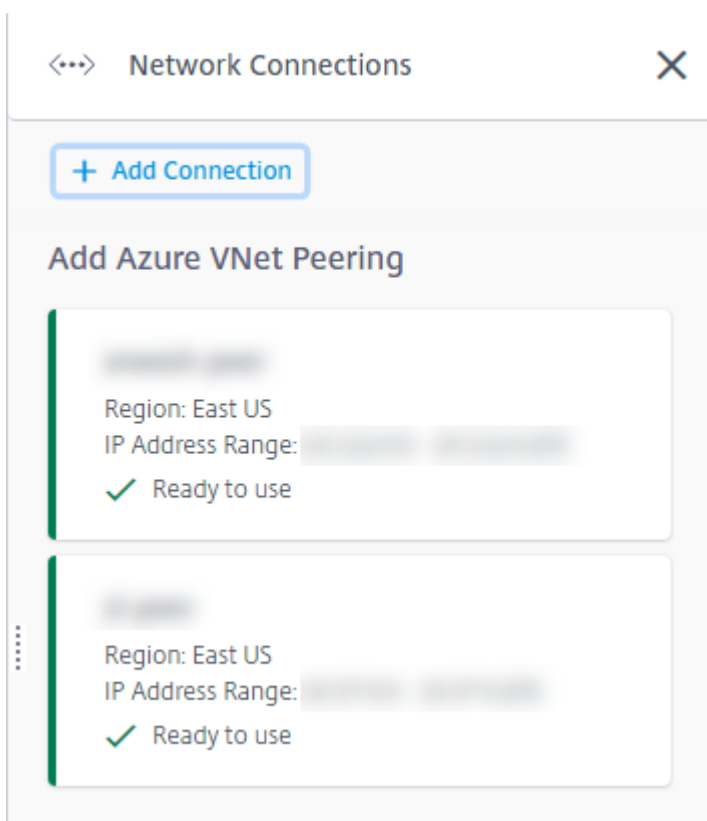
To learn about VNet peering, see the following Microsoft articles.

- [Virtual network peering](#)
- [Azure VPN Gateway](#)
- [Create a Site-to-Site connection in the Azure portal](#)
- [VPN Gateway FAQ](#) (search for "overlap")

### **Create an Azure VNet peering connection**

1. From **Manage > Quick Deploy**, expand **Network Connections** on the right. If you have already set up connections, they're listed.





2. Select **Add Connection**.
3. Click anywhere in the **Add Azure VNet Peering** box.


## Add a network connection

Choose how you want to connect to your local network:

**Add Azure VNet Peering**  
Easy setup for Azure customers – Seamlessly connect your Azure virtual network.

4. Select **Authenticate Azure Account**.

## Add Azure VNet Peering




### What's ahead

Virtual network peering seamlessly connects two Azure virtual networks (VNETs): yours and the Citrix Managed Desktops VNet. Peering also helps enable users to access files and other items from your on-premises networks.

You will need the following:

1. An Azure subscription, resource group, and virtual network (VNet).
2. Credentials for an Azure Resource Manager subscription owner.
3. An available IP address and network prefix (in CIDR format) that is unique among the network resources and the Azure VNETs being connected.
4. For domain-joined scenarios, we recommend that you have some form of Active Directory services running in the peered VNet.



5. The service automatically takes you to the Azure sign-in page to authenticate your Azure subscriptions. After you sign in to Azure (with the global administrator account credentials) and accept the terms, you are returned to the connection creation details dialog.

## Add Azure VNet Peering

Azure VNet peering name

VNet details to peer

Select Azure Subscription

Select Resource Group

Select VNet to Peer

✓ This VNet is in the West US region, which is supported

Is this VNet using an Azure Virtual Network Gateway?

No  Yes

IP address and network prefix to be used by VNet peering ?

⚠ The IP addresses cannot conflict with any existing IP addresses in your network.

/

?

✓ 10.2.0.0 - 10.2.0.255 (251 addresses available for machines)

Do you want to add routes? ?

No  Yes


6. Type a name for the Azure VNet peer.
7. Select the Azure subscription, resource group, and the VNet to peer.
8. Indicate whether the selected VNet uses an Azure Virtual Network Gateway. For information, see the Microsoft article [Azure VPN Gateway](#).
9. If you answered **Yes** in the previous step (the VNet uses an Azure virtual network gateway), indicate whether you want to enable virtual network gateway route propagation. When enabled, Azure automatically learns (adds) all routes through the gateway.

You can change this setting later on the connection's **Details** page. However, changing it can cause route pattern changes and VDA traffic interruptions. Also, if you disable it later, you must manually add routes to the networks that VDAs will use.


10. Type an IP address and select a network mask. The address range to be used is displayed, plus how many addresses that the range supports. Ensure that the IP range does not overlap any addresses that you use in your Azure and on-premises networks.
  - For example, if your Azure VNet has an address space of 10.0.0.0 /16, create the VNet peering connection in the Citrix service as something such as 192.168.0.0 /24.
  - In this example, creating a VNet peering connection with a 10.0.0.0 /24 IP range is considered an overlapping address range.

If addresses overlap, the VNet peering connection might not be created successfully. It also won't work correctly for site administration tasks.


11. Indicate whether you want to add custom routes to the VNet peering connection. If you select **Yes**, enter the following information:
  - a) Type a friendly name for the custom route.
  - b) Enter the destination IP address and network prefix. The network prefix must be between 16 and 24.
  - c) Select a next hop type for where you want traffic to be routed. If you select **Virtual appliance**, enter the internal IP address of the appliance.


Do you want to add routes? 


No  Yes


 Make sure your company's route tables are updated with the Citrix Managed Desktops VNet information to ensure end-to-end connectivity: 10.2.0.0/24 (provided above). Added routes override Azure default routing. Routes apply to all connections from machines using this VNet peering.


Route name


Destination IP address and network prefix 

 /  

 10.2.0.0 - 10.2.0.255

Next hop type 

Next hop address 

[+ Add route](#)

For more information about next hop types, see the [Custom routes](#) section in the Microsoft article *Virtual network traffic routing*.

d) To create another custom route for the connection, select **Add route**.

12. Select **Add VNet Peering**.

After the connection is created, it is listed under **Network Connections > Azure VNet Peers** on the right side of the **Manage > Quick Deploy** dashboard. When you create a catalog, this connection is included in the available network connections list.



## View Azure VNet peering connection details

XXXXXXXX-XXXX

Details Routes

Not in use



Catalogs

0

Machines

0

Images

0

Bastions

0

### Region

VNet 1  
East US

VNet 2 - CITRIX MANAGED  
East US

### Allocated Network Space

IP ADDRESS RANGE

IP ADDRESS AVAILABLE FOR MACHINES

DNS SERVERS

### Peered Virtual Network Details

VIRTUAL NETWORK

SUBSCRIPTION ID

RESOURCE GROUP

AZURE VIRTUAL NETWORK GATEWAY

Disabled

Delete Connection

1. From **Manage > Quick Deploy**, expand **Network Connections** on the right.
2. Select the Azure VNet peering connection you want to display.

Details include:

- The number of catalogs, machines, images, and bastions that use this connection.
- The region, allocated network space, and peered VNets.
- The routes currently configured for the VNet peering connection.

### **Manage custom routes for existing Azure VNet peer connections**

You can add new custom routes to an existing connection or modify existing custom routes, including disabling or deleting custom routes.

#### **Important:**

Modifying, disabling, or deleting custom routes changes the traffic flow of the connection and might disrupt any user sessions that might be active.

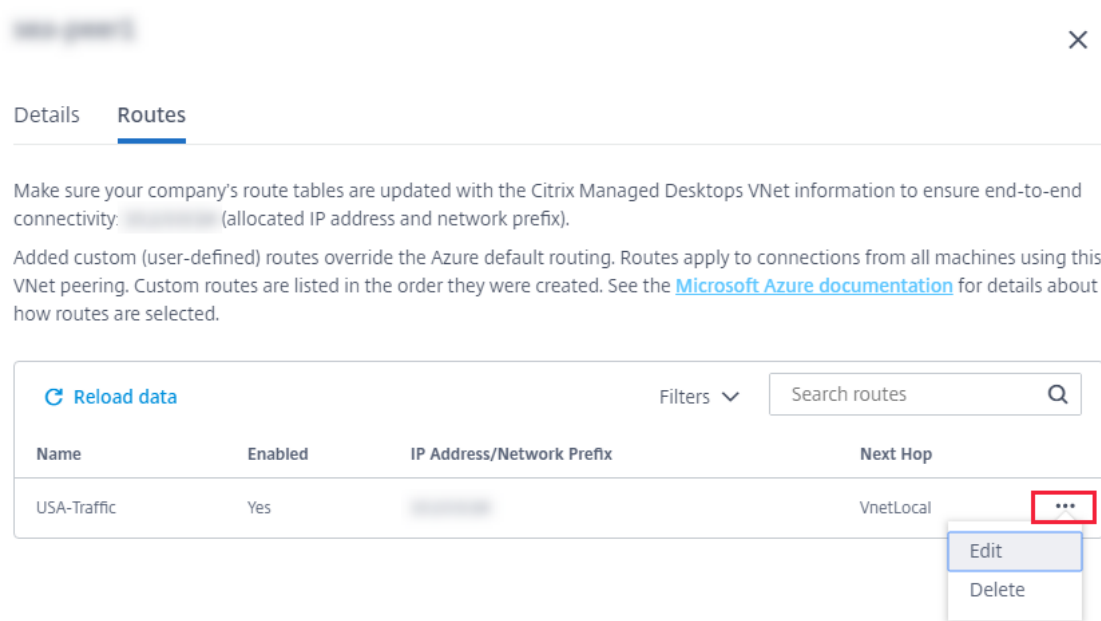
To add a custom route:

1. From **Manage > Quick Deploy**, expand **Network Connections** on the right.
2. Select the connection you want to delete.
3. From the connection details, select **Routes** and then select **Add Route**.
4. Enter a friendly name, the destination IP address and prefix, and the next hop type you want to use. If you select **Virtual Appliance** as the next hop type, enter the internal IP address of the appliance.
5. Indicate whether you want to enable the custom route. By default, the custom route is enabled.
6. Select **Add Route**.

To modify or disable a custom route:

1. From **Manage > Quick Deploy**, expand **Network Connections** on the right.
2. Select the connection you want to delete.
3. From the connection details, select **Routes** and then locate the custom route you want to manage.
4. From the ellipsis menu, select **Edit**.





Details **Routes**

Make sure your company's route tables are updated with the Citrix Managed Desktops VNet information to ensure end-to-end connectivity: [redacted] (allocated IP address and network prefix).

Added custom (user-defined) routes override the Azure default routing. Routes apply to connections from all machines using this VNet peering. Custom routes are listed in the order they were created. See the [Microsoft Azure documentation](#) for details about how routes are selected.

Name	Enabled	IP Address/Network Prefix	Next Hop
USA-Traffic	Yes	[redacted]	VnetLocal

5. Make any needed changes to the destination IP address and prefix or the next hop type, as needed.
6. To enable or disable a custom route, in **Enable this route?**, select **Yes** or **No**.
7. Select **Save**.

To delete a custom route:

1. From **Manage > Quick Deploy**, expand **Network Connections** on the right.
2. Select the connection you want to delete.
3. From the connection details, select **Routes** and then locate the custom route you want to manage.
4. From the ellipsis menu, select **Delete**.
5. Select **Deleting a route may disrupt active sessions** to acknowledge the impact of deleting the custom route.
6. Select **Delete Route**.

### Delete an Azure VNet peering connection

Before you can delete an Azure VNet peering connection, remove any catalogs associated with it. See [Delete a catalog](#).

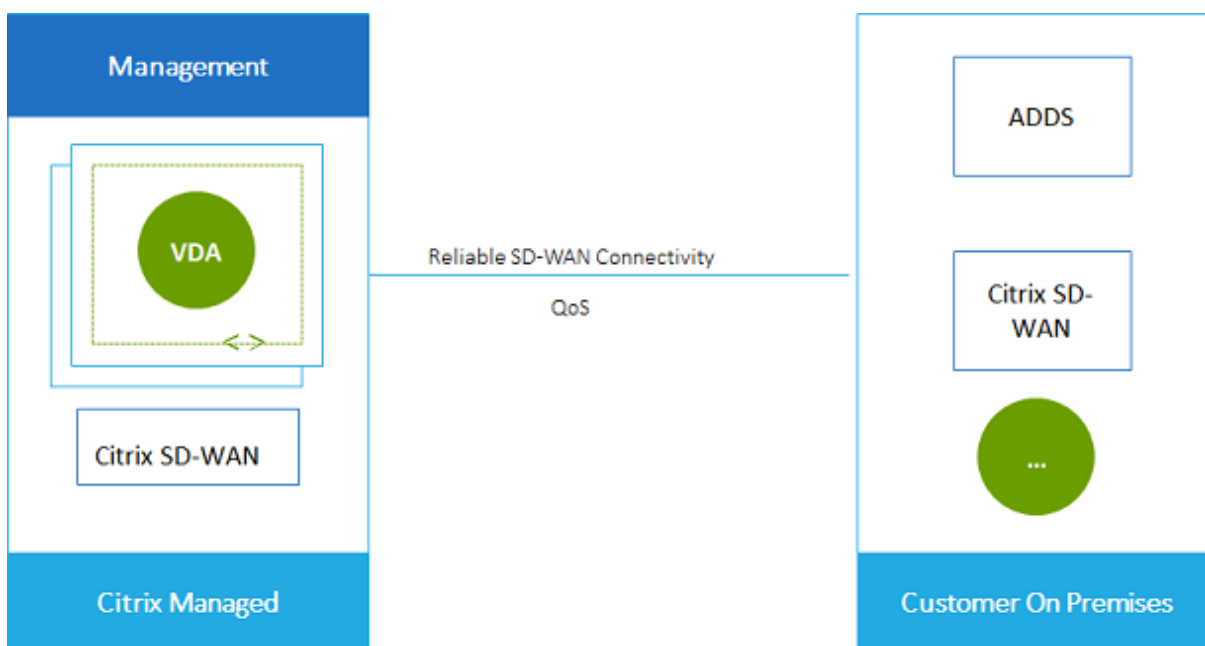
1. From **Manage > Quick Deploy**, expand **Network Connections** on the right.
2. Select the connection you want to delete.
3. From the connection details, select **Delete Connection**.

## About SD-WAN connections

Citrix SD-WAN optimizes all the network connections needed by the Citrix service. Working in concert with the HDX technologies, Citrix SD-WAN provides quality-of-service and connection reliability for ICA and out-of-band Citrix service traffic. Citrix SD-WAN supports the following network connections:

- Multi-stream ICA connection between users and their virtual desktops
- Internet access from the virtual desktop to websites, SaaS apps, and other cloud properties
- Access from the virtual desktop back to on-premises resources such as Active Directory, file servers, and database servers
- Real-time/interactive traffic carried over RTP from the media engine in the Workspace app to cloud-hosted Unified Communications services such as Microsoft Teams
- Client-side fetching of videos from sites like YouTube and Vimeo

As shown in the following graphic, you create an SD-WAN connection from the Citrix Managed Azure subscription to your sites. During connection creation, SD-WAN VPX appliances are created in the Citrix Managed Azure subscription. From the SD-WAN perspective, that location is treated as a branch.



## SD-WAN connection requirements and preparation

- If the following requirements are not met, the SD-WAN network connection option is not available.
  - Citrix Cloud service entitlements: Citrix Virtual Apps and Desktops service and SD-WAN Orchestrator.
  - An installed and configured SD-WAN deployment. The deployment must include a Master Control Node (MCN), whether in the cloud or on-premises, and be managed with SD-WAN

Orchestrator.

- VNet IP range: Provide an available CIDR address space (IP address and network prefix) that is unique among the network resources being connected. This is the IP range assigned to the VMs within the Citrix service's VNet.

Ensure that you specify an IP range that does not overlap any addresses that you use in your cloud and on-premises networks.

- For example, if your network has an address space of 10.0.0.0 /16, create the connection in the Citrix service as something such as 192.168.0.0 /24.
- In this example, creating a connection with a 10.0.0.0 /24 IP range would be considered an overlapping address range.

If addresses overlap, the connection might not be created successfully. It also does not work correctly for site administration tasks.

- The connection configuration process includes tasks that you (the service administrator) and the SD-WAN Orchestrator administrator must complete. Also, to complete your tasks, you need information provided by the SD-WAN Orchestrator administrator.

We recommend that you both review the guidance in this document, plus the SD-WAN documentation, before actually creating a connection.

## Create an SD-WAN connection

### Important:

For details about SD-WAN configuration, see [SD-WAN configuration for Citrix Virtual Apps and Desktops integration](#).

1. From **Manage > Quick Deploy**, expand **Network Connections** on the right.
2. Select **Add Connection**.
3. On the **Add a network connection** page, click anywhere in the SD-WAN box.
4. The next page summarizes what's ahead. When you're done reading, select **Start Configuring SD-WAN**.
5. On the **Configure SD-WAN** page, enter the information provided by your SD-WAN Orchestrator administrator.
  - **Deployment mode:** If you select **High availability**, two VPX appliances are created (recommended for production environments). If you select **Standalone**, one appliance is created. You cannot change this setting later. To change to the deployment mode, you'll have to delete and re-create the branch and all associated catalogs.
  - **Name:** Type a name for the SD-WAN site.

- **Throughput and number of offices:** This information is provided by your SD-WAN Orchestrator administrator.
  - **Region:** The region where the VPX appliances will be created.
  - **VDA subnet and SD-WAN subnet:** This information is provided by your SD-WAN Orchestrator administrator. See SD-WAN connection requirements and preparation for information about avoiding conflicts.
6. When you're done, select **Create Branch**.
  7. The next page summarizes what to look for on the **Manage > Quick Deploy** dashboard. When you're done reading, select **Got it**.
  8. From **Manage > Quick Deploy**, the new SD-WAN entry under **Network Connections** shows the progress of the configuration process. When the entry turns orange with the message *Awaiting activation by SD-WAN administrator*, notify your SD-WAN Orchestrator administrator.
  9. For SD-WAN Orchestrator administrator tasks, see the SD-WAN Orchestrator [product documentation](#).
  10. When the SD-WAN Orchestrator administrator finishes, the SD-WAN entry under **Network Connections** turns green, with the message *You can create catalogs using this connection*.

### View SD-WAN connection details

1. From **Manage > Quick Deploy**, expand **Network Connections** on the right.
2. Select **SD-WAN** if it's not the only selection.
3. Select the connection you want to display.

The display includes:

- **Details tab:** Information you specified when configuring the connection.
- **Branch Connectivity tab:** Name, cloud connectivity, availability, bandwidth tier, role, and location for each branch and MCN.

### Delete an SD-WAN connection

Before you can delete an SD-WAN connection, remove any catalogs associated with it. See [Delete a catalog](#).

1. From **Manage > Quick Deploy**, expand **Network Connections** on the right.
2. Select SD-WAN if it's not the only selection.
3. Select the connection you want to delete, to expand its details.
4. On the **Details** tab, select **Delete Connection**.
5. Confirm the deletion.

## Users and authentication in Quick Deploy

July 21, 2021

### User authentication methods

Users must authenticate when they log in to Citrix Workspace to start their desktop or apps.

Quick Deploy supports the following user authentication methods:

- **Managed Azure AD:** Managed Azure AD is an Azure Active Directory (AAD) provided and managed by Citrix. You don't need to provide your own Active Directory structure. Just add your users to the directory.
- **Your identity provider:** You can use any available authentication method in Citrix Cloud.

#### Note:

- Remote PC Access deployments use only Active Directory. For details, see [Remote PC Access](#).
- If you use Azure AD Domain Services: Workspace logon UPNs must contain the domain name that was specified when enabling Azure AD Domain Services. Logons cannot use UPNs for a custom domain you create, even if that custom domain is designated as primary.

Setting up user authentication includes the following procedures:

1. Configure the user authentication method in Citrix Cloud and Workspace Configuration.
2. If you're using Managed Azure AD for user authentication, add users to the directory.
3. Add users to a catalog.

### Configure user authentication in Citrix Cloud

To configure user authentication in Citrix Cloud:

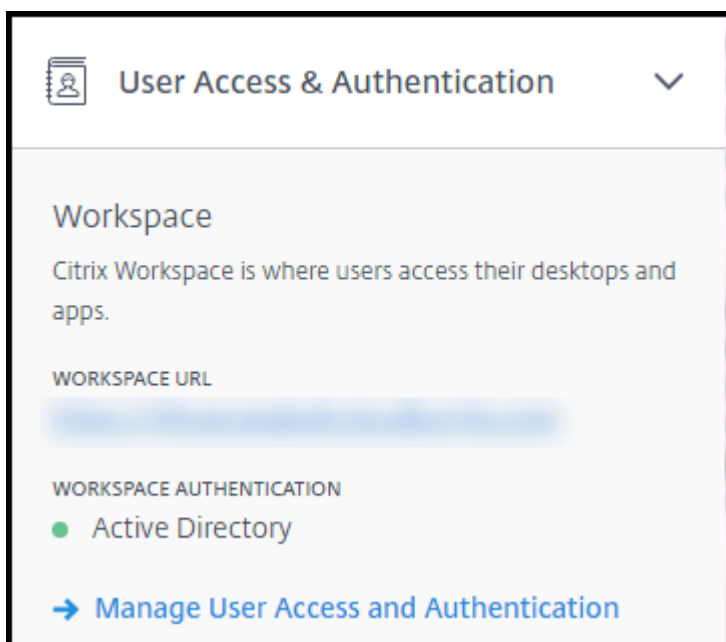
- Connect to the user authentication method you want to use. (In Citrix Cloud, you "connect" or "disconnect" from an authentication method.)
- In Citrix Cloud, set Workspace authentication to use the connected method.

#### Note:

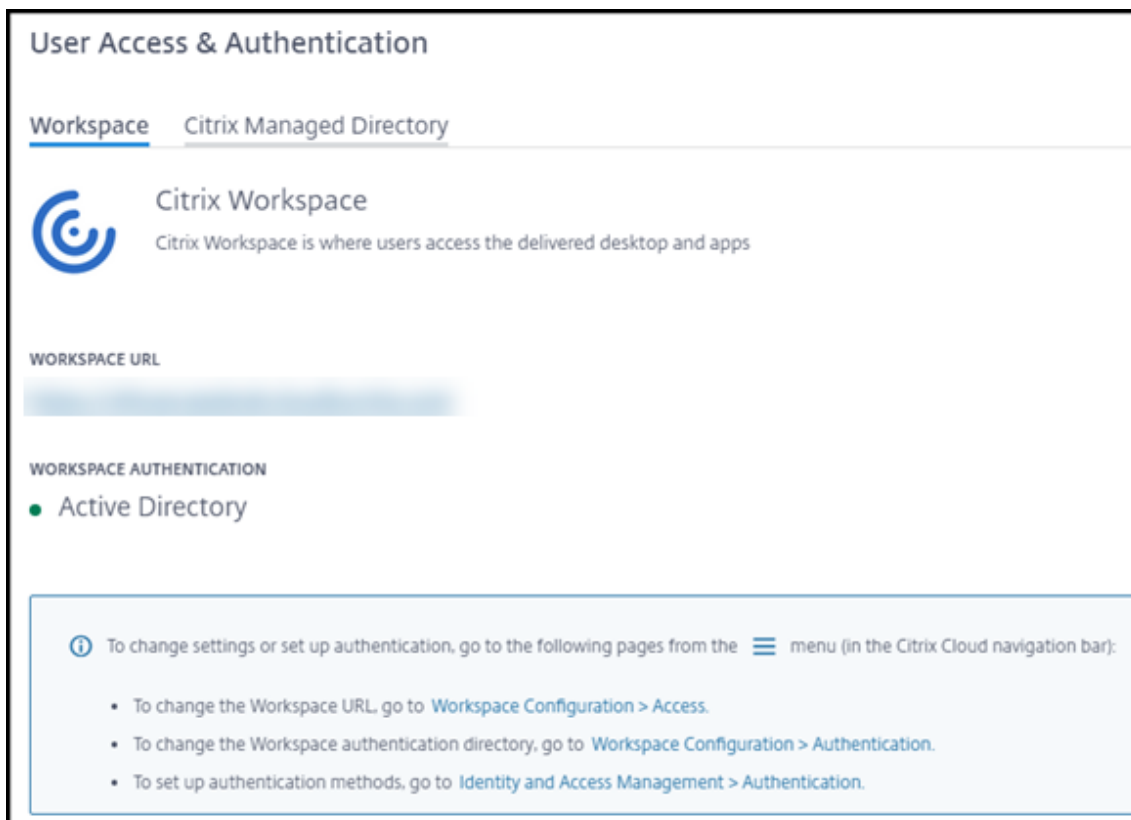
The Managed Azure AD authentication method is configured by default. That is, it is automatically connected in Citrix Cloud, and Workspace authentication is automatically set to use Managed Azure AD for this service. If you want to use this method (and have not previously configured a different method), continue with Add and delete users in Managed Azure AD.

To change the authentication method:

1. From **Manage > Quick Deploy**, select **User Access & Authentication** on the right.



2. Select **Manage User Access and Authentication**. Select the **Workspace** tab, if it isn't already selected. (The other tab indicates which user authentication method is currently configured.)



3. Follow the link **To set up authentication methods**. That link takes you to Citrix Cloud. Select **Connect** in the ellipsis menu for the method you want.
4. While still in Citrix Cloud, select **Workspace Configuration** in the upper left menu. On the **Authentication** tab, select the method you want.

What to do next:

- If you're using Managed Azure AD, add users to the directory.
- For all authentication methods, add users to the catalog.

### **Add and delete users in Managed Azure AD**

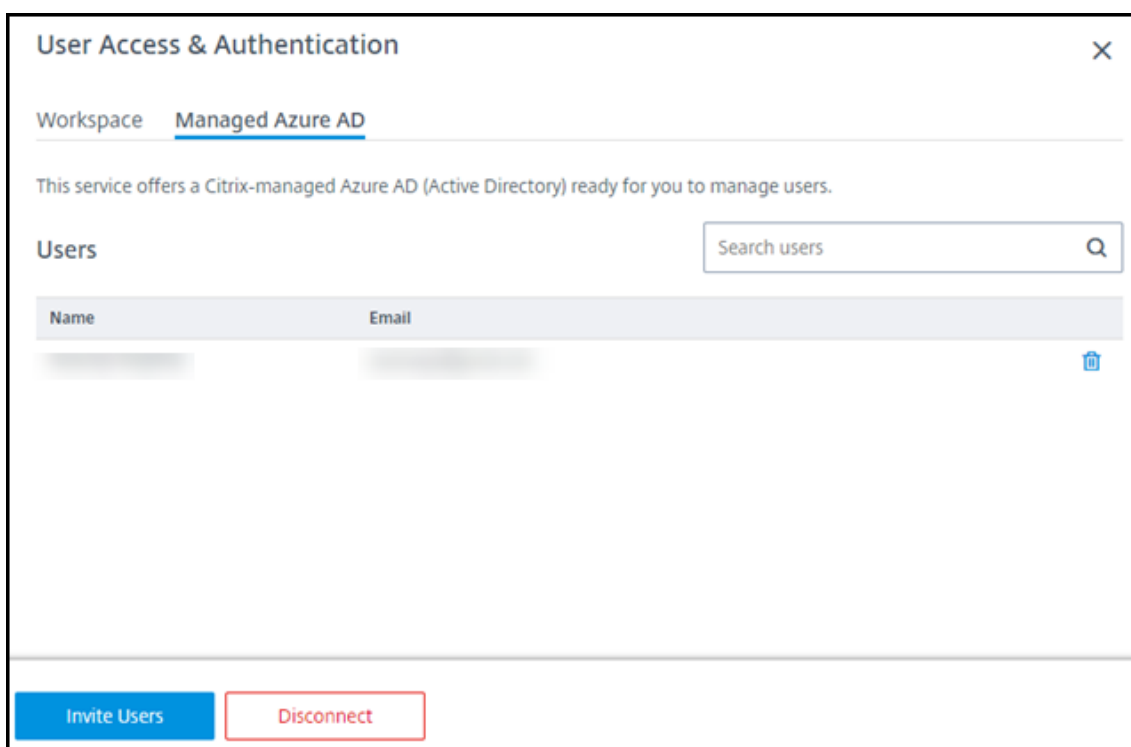
Complete this procedure only if you're using Managed Azure AD for user authentication to Citrix Workspace.

You provide your users' name and email addresses. Citrix then emails an invitation to each of them. The email instructs users to select a link that joins them to the Citrix Managed Azure AD.

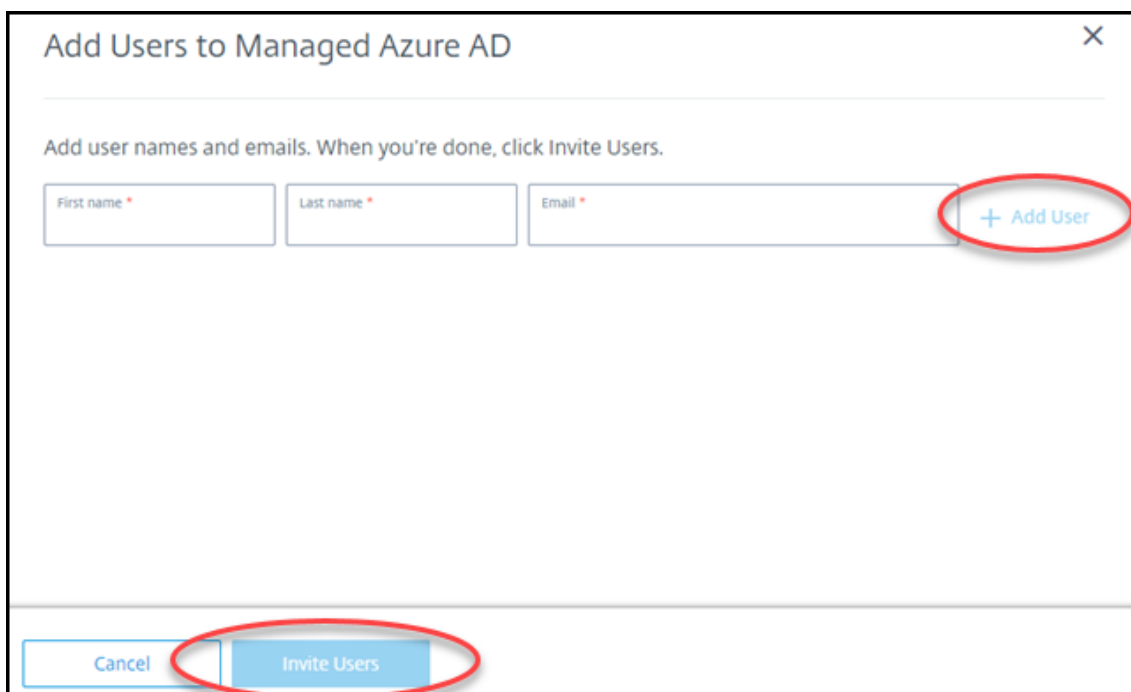
- If the user already has a Microsoft account with the email address you provided, that account is used.
- If the user does not have a Microsoft account with the email address, Microsoft creates an account.

To add and invite users to Managed Azure AD:

1. From **Manage > Quick Deploy**, expand **User Access & Authentication** on the right. Select **Manage User Access and Authentication**.
2. Select the **Managed Azure AD** tab.
3. Select **Invite Users**.



4. Type the name and email address of a user, and then select **Add User**.



5. Repeat the preceding step to add other users.
6. When you're done adding user information, select **Invite Users** at the bottom of the card.

To delete a user from Managed Azure AD, select the trash icon next to the name of the user you want



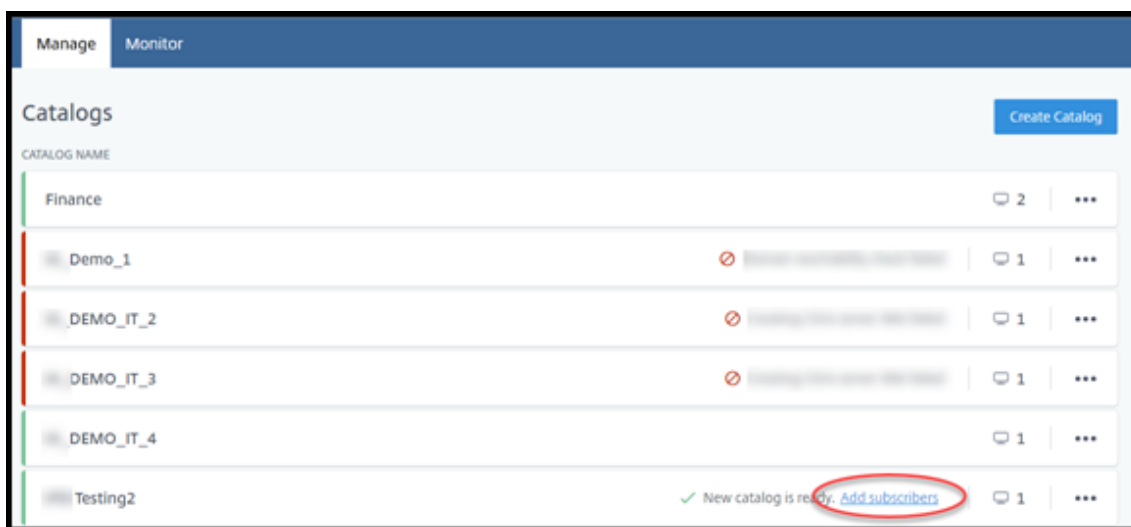
to delete from the directory. Confirm the deletion.

What to do next: Add users to the catalog

### Add or remove users in a catalog

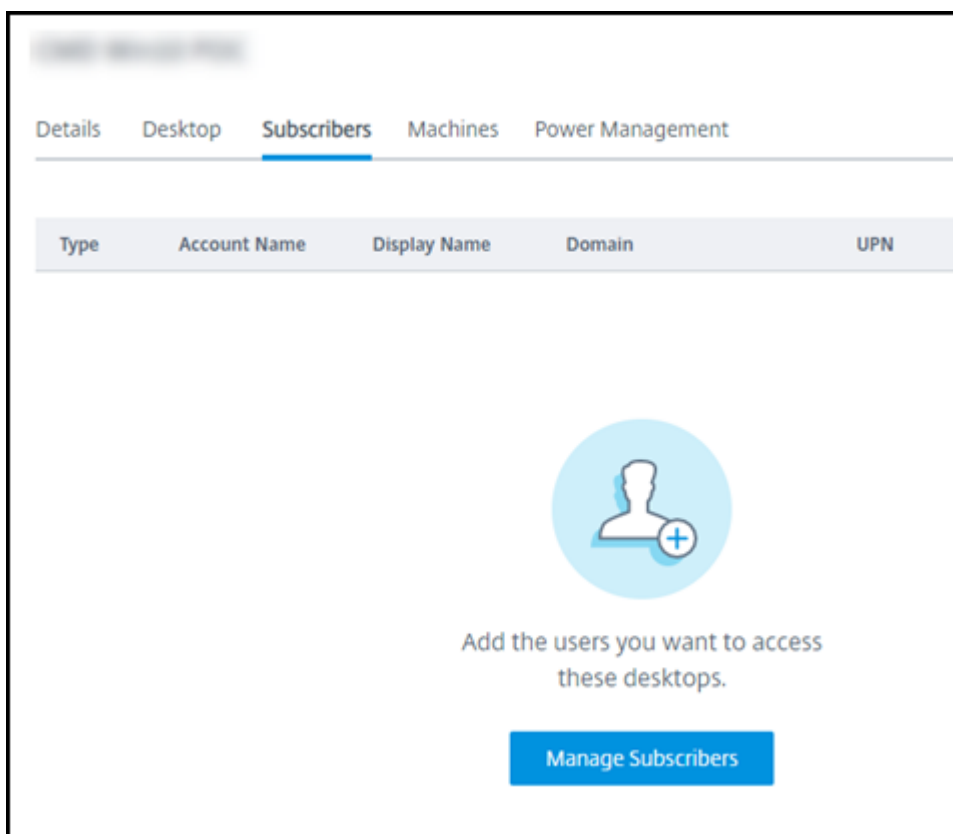
Complete this procedure regardless of which authentication method you use.

1. From **Manage > Quick Deploy**, if you haven't added any users to a catalog, select **Add subscribers**.

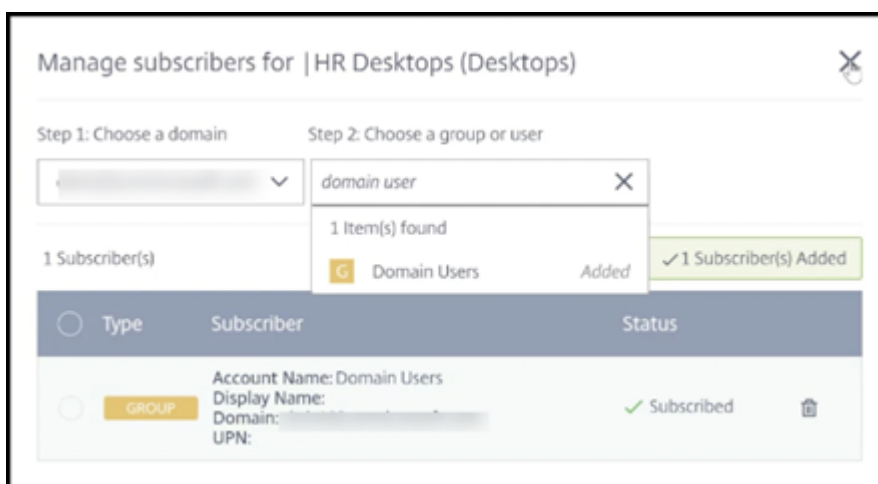


To add users to a catalog that already has users, click anywhere in the catalog's entry.

2. On the **Subscribers** tab, select **Manage Subscribers**.



3. Select a domain. (If you're using Managed Azure AD for user authentication, there's only one entry in the domain field.) Then select a user.



4. Select other users, as needed. When you're done, select the **X** in the upper right corner.

To remove users from a catalog, follow steps 1 and 2. In step 3, select the trash icon next to the name you want to delete (instead of selecting a domain and group/user). This action removes the user from the catalog, not from the source (such as Managed Azure AD or your own AD or AAD).

What to do next:

- For a catalog with multi-session machines, [add applications](#), if you haven't already.
- For all catalogs, [send the Citrix Workspace URL to your users](#).

## More information

For more information about authentication in Citrix Cloud, see [Identity and access management](#).

## Remote PC Access in Quick Deploy

July 22, 2021

### Introduction

Citrix Remote PC Access enables users to remotely use physical Windows or Linux machines located in the office. Users receive the best user experience by using Citrix HDX to deliver their office PC session.

Remote PC Access supports domain-joined machines.

This article describes how to create a Remote PC Access deployment using the Quick Deploy interface. To create a Remote PC Access deployment using the Full Configuration interface, see [Remote PC Access](#).

### Differences from delivering virtual desktops and apps

If you're familiar with delivering virtual desktops and apps, the Remote PC Access feature has several differences:

- A Remote PC Access catalog usually contains existing physical machines. So, you don't have to prepare an image or provision machines to use Remote PC Access. Delivering desktops and apps usually uses virtual machines (VMs), and an image is used as a template to provision the VMs.
- When a machine in a Remote PC Access random pooled catalog is powered off, it is not reset to the original state of the image.
- For Remote PC Access static user assignment catalogs, the assignment occurs after a user logs in (either at the machine or via RDP). When delivering desktops and apps, a user is assigned if a machine is available.

### Installation and configuration summary

Review this section before starting the tasks.

1. Before you start:
  - a) Review the requirements and considerations.
  - b) Complete the preparation tasks.
2. From Citrix Cloud:
  - a) [Set up a Citrix Cloud account and subscribe to the Citrix Virtual Apps and Desktops service.](#)
  - b) Set up a resource location that can access your Active Directory resources. Install at least two Cloud Connectors in the resource location. The Cloud Connectors communicate with Citrix Cloud.  
  
Follow the guidance for [creating a resource location and installing Cloud Connectors in it](#). This information includes system requirements, preparation, and procedures.
  - c) [Connect your Active Directory to Citrix Cloud.](#)
3. Install a Citrix Virtual Delivery Agent (VDA) on each machine that users will access remotely. VDAs communicate with Citrix Cloud through the Cloud Connectors in the resource location.
4. From **Manage > Quick Deploy**:
  - a) Create a Remote PC Access catalog. In this procedure, you specify the location of your resource location and select the user assignment method.
  - b) [Add subscribers \(users\) to the catalog](#), if needed. Add users to a catalog if the catalog uses either the static autoassigned or random pooled user assignment method. You do not need to add users to a static preassigned catalog.
5. [Send the workspace URL to users](#). From their workspace, users can log on to their machines in the office.

## Requirements and considerations

References to machines in this section refer to the machines that users access remotely.

### General

- The machines must be running a single-session Windows 10 or Linux (Red Hat Enterprise Linux and Ubuntu) operating system.
- The machine must be joined to an Active Directory Domain Services domain.
- If you are familiar with using Remote PC Access with Citrix Virtual Apps and Desktops, the Wake-on-LAN feature is not available in this service.

## Network

- The machine must have an active network connection. A wired connection is preferred for greater reliability and bandwidth.
- If using Wi-Fi:
  - Set the power settings to leave the wireless adapter turned on.
  - Configure the wireless adapter and network profile to allow automatic connection to the wireless network before the user logs on. Otherwise, the VDA does not register until the user logs on. The machine isn't available for remote access until a user logs on.
  - Ensure that the Cloud Connectors can be reached from the Wi-Fi network.

## Devices and peripherals

- The following devices are not supported:
  - KVM switches or other components that can disconnect a session.
  - Hybrid PCs, including All-in-One and NVIDIA Optimus laptops and PCs.
- Connect the keyboard and mouse directly to the machine. Connecting to the monitor or other components that can be turned off or disconnected, can make these peripherals unavailable. If you must connect the input devices to components such as monitors, do not turn those components off.
- For laptop and Surface Pro devices: Ensure that the laptop is connected to a power source instead of running on the battery. Configure the laptop power options to match the options of a desktop machine. For example:
  - Disable the hibernate feature.
  - Disable the sleep feature.
  - Set the close lid action to **Do Nothing**.
  - Set the **press the power button action** to **Shut Down**.
  - Disable video card and NIC energy-saving features.

When using a docking station, you can undock and redock laptops. When you undock the laptop, the VDA reregisters with the Cloud Connectors over Wi-Fi. However, when you redock the laptop, the VDA doesn't switch to use the wired connection unless you disconnect the wireless adapter. Some devices provide built-in functionality to disconnect the wireless adapter upon establishing a wired connection. Other devices require custom solutions or third-party utilities to disconnect the wireless adapter. Review the Wi-Fi considerations mentioned previously.

To enable docking and undocking for Remote PC Access devices:

- In **Start > Settings > System > Power & Sleep**, set **Sleep** to **Never**.
- In **Device Manager > Network adapters > Ethernet adapter**, go to **Power Management** and clear **Allow the computer to turn off this device to save power**. Ensure that **Allow**

**this device to wake the computer** is selected.

## Linux VDA

- Use the Linux VDA on physical machines only in non-3D mode. Due to limitations on NVIDIA's driver, the local screen of the PC cannot be blacked out and displays the activities of the session when HDX 3D mode is enabled. Showing this screen is a security risk.
- Catalogs with Linux machines must use the static preassigned user assignment method. Catalogs with Linux machines cannot use either the static autoassigned or random pooled assignment methods.

## Workspace considerations

- Multiple users with access to the same office PC see the same icon in Citrix Workspace. When a user signs in to Citrix Workspace, that machine appears as unavailable if it is already in use by another user.

## Prepare

- Decide how to install the VDA on the machines. Several methods are available:
  - Manually install the VDA on each machine.
  - Push the VDA installation using Group Policy, [using a script](#).
  - Push the VDA installation using an Electronic Software Distribution (ESD) tool such as Microsoft System Center Configuration Manager (SCCM). For details, see [Install VDAs using SCCM](#).
- Learn about user assignment methods and decide which method you'll use. You specify the method when creating a Remote PC Access catalog.
- Decide how the machines (actually the VDAs you install on the machines) will register with Citrix Cloud. A VDA must register to establish communications with the session broker in Citrix Cloud.

VDAs register through the Cloud Connectors in their resource location. You can specify Cloud Connector addresses when you install a VDA, or later.

For a VDA's first (initial) registration, Citrix recommends using policy-based GPO or LGPO. After the initial registration, Citrix recommends using auto-update, which is enabled by default. [Learn more about VDA registration](#).

## Install a VDA

Download and install a VDA on each physical machine that users will access remotely.

## Download a VDA

- To download a Windows VDA:
  1. Using your Citrix Cloud account credentials, browse to the [Citrix Virtual Apps and Desktops service download page](#).
  2. Download the latest VDA. Two types of installation packages are available. The year and month values in the VDA title vary.
- To download a Linux VDA for Remote PC Access, follow the guidance in the [Linux VDA documentation](#).

## Windows VDA installation package types

The Citrix download site provides two Windows VDA installation package types that can be used for Remote PC Access machines:

- Single-session core VDA installer (*release is yymm*): `VDAWorkstationCoreSetup_release.exe`

The single-session core VDA installer is tailored specifically for Remote PC Access. It's lightweight and easier to deploy (than other VDA installers) over the network to all machines. It does not include components that typically aren't needed in these deployments, such as Citrix Profile Management, Machine Identity Service, and the user personalization layer.

However, without Citrix Profile Management installed, the displays for Citrix Analytics for Performance and some Monitor details aren't available. For details about those limitations, see the blog post [Monitor and troubleshoot Remote PC Access machines](#).

If you want full analytics and monitoring displays, use the single-session full VDA installer.

- Single-session full VDA installer (*release is yymm*): `VDAWorkstationSetup_release.exe`

Although the single-session full VDA installer is a larger package than the single-session core VDA installer, you can tailor it to install only the components you need. For example, you can install the components that support Profile Management.

## Install a Windows VDA for Remote PC Access interactively

1. Double-click the VDA installation file that you downloaded.
2. On the **Environment** page, select **Enable Remote PC Access**, and then click **Next**.
3. On the **Delivery Controller** page, select one of the following:
  - If you know the addresses of your Cloud Connectors, select **Do it manually**. Enter the FQDN of a Cloud Connector and click **Add**. Repeat for the other Cloud Connectors in your resource location.

- If you know where you installed the Cloud Connectors in your AD structure, select **Choose locations from Active Directory**, and then navigate to that location. Repeat for the other Cloud Connectors.
- If you want to specify the Cloud Connector addresses in Citrix Group Policy, select **Do it later (Advanced)**, and then confirm that selection when prompted.

When you're done, click **Next**.

4. If you're using the single-session full VDA installer, on the **Additional Components** page, select the components you want to install, such as Profile Management. (This page does not appear if you're using the single-session core VDA installer.)
5. On the **Features** page, click **Next**.
6. On the **Firewall** page, select **Automatically** (if it isn't already). Then click **Next**.
7. On the **Summary** page, click **Install**.
8. On the **Diagnose** page, click **Connect**. Make sure that the check box is selected. When prompted, enter your Citrix account credentials. After your credentials are validated, click **Next**.
9. On the **Finish** page, click **Finish**.

For full installation information, see [Install VDAs](#).

### Install a Windows VDA for Remote PC Access using a command line

- If you're using the single-session core VDA installer: Run `VDAWorkstationCoreSetup.exe`, and include the `/quiet`, `/enable_hdx_ports`, and `/enable_hdx_udp_ports` options. To specify Cloud Connector addresses, use the `/controllers` option.

For example, the following command installs a single-session core VDA. Citrix Workspace app and other non-core services are not installed. The FQDNs of two Cloud Connectors are specified, and ports in the Windows Firewall Service will be opened automatically. The administrator will handle restarts.

```
VDAWorkstationCoreSetup .exe /quiet /controllers "Connector-East.domain.com" "Connector-East2.domain.com" /enable_hdx_ports /noreboot
```

- If you're using the single-session full VDA installer and want to include Profile Management (or other optional components): Run `VDAWorkstationSetup.exe` and include the `/remotepc` and `/includeadditional` options. The `/remotepc` option prevents installation of most additional components. The `/includeadditional` option specifies exactly which additional components you want to install.

For example, the following command prevents installation of all optional additional components except Profile Management.



```
VDAWorkstationSetup.exe /quiet /remotepc /includeadditional "Citrix  
User Profile Manager" , "Citrix User Profile Manager WMI Plugin" /  
controllers "connector.domain.com" "connector2.domain.com" /enable_hdx_ports  
/noresume /noreboot
```

For details, see [Command-line options to install a VDA](#).

## Install a Linux VDA

Follow the guidance in the [Linux documentation](#) for installing a Linux VDA interactively or using the command line.

## Create a Remote PC Access catalog

A resource location containing at least two Cloud Connectors must exist before you can successfully create a catalog.

### Important:

A machine can belong to only one catalog at a time. This restriction is not enforced when you specify the machines to be added to a catalog. However, ignoring the restriction can cause problems later.

1. Sign in to [Citrix Cloud](#).
2. In the upper left menu, select **My Services > Virtual Apps and Desktops**.
3. If you haven't created any catalogs yet, click **Get Started** on the **Welcome** page.
4. Select **Manage > Quick Deploy**.
5. Select **Create Catalog**.
6. On the **Remote PC Access** tab, select a method for assigning users to machines.
7. Enter a name for the catalog and select the resource location you created.
8. Add machines.
9. Click **Create Catalog**.
10. On the **Your Remote PC Access catalog is being created** page, click **Done**.
11. An entry for the new catalog appears on the **Manage > Quick Deploy** dashboard.

After the catalog is successfully created, click one of the links to [add subscribers \(users\) to the catalog](#). This step applies if the catalog uses either the static autoassigned or random pool unassigned user assignment method.

After you create a catalog and add users (if needed), [send the Workspace URL](#) to your users.

## User assignment methods

The user assignment method that you choose when creating a catalog indicates how users are assigned to machines.

- **Static autoassigned:** User assignment occurs when a user logs on to the machine (not using Citrix, for example, in-person or RDP), after a VDA is installed on the machine. Later, if other users log on to that machine (not using Citrix), they are also assigned. Only one user can use the machine at a time. This is a typical setup for either office workers or shift workers who share a computer.

This method is supported for Windows machines. It cannot be used with Linux machines.

- **Static preassigned:** Users are preassigned to machines. (This is usually configured by uploading a CSV file containing machine-user mapping.) There is no need for user logon to establish assignment after the VDA is installed. There is also no need to assign users to the catalog after it's created. This is best for office workers.

This method is supported for Windows and Linux machines.

- **Random pool unassigned:** Users are randomly assigned to an available machine. Only one user can use the machine at a time. This is ideal for computing labs in schools.

This method is supported for Windows machines. It cannot be used with Linux machines.

## Methods for adding machines to a catalog

Remember: Each machine must have a VDA installed on it.

When creating or editing a catalog, there are three ways you can add machines to a catalog:

- Select machine accounts one by one.
- Select OUs.
- Add in bulk using a CSV file. A template is available for you to use for the CSV file.

### Add machine names

This method adds machine accounts one by one.

1. Select your domain.
2. Search for the machine account.
3. Click **Add**.
4. Repeat to add more machines.
5. When you finish adding machines, click **Done**.

## Add OUs

This method adds machine accounts according to the Organizational Unit where they reside. When selecting OUs, choose lower-level OUs for greater granularity. If that granularity is not required, you can choose higher-level OUs.

For example, in the case of `Bank/Officers/Tellers`, select `Tellers` for greater granularity. Otherwise, you can select `Officers` or `Bank`, based on the requirement.

Moving or deleting OUs after they're assigned to a Remote PC Access catalog affects VDA associations and causes issues with future assignments. Ensure that your AD change plan accounts for OU assignment updates for catalogs.

To add OUs:

1. Select your domain.
2. Select the OUs that contain the machines accounts you want to add.
3. Indicate in the check box whether to include subfolders included in your selections.
4. When you finish selecting OUs, click **Done**.

## Add in bulk

1. Click **Download CSV Template**.
2. In the template, add the machine account information (up to 100 entries). The CSV file can also contain the names of users assigned to each machine.
3. Save the file.
4. Either drag the file on to the **Add machines in bulk** page or browse to the file.
5. A preview of the file's content is displayed. If that's not the file you want, you can create another file and then drag or browse to it.
6. When you're finished, click **Done**.

## Manage Remote PC Access catalogs

To display or change a Remote PC Access catalog's configuration information, select the catalog from the **Manage > Quick Deploy** dashboard (click anywhere in the catalog's entry).

- From the **Details** tab, you can add or remove machines.
- From the **Subscribers** tab, you can add or remove users.
- From the **Machines** tab, you can:
  - Add or remove machines: **Add or remove machines** button.
  - Change user assignments: **Remove assignment** trash icon, **Edit machine assignment** in ellipsis menu.
  - See which machines are registered, and place machines in or out of maintenance mode.

## Monitor in Quick Deploy

July 21, 2021

From the **Monitor** dashboard, you can view desktop usage, sessions, and machines in your Citrix Virtual Apps and Desktops deployment. You can also control sessions, power-manage machines, end running applications, and end running processes.

To access the **Monitor** dashboard:

1. Sign in to [Citrix Cloud](#), if you haven't already. In the upper left menu, select **My Services > Virtual Apps and Desktops**.
2. From the **Manage > Quick Deploy** dashboard, select the **Monitor** tab.

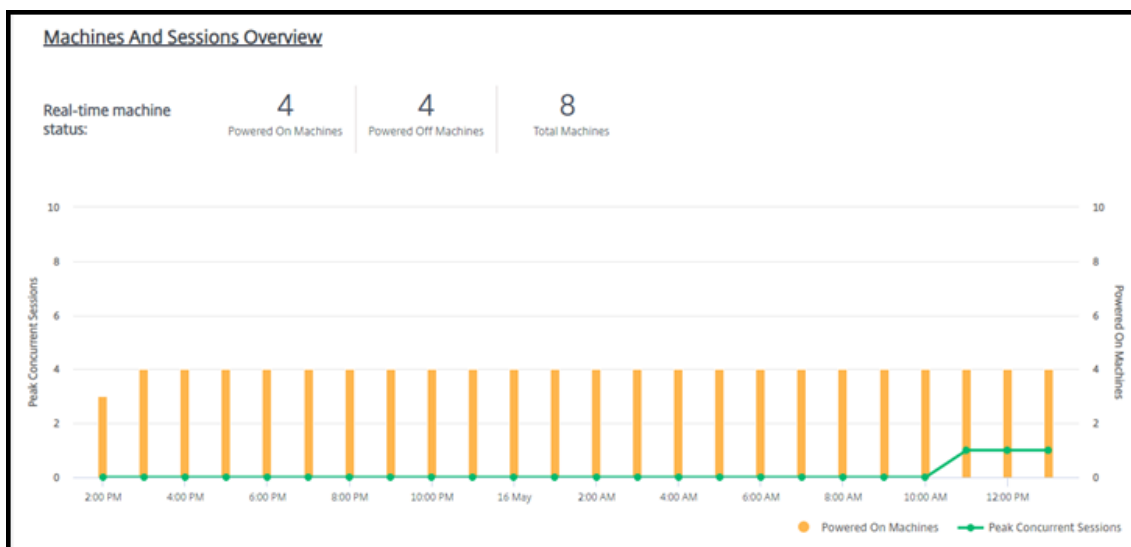
### Monitor desktop usage

The service usage page appears by default when you select the **Monitor** tab. Displays on this page refresh every five minutes.

- **Machine and Sessions Overview:** You can tailor the display to show information about all catalogs (default) or a selected catalog. You can also tailor the time period: the last day, week, month, or three months.

Counts at the top of the display indicate the total number of machines, plus the number of machines that are powered-on and powered-off. Hover over a value to display how many are single-session and multi-session.

The graph below the counts shows the number of powered-on machines and peak concurrent sessions at regular points during the time period you selected. Hover on a point the graph to display the counts at that point.



- **Top 10s:** To tailor a top 10 display, select a time period: the past week (default), month, or three months. You can also tailor the display to show only information about activity involving single-session machines, multi-session machines, or applications.
  - **Top 10 Active Users:** Lists the users who started desktops most frequently during the time period. Hovering on a line displays the total launches.
  - **Top 10 Active Catalogs:** Lists the catalogs with the longest duration during the selected time period. Duration is the sum of all user sessions from that catalog.

### Desktop usage report

To download a report containing information about machine launches during the last month, select **Launch Activity**. A message indicates that the request is being processed. The report downloads automatically to the default download location on the local machine.

### Filter and search to monitor machines and sessions

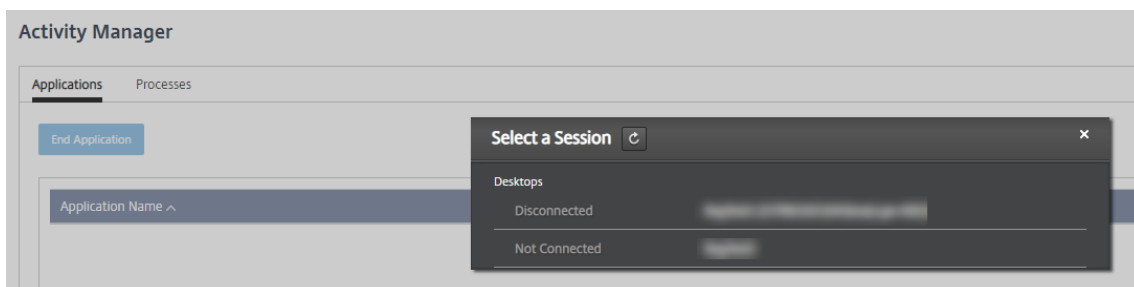
When you're monitoring session and machine information, all machines or sessions are displayed by default. You can:

- Filter the display by machines, sessions, connections, or applications.
- Refine the display of sessions or machines by choosing the criteria you want, building a filter by using expressions.
- Save the filters that you build, for reuse.

### Control a user's applications


You can display and manage applications and processes for a user that has a running session or an assigned desktop.

1. From the **Monitor** dashboard in the service, select **Search** and enter the user name (or the beginning characters of the user name), machine, or endpoint. From the search results, select the item you're looking for. (To collapse the search box without searching, select **Search** again.)
2. Select a session.



The Activity Manager lists the applications and processes for the user's session.

3. To end an application, on the **Applications** tab in Activity Manager, select in the application's row to select that application, and then select **End Application**.



4. To end a process, on the **Processes** tab in Activity Manager, select in the process's row to select that process, and then select **End Process**.
5. To display session details, select **Details** in the upper right. To return to the applications and processes display, select Activity Manager in the upper right.
6. To control the session, select **Session Control > Log Off** or **Session Control > Disconnect**.

## Shadow users

Use the shadow feature to view or work directly on a user's virtual machine or session. You can shadow Windows and Linux VDAs. The user must be connected to the machine that you want to shadow. Verify this by checking the machine name listed in the [User](#) title bar.

Shadowing launches in a new browser tab. Ensure that your browser allows pop-ups from the Citrix Cloud URL.

Shadowing is supported only for users on domain-joined machines. To shadow a non-domain-joined machine, you must set up a bastion machine. For details, see [Bastion access](#).

Shadowing must be initiated from a machine on the same virtual network as the domain-joined machines, and also meet any port requirements.

## Enable shadowing

1. From **Manage > Quick Deploy > Monitor**, go to the **User Details** view.
2. Select the user session, and then select **Shadow** in the **Activity Manager** view or the **Session Details** panel.

## Shadow Linux VDAs

Shadowing is available for Linux VDAs Version 7.16 or and later running the RHEL7.3 or Ubuntu Version 16.04 Linux distributions.

Monitor uses the FQDN to connect to the target Linux VDA. Ensure that the Monitor client can resolve the FQDN of the Linux VDA.

- The VDA must have the `python-websocketify` and `x11vnc` packages installed.
- `noVNC` connection to the VDA uses the WebSocket protocol. By default, `ws://` WebSocket protocol is used. For security reasons, Citrix recommends that you use the secure `wss://` protocol. Install SSL certificates on each Monitor client and Linux VDA.

Follow the instructions in Session Shadowing to configure your Linux VDA for shadowing.

1. After you enable shadowing, the shadowing connection initializes and a confirmation prompt appears on the user device.
2. Instruct the user to select **Yes** to start the machine or session sharing.
3. The administrator can view only the shadowed session.

## Shadow Windows VDAs

Windows VDA sessions are shadowed using Windows Remote Assistance. Enable the `Use Windows Remote Assistance` feature when installing the VDA.

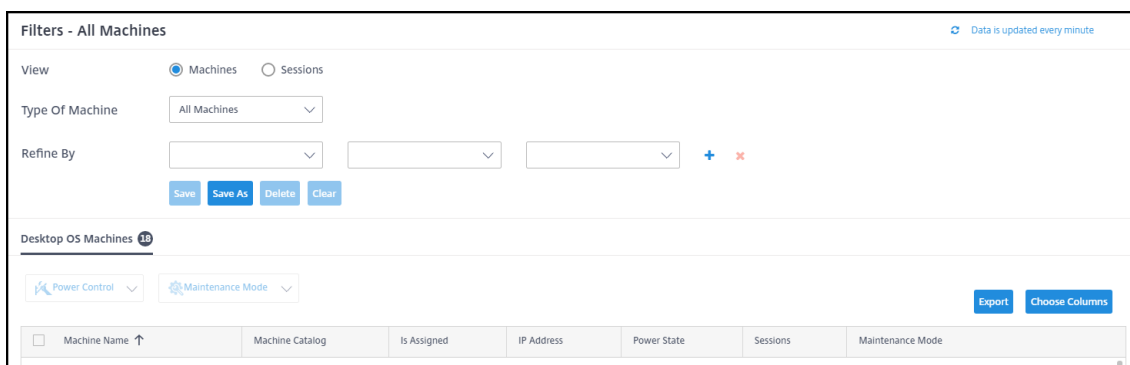
1. After you enable shadowing, the shadowing connection initializes and a dialog box prompts you to open or save the `.msrc incident` file.
2. Open the incident file with the Remote Assistance Viewer, if it's not already selected by default. A confirmation prompt appears on the user device.
3. Instruct the user to select **Yes** to start the machine or session sharing.
4. For more control, ask the user to share keyboard and mouse control.

## Monitor and control sessions


Session displays are updated every minute.

In addition to viewing sessions, you can disconnect one or more sessions or log off users from sessions.

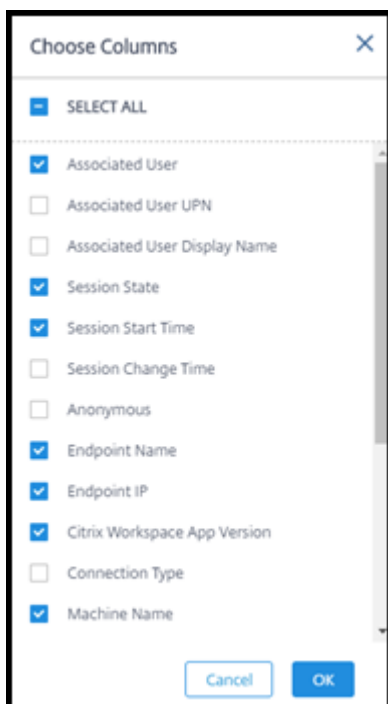
1. From **Manage > Quick Deploy > Monitor**, select **Filters**.



2. Select the **Sessions** view.



3. To tailor the display, select **Choose Columns** and select the check boxes of items you want to appear. When you're done, select **OK**. The sessions display refreshes automatically.



4. Select the check box to the left of each session you want to control.

5. To log off or disconnect the session, elect either **Session Control > Log Off** or **Session Control > Disconnect**.

Remember that the power management schedule for the catalog can also control disconnecting sessions and logging off users from disconnected sessions.



As an alternative to the above procedure you can also **Search** for a user, select the session you want to control, and then display session details. The log off and disconnect options are available there, too.

### Session information report

To download session information, select **Export** on the sessions display. A message indicates that the request is being processed. The report downloads automatically to the default download location on the local machine.

### Monitor and power control machines

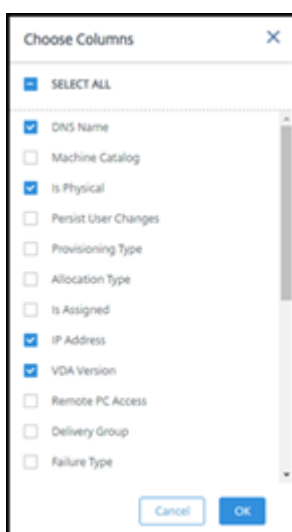
Machine displays are updated every minute.

1. From **Manage > Quick Deploy > Monitor**, select **Filters**.
2. Select the **Machines** view.


<input type="checkbox"/>	Machine Name ↑	Is Assigned	IP Address	Delivery Group	Failure Type	Failure Reason	Failure Time	Power State	Sessions	Maintenance Mode
<input type="checkbox"/>	[blurred]	[blurred]	[blurred]	[blurred]	n/a	None		On	0	Off
<input type="checkbox"/>	[blurred]	[blurred]	[blurred]	[blurred]	n/a	None		On	0	Off
<input type="checkbox"/>	[blurred]	[blurred]	[blurred]	[blurred]	n/a	None		Off	0	Off

By default, the display lists single-session OS machines. Alternatively, you can display multi-session machines.

3. To tailor the display, select **Choose Columns** and select the check boxes of items you want to appear. When you're done, select **OK**. The machines display refreshes automatically.



4. To power-control machines or place them in or out of maintenance mode, select the check box to the left of each machine you want to control.
5. To power-control the selected machines, select **Power Control** and select an action.



6. To place the selected machines in or out of maintenance mode, select **Maintenance Mode > ON** or **Maintenance Mode > OFF**.

When you use the search feature to find and select a machine, you see machine details, utilization, historical utilization (from the last seven days), and average IOPS.

### Machine information report

To download session information, select **Export** on the machines display. A message indicates that the request is being processed. The report downloads automatically to the default download location on the local machine.

### Checking app and desktop health

Probing automates the process of checking the health of published apps and desktops. The health check results are available through the **Monitor** dashboard. For details, see:

- [Application probing](#)
- [Desktop probing](#)

## Troubleshoot in Quick Deploy

July 21, 2021

## Introduction

Resource locations contain the machines that deliver desktops and apps. Those machines are created in catalogs, so the catalogs are considered part of the resource location. Each resource location also contains Cloud Connectors. Cloud Connectors enable Citrix Cloud to communicate with the resource location. Usually, Citrix installs and updates the Cloud Connectors.

Optionally, you can initiate several Cloud Connector and resource location actions. See:

- [Resource location actions](#)
- [Resource location settings when creating a catalog](#)

This service has troubleshooting and supportability tools that can help resolve configuration and communication issues with the machines that deliver desktops and apps (the VDAs). For example, creating a catalog might fail, or users might be unable to start their desktop or apps.

This troubleshooting includes gaining access to your Citrix Managed Azure subscription through a bastion machine or direct RDP. After gaining access to the subscription, you can use Citrix supportability tools to locate and resolve issues. For details, see:

- [VDA troubleshooting using a bastion or direct RDP](#)
- [Bastion access](#)
- [Direct RDP access](#)

## VDA troubleshooting using a bastion or direct RDP

The supportability features are for people who have experience with troubleshooting Citrix issues. This includes:

- Citrix Service Providers (CSPs) and others who have the technical knowledge and troubleshooting experience with Citrix Virtual Apps and Desktops products.
- Citrix Support personnel.

If you're not familiar or comfortable with troubleshooting Citrix components, you can request help from Citrix Support. Citrix Support representatives might ask you to set up one of the access methods described in this section. However, the Citrix representatives do the actual troubleshooting, using Citrix tools and technologies.

### Important:

These supportability features are valid only for domain-joined machines. If the machines in your catalogs are not domain joined, you're guided to request troubleshooting help from Citrix Support.

## Access methods

These access methods are valid only for the Citrix Managed Azure subscription. For more information, see [Azure subscriptions](#).

Two supportability access methods are provided.

- Access your resources through a bastion machine in the customer's dedicated Citrix Managed Azure subscription. The bastion is a single point of entry that allows access to the machines in the subscription. It provides a secure connection to those resources by allowing remote traffic from IP addresses in a specified range.

The steps in this method include:

- Create the bastion machine
- Download an RDP agent
- RDP to the bastion machine
- Connect from the bastion machine to the other Citrix machines in your subscription

The bastion machine is intended for short-term use. This method is intended for issues involving the creation of catalogs or image machines.

- Direct RDP access to the machines in the customer's dedicated Citrix Managed Azure subscription. To permit RDP traffic, port 3389 must be defined in the Network Security Group.

This method is intended for catalog issues other than creation, such as users unable to start their desktops.

Remember: As an alternative to these two access methods, contact Citrix Support for help.

## Bastion access

1. From **Manage > Quick Deploy**, expand **Troubleshoot & Support** on the right.
2. Click **View troubleshooting options**.
3. On the **Troubleshoot** page, select either of the first two issue types, and then click **Use our troubleshooting machine**.
4. On the **Troubleshoot with Bastion Machine** page, select the catalog.
  - If the machines in the selected catalog are not domain joined, you're instructed to contact Citrix Support.
  - If a bastion machine has already been created with RDP access to the selected catalog's network connection, skip to step 8.
5. The RDP access range is displayed. If you want to restrict RDP access to a smaller range than allowed by the network connection, select the **Restrict RDP access to only computers in IP address range** check box and then enter the desired range.

6. Type a username and password that you'll use to log in when you RDP to the bastion machine. [Password requirements](#).

Do not use Unicode characters in the username.

7. Click **Create Bastion Machine**.

When the bastion machine is successfully created, the page title changes to **Bastion – connection**.

If the bastion machine creation fails (or if it fails during operation), click **Delete** at the bottom of the failure notification page. Try to create the bastion machine again.

You can change the RDP range restriction after the bastion machine is created. Click **Edit**. Enter the new value and then click the check mark to save the change. (Click **X** to cancel the change.)

8. Click **Download RDP File**.
9. RDP to the bastion, using the credentials you specified when creating the bastion. (The bastion machine's address is embedded in the RDP file you downloaded.)
10. Connect from the bastion machine to the other Citrix machines in the subscription. You can then collect logs and run diagnostics.

Bastion machines are powered on when they are created. To save costs, machines are powered off automatically if they remain idle after startup. The machines are deleted automatically after several hours.

You can power manage or delete a bastion machine, using the buttons at the bottom of the page. If you choose to delete a bastion machine, you must acknowledge that any active sessions on the machine will end automatically. Also, any data and files that were saved on the machine will be deleted.

### Direct RDP access

1. From **Manage > Quick Deploy**, expand **Troubleshoot & Support** on the right.
2. Click **View troubleshooting options**.
3. On the **Troubleshoot** page, select **Other catalog issue**.
4. On the **Troubleshoot with RDP Access** page, select the catalog.  
If RDP has already been enabled to the selected catalog's network connection, skip to step 7.
5. The RDP access range is displayed. If you want to restrict RDP access to a smaller range than permitted by the network connection, select the **Restrict RDP access to only computers in IP address range** check box and then enter the desired range.
6. Click **Enable RDP Access**.

When RDP access is successfully enabled, the page title changes to **RDP Access – connection**.

If RDP access is not successfully enabled, click **Retry Enabling RDP** at the bottom of the failure notification page.

7. Connect to machines using your Active Directory administrator credentials. You can then collect logs and run diagnostics.

## Get help

If you still have problems, open a ticket by following the instructions in [How to Get Help and Support](#).

## Quick Deploy reference

July 21, 2021

### Catalog tabs on the Quick Deploy dashboard

From the **Manage > Quick Deploy** dashboard in the service, click anywhere in the catalog's entry. The following tabs contain information about the catalog:

- **Details:** Lists the information specified when the catalog was created (or its most recent edit). It also contains information about the image that was used to create the catalog.

From this tab, you can:

- [Change the image](#) that is used in the catalog.
  - [Delete the catalog](#).
  - Access the page containing details for the resource location used by the catalog.
- **Desktop:** Available only for catalogs containing single-session (static or random) machines. From this tab, you can change the name and description of the catalog.
  - **Desktop and Apps:** The **Desktops and Apps** tab is available only for catalogs containing multi-session machines. From this tab, you can:
    - [Add](#), [edit](#), or [remove](#) applications that the catalog's users can access in Citrix Workspace.
    - Change the name and description of the catalog.
  - **Subscribers:** Lists all users, including their type (user or group), account name, display name, plus their Active Directory domain and user principal name.

From this tab, you can [add or remove users](#) for a catalog.

- **Machines:** Shows the total number of machines in the catalog, plus the number of registered machines, unregistered machines, and machines that have maintenance mode turned on.

For each machine in the catalog, the display includes each machine's name, power state (on/off), registration state (registered/unregistered), assigned users, session count (0/1), and maintenance mode status (an icon indicating on or off).

From this tab, you can:

- Add or delete a machine
- Start, restart, force restart, or shut down a machine
- Turn a machine's maintenance mode on or off

For details, see [Manage catalogs](#). Many of the machine actions are also available from the **Monitor** tab on the Quick Deploy dashboard. See [Monitor and power control machines](#).

- **Power Management:** Enables you to manage when machines in the catalog are powered on and off. A schedule also indicates when idle machines are disconnected.

You can configure a power schedule when you create a custom catalog or later. If no schedule is explicitly set, a machine powers off when a session ends.

When creating a catalog using quick create, you cannot select or configure a power schedule. By default, quick create catalogs use the Cost Saver preset schedule. However, you can edit that catalog later and change the schedule.

For details, see [Manage power management schedules](#).

## DNS servers

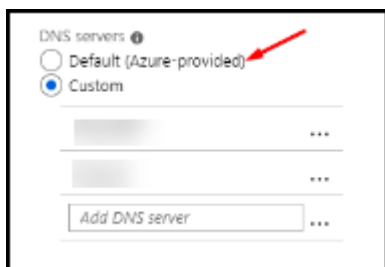
This section applies to all deployments that contain domain-joined machines. You can ignore this section if you use only non-domain-joined machines.

1. Before creating a domain-joined catalog (or a connection, if you're using a Citrix Managed Azure subscription), check whether you have DNS server entries that can resolve public and private domain names.

When the service creates a catalog or a connection, it looks for at least one valid DNS server entry. If no valid entries are found, the creation operation fails.

Where to check:

- If you are using your own Azure subscription, check the **DNS servers** entry in your Azure.
  - If you are using a Citrix Managed Azure subscription and creating an Azure VNet peering connection, check the **DNS servers** entry in the Azure VNet that you're peering.
  - If you are using a Citrix Managed Azure subscription and creating an SD-WAN connection, check the DNS entries in the [SD-WAN Orchestrator](#).
2. In Azure, the **Custom** setting must have at least one valid entry. This service cannot be used with the **Default (Azure-provided)** setting.



- If **Default (Azure-provided)** is enabled, change the setting to **Custom**, and add at least one DNS server entry.
  - If you already have DNS server entries under **Custom**, verify that the entries you want to use with this service can resolve public and private domain IP names.
  - If you do not have any DNS servers that can resolve domain names, Citrix recommends adding an Azure-provided DNS server that has those capabilities.
3. If you change any DNS server entries, restart all machines that are connected to the virtual network. The restart assigns the new DNS server settings. (The VMs continue using their current DNS settings until the restart.)

If you want to change DNS addresses later, after a connection is created:

- When using your own Azure subscription, you can change them in Azure (as described in the preceding steps). Or, you can change them in this service.
- When using a Citrix Managed Azure subscription, this service does not synchronize DNS address changes that you make in Azure. However, you can change DNS settings for the connection in this service.

Keep in mind that changing DNS server addresses can potentially cause connectivity issues for machines in catalogs that use that connection.

### Adding DNS servers through this service

Before adding a DNS server address to a connection, make sure that the DNS server can resolve public and internal domain names. Citrix recommends that you test connectivity to a DNS server before adding it.

1. To add, change, or remove a DNS server address when you're creating a connection, select **Edit DNS servers** on the **Add connection type** page. Or, if a message indicates that no DNS server addresses were found, select **Add DNS Servers**. Continue with step 3.
2. To add, change, or remove a DNS server address for an existing connection:
  - a) From **Manage > Quick Deploy**, expand **Network Connections** on the right.
  - b) Select the connection you want to edit.
  - c) Select **Edit DNS servers**.
3. Add, change, or remove addresses.



- a) To add an address, select **Add DNS server** and then enter the IP address.
  - b) To change an address, click inside the address field and change the numbers.
  - c) To remove an address, select the trash icon next to the address entry. You cannot remove all DNS server addresses. The connection must have at least one.
4. When you're done, select **Confirm Changes** at the bottom of the page.
  5. Restart all machines that use that connection. The restart assigns the new DNS server settings. (The VMs continue using their current DNS settings until the restart.)

## Policies

### Set group policies for non-domain-joined machines

1. RDP to the machine that is being used for the image.
2. Make sure that Visual Studio 2017 is installed on the machine.
3. Install Citrix Group Policy Management:
  - a) Browse to [CTX220345](#). Download the attachment.
  - b) Double-click the downloaded file. In the `Group Policy Templates 1912 > Group Policy Management` folder, double-click `CitrixGroupPolicyManagement_x64.msi`.
4. Using the **Run** command, launch `gpedit.msc` to open the Group Policy Editor.
5. In `User Configuration Citrix Policies > Unfiltered`, select **Edit Policy**.
6. Enable policy settings as needed. For example:
  - When working in **Computer Configuration** or **User Configuration** (depending on what you want to configure) on the **Settings** tab, in `Category > ICA / Printing`, select **Auto-create PDF Universal Printer** and set to `Enabled`.
  - If you want logged-in users to be administrators of their desktop, add the **Interactive User** group to the built-in administrators group.
7. When you're done, save the image.
8. Either [update the existing catalog](#) or [create a new catalog](#) using the new image.

### Set group policies for domain-joined machines

1. Ensure that the machine has minimum Visual Studio 2017 installed.
2. Ensure that the Group Policy Management feature is installed.
  - On a Windows multi-session machine, add the Group Policy Management feature, using the Windows tool for adding roles and features (such as **Add Roles and Features**).

- On a Windows single-session machine, install the Remote Server Administration Tools for the appropriate OS. (This installation requires a domain admin account.) After that installation, the Group Policy Management console is available from the **Start** menu.
3. Download and install the Citrix Group Policy management package from the Citrix [download page](#), and then configure policy settings as needed. Follow the procedure in Set group policies for non-domain-joined machines, step 2 through the end.

See the [Policy settings reference](#) articles to learn about what's available. All policy features are available from the service's Full Configuration interface.

## Resource location actions

Citrix automatically creates a resource location and two Cloud Connectors when you create the first catalog for publishing desktops and apps. You can specify some information related to the resource location when you create a catalog. See [Resource location settings when creating a catalog](#).

For Remote PC Access, you create the resource location and Cloud Connectors.

This section describes available actions after a resource location is created.

1. From **Manage > Quick Deploy**, expand **Cloud Subscriptions** on the right.
2. Select the subscription.
  - The **Details** tab shows the number and names of catalogs and images in the subscription. It also indicates the number of machines that can deliver desktops or apps. That count does not include machines used for other purposes, such as images, Cloud Connectors, or RDS license servers.
  - The **Resource Locations** tab lists each resource location. Each resource location entry includes the status and address of each Cloud Connector in the resource location.

The ellipsis menu in a resource location's entry contains the following actions.

### Run Health Check

Selecting **Run Health Check** starts the connectivity check immediately. If the check fails, the Cloud Connector's state is unknown, because it is not communicating with Citrix Cloud. You might want to restart the Cloud Connector.

### Restart Connectors

Citrix recommends restarting only one Cloud Connector at a time. Restarting takes the Cloud Connector offline, and disrupts user access and machine connectivity.

Select the check box for the Cloud Connector you want to restart. Select **Restart**.

### **Add Connectors**

Adding a Cloud Connector typically takes 20 minutes to complete.

Provide the following information:

- How many Cloud Connectors to add.
- Domain service account credentials, which are used to join the Cloud Connector machines to the domain.
- Machine performance.
- Azure resource group. The default is the resource group last used by the resource location.
- Organizational Unit (OU). The default is the OU last used by the resource location.
- Whether your network requires a proxy server for internet connectivity. If you indicate **Yes**, provide the proxy server FQDN or IP address, and port number.

When you're done, select **Add Connectors**.

### **Delete Connectors**

If a Cloud Connector cannot communicate with Citrix Cloud, and a restart does not resolve the issue, Citrix Support might recommend deleting that Cloud Connector.

Select the check box for the Cloud Connector you want to delete. Then select **Delete**. When prompted, confirm the deletion.

You can also delete an available Cloud Connector. However, if deleting that Cloud Connector would result in fewer than two available Cloud Connectors in the resource location, you're not allowed to delete the selected Cloud Connector.

### **Select Update Time**

Citrix automatically provides software updates for the Cloud Connectors. During an update, one Cloud Connector is taken offline and updated, while other Cloud Connectors remain in service. When the first update completes, another Cloud Connector is taken offline and updated. This process continues until all Cloud Connectors in the resource location are updated. The best time to start updates is usually outside your typical business hours.

Choose the time to begin updates, or indicate that you want updates to start when an update is available. When you're done, select **Save**.

### **Rename**

Enter the new name for the resource location. Select **Save**.

## Configure Connectivity

Indicate whether users can access desktops and apps through the Citrix Gateway service, or only from within your corporate network.

## Profile Management

[Profile Management](#) ensures that personal settings apply to users' virtual applications, regardless of the location of the user device.

Configuring Profile Management is optional.

You can enable Profile Management with the profile optimization service. This service provides a reliable way for managing these settings in Windows. Managing profiles ensures a consistent experience by maintaining a single profile that follows the user. It consolidates automatically and optimizes user profiles to minimize management and storage requirements. The profile optimization service requires minimal administration, support, and infrastructure. Also, profile optimization provides users with an improved logon and logoff experience.

The profile optimization service requires a file share where all the personal settings persist. You manage the file servers. We recommend setting up network connectivity to allow access to these file servers. You must specify the file share as a UNC path. The path can contain system environment variables, Active Directory user attributes, or Profile Management variables. To learn more about the format of the UNC text string, see [Specify the path to the user store](#).

When enabling Profile Management, consider further optimizing the user's profile by configuring folder redirection to minimize the effects of the user profile size. Applying folder redirection complements the Profile Management solution. For more information, see [Microsoft Folder Redirection](#).

## Configure the Microsoft RDS License Server for Windows Server workloads

This service accesses Windows Server remote session capabilities when delivering a Windows Server workload, such as Windows 2016. This typically requires a Remote Desktop Services client access license (RDS CAL). The VDA must be able to contact an RDS license server to request RDS CALs.

Install and activate the license server. For more information, see the Microsoft document [Activate the Remote Desktop Services License Server](#). For proof of concept environments, you can use the grace period provided by Microsoft.

With this method, you can have the Citrix service apply the license server settings. You can configure the license server and per user mode in the RDS console on the image. You can also configure the license server using Microsoft Group Policy settings. For more information, see the Microsoft document [License your RDS deployment with client access licenses \(CALs\)](#).

To configure the RDS license server using Group Policy settings

1. Install a Remote Desktop Services License Server on one of the available VMs. The VM must always be available. The Citrix service workloads must be able to reach this license server.
2. Specify the license server address and per-user license mode using Microsoft Group Policy. For details, see the Microsoft document [Specify the Remote Desktop Licensing Mode for an RD Session Host Server](#).

Windows 10 workloads require appropriate Windows 10 license activation. We recommend that you follow Microsoft documentation to activate Windows 10 workloads.

## Consumption commitment usage

### Note:

This feature is in preview.

From **Manage > Quick Deploy**, select the **General** card. The **Consumption** value indicates how much consumption has been used in the current calendar month. That value includes monthly and term commitments.

When you select **General**, the **Notifications** tab includes:

- Total consumption used for the month (monthly and term).
- Number of units of monthly consumption commitment.
- Percentage of term consumption commitment.

The values and progress bars can alert you to potential or actual usage overages.

Actual data can take 24 hours to appear. Usage and billing data are considered final 72 hours after the end of a calendar month.

For more usage information, see [Monitor licenses and active usage](#).

You can optionally request notifications to appear in the **Manage > Quick Deploy** dashboard when consumption usage (for monthly, term, or both commitments) reaches a specified level. By default, notifications are disabled.

1. On the **Notifications** tab, select **Edit Notification Preferences**.
2. To enable notifications, click the slider so that the check mark appears.
3. Enter a value. Repeat for the other consumption type, if needed.
4. Select **Save**.

To disable notifications, click the slider so that the check mark no longer appears, and then select **Save**.

## Monitor Citrix license usage

To view your Citrix license usage information, follow the guidance in [Monitor licenses and active usage](#). You can view:

- Licensing summary
- Usage reports
- Usage trends and license activity
- Licensed users

You can also release licenses.

## Load balancing

Load balancing applies to multi-session machines, not single-session machines.

### Important:

Changing the load balancing method affects all catalogs in your deployment. That includes all catalogs created using any supported host type, cloud-based and on-premises, regardless of interface used to create them (such as Studio or Quick Deploy).

Make sure you have maximum session limits configured for all catalogs before proceeding.

- In Quick Deploy, that setting is located on each catalog's **Details** tab.
- In Full Configuration, see [Load balance machines](#).

Load balancing measures the machine load, and determines which multi-session machine to select for an incoming user session under the current conditions. This selection is based on the configured load balancing method.

You can configure one of two load balancing methods: horizontal or vertical. The method applies to all multi-session catalogs (and therefore, all multi-session machines) in your service deployment.

- **Horizontal load balancing:** An incoming user session is assigned to the least-loaded powered-on machine available.

Simple example: You have two machines configured for 10 sessions each. The first machine handles five concurrent sessions. The second machine handles five.

Horizontal load balancing offers high user performance, but it can increase costs as more machines are kept powered-on and busy.

This method is enabled by default.

- **Vertical load balancing:** An incoming user session is assigned to the powered-on machine with the highest load index. (The service calculates and then assigns a load index for every multi-session machine. The calculation considers factors such as CPU, memory, and concurrency.)

This method saturates existing machines before moving on to new machines. As users disconnect and free up capacity on existing machines, new load is assigned to those machines.

Simple example: You have two machines configured for 10 sessions each. The first machine handles the first 10 concurrent sessions. The second machine handles the eleventh session.

With vertical load balancing, sessions maximize powered-on machine capacity, which can save machine costs.

To configure the load balancing method:

1. From **Manage > Quick Deploy**, expand **General** on the right.
2. Under **Global Settings**, select **View All**.
3. On the **Global Settings** page, under **Multi-Session Catalog Load Balancing**, choose the load balancing method.
4. Select **Confirm**.

### Create a catalog in a network that uses a proxy server

Follow this procedure if your network requires a proxy server for Internet connectivity, and you are using your own Azure subscription. (Using a Citrix Managed Azure subscription with a network requiring a proxy server is not supported.)

1. From **Manage > Quick Deploy**, start the [catalog creation process](#) by providing the required information and then selecting **Create Catalog** at the bottom of the page.
2. The catalog creation fails because of the proxy requirement. However, a resource location is created. That resource location's name begins with "DAS", unless you provided a resource location name when creating the catalog. On the **manage > Quick Deploy** dashboard, expand **Cloud Subscriptions** on the right. On the **Resource Locations** tab, check whether the newly created resource location has any Cloud Connectors in it. If it does, delete them.
3. In Azure, create two VMs (see [Cloud Connector system requirements](#)). Join those machines to the domain.
4. From the Citrix Cloud console, [install a Cloud Connector](#) on each VM. Make sure that the Cloud Connectors are in the same resource location that was created earlier. Follow the guidance in:
  - [Cloud Connector proxy and firewall configuration](#)
  - [System and connectivity requirements](#)
5. From **Manage > Quick Deploy**, repeat the catalog creation process. When the catalog is created, it uses the resource location and Cloud Connectors you created in the preceding steps.

### Get help

- Review [Troubleshoot](#).
- If you need further assistance with this service, open a ticket by following the guidance in [How to Get Help and Support](#).

## Create delivery groups

September 2, 2021

### Important:

If you are using an on-premises StoreFront with the Virtual Apps and Desktops service, do not use Library to assign resources when creating delivery groups. Instead, use **Manage > Full Configuration** to assign resources to users. If you use Library in this scenario, resources might not be enumerated to users.

When creating a delivery group in **Manage > Full Configuration**, on the **Users** page, do not select **Leave user management to Citrix Cloud**. Instead, select a different option (**Allow any authenticated users to use this delivery group** or **Restrict use of this delivery group to the following users**).

### Introduction

A delivery group is a collection of machines selected from one or more machine catalogs. The delivery group can also specify which users can use those machines, plus the applications and desktops available to those users.

Creating a delivery group is the next step in configuring your deployment after creating a machine catalog. Later, you can change the initial settings in the first delivery group and create other delivery groups. There are also features and settings you can configure only when editing a delivery group, not when creating it.

Before creating a delivery group:

- Review this section to learn about the choices you make and information you supply.
- Ensure that you have created a connection to the hypervisor, cloud service, other resource that hosts your machines.
- Ensure that you have created a machine catalog containing virtual or physical machines.

To launch the delivery group creation wizard:

1. Sign in to [Citrix Cloud](#). In the upper left menu, select **My Services > Virtual Apps and Desktops**.
2. Select **Manage**.
3. If this is the first delivery group being created, the console guides you to the correct selection (such as “Set up delivery groups to be displayed as services”). The delivery group creation wizard opens and walks you through the process.
4. If you already created a delivery group and want to create another, from **Manage > Full Configuration**, select **Delivery Groups** in the left pane. Then select **Create Delivery Group** in the action bar.



The wizard walks you through the pages described below. The wizard pages you see might be different, depending on the selections you make.

## Step 1. Machines

Select a machine catalog and select the number of machines you want to use from that catalog.

Good to know:

- At least one machine must remain unused in a selected catalog.
- A catalog can be specified in more than one delivery group. However, a machine can be used in only one delivery group.
- A delivery group can use machines from more than one catalog. However, those catalogs must contain the same machine types (multi-session OS, single-session OS, or Remote PC Access). In other words, you cannot mix machine types in a delivery group. Similarly, if your deployment has catalogs of Windows machines and catalogs of Linux machines, a delivery group can contain machines from either OS type, but not both.
- A MCS delivery group can only add a MCS type catalog.
- Citrix recommends that you install or upgrade all VDAs with the latest version, and then upgrade machine catalogs and delivery groups as needed. When creating a delivery group, if you select machines that have different VDA versions installed, the delivery group will be compatible with the earliest VDA version. For example, if one of the machines you select has VDA version 7.1 installed and other machines have a later version, all machines in the group can use only those features that were supported in VDA 7.1. This means that some features that require newer VDA versions might not be available in that delivery group.
- The following compatibility checks are performed:
  - MinimumFunctionalLevel must be compatible
  - SessionSupport must be compatible
  - AllocationType must be compatible for SingleSession
  - ProvisioningType must be compatible
  - PersistChanges must be compatible for MCS and Citrix Provisioning
  - RemotePC catalog is only compatible with RemotePC catalog
  - AppDisk related check

## Step 2. Delivery type

This page appears only if you chose a machine catalog containing static (assigned) single-session OS machines. Choose either **Applications** or **Desktops**. You cannot enable both.

(If you selected machines from a multi-session OS or single-session OS random (pooled) catalog, the delivery type is assumed to be applications and desktops. You can deliver applications, desktops, or both.

### Step 3. AppDisks

Ignore this page. Select **Next**.

### Step 4. Users

Specify the users and user groups who can use the applications and desktops in the delivery group.

As an alternative to specifying applications in the delivery group wizard (as described in this section), you can configure them through the Citrix Cloud library. (See important exception at top of this article if using an on-premises StoreFront.)

#### Where user lists are specified

Active Directory user lists are specified when you create or edit the following:

- A deployment's user access list, which is not configured through this console. By default, the application entitlement policy rule includes everyone. See the PowerShell SDK [BrokerAppEntitlementPolicyRule](#) cmdlets for details.
- Delivery groups.
- Applications.

The list of users who can access an application is formed by the intersection of the above user lists.

#### Authenticated and unauthenticated users

There are two types of users: authenticated and unauthenticated (unauthenticated is also called anonymous). You can configure one or both types in a delivery group.

- **Authenticated:** To access applications and desktops, the users and group members you specify by name must present credentials such as smart card or user name and password to StoreFront or Citrix Workspace app. (For delivery groups containing single-session OS machines, you can import user data (a list of users) later by editing the delivery group.)
- **Unauthenticated (anonymous):** For delivery groups containing multi-session OS machines, you can allow users to access applications and desktops without presenting credentials to StoreFront or Citrix Workspace app. For example, at kiosks, the application might require credentials, but the Citrix access portal and tools do not. An Anonymous Users Group is created when you install the first Delivery Controller.

To grant access to unauthenticated users, each machine in the delivery group must have a multi-session OS VDA installed. When unauthenticated users are enabled, you must have an unauthenticated StoreFront store.

Unauthenticated user accounts are created on demand when a session is launched, and named AnonXYZ, in which XYZ is a unique three-digit value.

Unauthenticated user sessions have a default idle timeout of 10 minutes, and are logged off automatically when the client disconnects. Reconnection, roaming between clients, and Workspace Control are not supported.

The following table describes your choices on the **Users** page:

Enable access for	Add/assign users and user groups?	Enable the “Give access to unauthenticated users” check box?
Only authenticated users	Yes	No
Only unauthenticated users	No	Yes
Both authenticated and unauthenticated users	Yes	Yes

## Step 5. Applications

Good to know:

- You cannot add applications to Remote PC Access delivery groups.
- By default, new applications you add are placed in a folder named Applications. You can specify a different folder. For details, see the Applications article.
- You can change the properties for an application when you add it to a delivery group, or later. For details, see the Applications article.
- If you try to add an application and one with the same name already exists in that folder, you are prompted to rename the application you are adding. If you decline, the application is added with a suffix that makes it unique within that application folder.
- When you add an application to more than one delivery group, a visibility issue can occur if you do not have permission to view the application in all those delivery groups. In such cases, either consult an administrator with greater permissions or have your scope extended to include all the delivery groups to which the application was added.
- If you publish two applications with the same name to the same users, change the Application name (for user) property. Otherwise, users see duplicate names in Citrix Workspace app.

Select the **Add** menu to display the application sources.

- **From Start menu:** Applications that are discovered on a machine created from the image in the selected catalog. When you select this source, a new page launches with a list of discovered applications; select those you want to add and then select **OK**.

- **Manually defined:** Applications located in the deployment or elsewhere in your network. When you select this source, a new page launches where you type the path to the executable, working directory, optional command line arguments, and display names for administrators and users. After entering this information, select **OK**.
- **Existing:** Applications previously added to the deployment, perhaps in another delivery group. When you select this source, a new page launches with a list of discovered applications; select those you want to add and then select **OK**.
- **App-V:** Applications in App-V packages. When you select this source, a new page launches where you select the App-V server or the Application Library. Select the applications you want to add from the resulting display and then select **OK**.

If an application source or application is not available or valid, it is either not visible or cannot be selected. For example, the **Existing** source is not available if no applications have been added to the deployment. Or, an application might not be compatible with the supported session types on machines in the selected machine catalog.

As an alternative to specifying applications in the delivery group wizard (as described in this section), you can configure them through the Citrix Cloud library. (See important exception at top of this article if using an on-premises StoreFront.)

## Step 6. Desktops (or Desktop Assignment Rules)

The title of this page depends on the machine catalog you chose earlier in the wizard:

- If you chose a catalog containing pooled machines, this page is titled **Desktops**.
- If you chose a catalog containing assigned machines and specified “Desktops” on the **Delivery Type** page, this page is titled **Desktop User Assignments**.
- If you chose a catalog containing assigned machines and specified “Applications” on the **Delivery Type** page, this page is titled **Application Machine User Assignments**.

Select **Add**. In the dialog box:

- In the **Display name** and **Description** fields, type the information to be displayed in Citrix Workspace app.
- To add a tag restriction to a desktop, select **Restrict launches to machines with this tag** and then select the tag from the menu.
- Using the radio buttons, indicate who can launch a desktop (for groups with pooled machines) or who can be assigned a machine when they launch the desktop (for groups with assigned machines). The users can be either everyone who can access this delivery group, or specific users and user groups.
- If the group contains assigned machines, specify the maximum number of desktops per user. This must be a value of one or greater.

- Enable or disable the desktop (for pooled machines) or desktop assignment rule (for assigned machines). Disabling a desktop stops desktop delivery. Disabling a desktop assignment rule stops desktop auto-assignment to users.
- When you are finished with the dialog box, select **OK**.

As an alternative to specifying desktops in the delivery group wizard (as described in this section), you can configure them through Citrix Cloud library. (See important exception at top of article if using an on-premises StoreFront.)

### Step 7. Summary

Enter a name for the delivery group. You can also (optionally) enter a description, which appears in Workspace app and in the Full Configuration management interface.

Review the summary information and then select **Finish**. If you did not select any applications or specify any desktops to deliver, you are asked if you want to continue.

### If you don't specify users or applications in the wizard

As an alternative to specifying users and applications in a delivery group, you can specify them in the Citrix Cloud console. (See important exception at top of this article if using an on-premises StoreFront.)

1. In the Citrix Cloud Console, select **Library**.
2. Find the card containing the resources (applications or desktops) you want. Hover over the ellipsis menu in the upper right corner and select **Manage Subscribers**.
3. In the **Manage subscribers** dialog, under **Add Subscribers** in the left menu, select subscribers (users). If you have multiple subscribers, you might need to type one or more characters of the domain group containing those users in the right search field. Matches appear in the table below the two fields. Select the correct match. (If there's only one match, it's automatically selected.) When the **Status** field indicates **Ready**, select the **X** in the upper right corner to close the dialog.
4. Refresh the **Resources** page. The lower left corner of the resource card contains a value that indicates domain users have been selected.

For more information about the Library, see [Assign users and groups to service offerings using Library](#).

### More information

- [Manage delivery groups](#)
- [Applications](#)

## Manage delivery groups

September 9, 2021

### Introduction

This article describes procedures for managing delivery groups from the management console. In addition to changing the settings specified when creating the group, you can configure other settings that are not available when you create a delivery group.

The procedures are organized by categories: general, users, machines, and sessions. Some tasks span more than one category. For example, “Prevent users from connecting to machines” is described in the machines category, but it also affects users. So, if you can’t find a task in one category, check a related category.

Other articles also contain related information:

- [Applications](#) contains information about managing applications in delivery groups.
- Managing delivery groups requires the Delivery Group Administrator built-in role permissions. For details, see [Delegated administration](#).

### General

- Change the delivery type
- Change StoreFront addresses
- Upgrade a delivery group
- Manage Remote PC Access delivery groups

### Change the delivery type of a delivery group

The delivery type indicates what the group can deliver: applications, desktops, or both.

Before changing an **application only** or **desktops and applications** type to the **desktops only** type, delete all applications from the group.

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **Delivery Type** page, select the delivery type you want.
4. Select **Apply** to apply any changes you made and keep the window open. Or, select **OK** to apply changes and close the window.

### Change StoreFront addresses

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **StoreFront** page, select or add StoreFront URLs that are used by the Citrix Workspace app, which is installed on each machine in the delivery group.
4. Select **Apply** to apply any changes you made and keep the window open. Or, select **OK** to apply changes and close the window.

You can also specify StoreFront server addresses by selecting **StoreFront** in the left pane.

### Upgrade a delivery group or revert an upgrade

Upgrade a delivery group after you upgrade the VDAs on its machines and the machine catalogs containing the machines used in the delivery group.

Before you start the delivery group upgrade:

- If you use Citrix Provisioning (formerly Provisioning Services), upgrade the VDA version in the Citrix Provisioning console.
- Start the machines containing the upgraded VDA so that they can register with a Delivery Controller. This process tells the console about what needs upgrading in the delivery group.
- If you must continue to use earlier VDA versions, newer product features might not be available. For more information, see the upgrade documentation.

To upgrade a delivery group:

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Upgrade Delivery Group** in the action bar. The **Upgrade Delivery Group** action appears only if upgraded VDAs are detected.

The display indicates you which, if any, machines cannot be upgraded and why. You can then cancel the upgrade, resolve the machine issues, and then start the upgrade again.

After the upgrade completes, you can revert the machines to their previous states by selecting the delivery group and then selecting **Undo** in the action bar.

### Manage Remote PC Access delivery groups

If a machine in a Remote PC Access machine catalog is not assigned to a user, the machine is temporarily assigned to a delivery group associated with that catalog. This temporary assignment enables the machine to be assigned to a user later.

The delivery group-to-machine catalog association has a priority value. Priority determines which delivery group that machine is assigned to when it registers with the system or when a user needs a

machine assignment: the lower the value, the higher the priority. If a Remote PC Access machine catalog has multiple delivery group assignments, the software selects the match with the highest priority. Use the PowerShell SDK to set this priority value.

When first created, Remote PC Access machine catalogs are associated with a delivery group. This means that machine accounts or Organizational Units added to the catalog later can be added to the delivery group. This association can be switched off or on.

To add or remove a Remote PC Access machine catalog association with a delivery group:

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a Remote PC Access group.
3. In the **Details** section, select the **Machine Catalogs** tab and then select a Remote PC Access catalog.
4. To add or restore an association, select **Add Desktops**. To remove an association, select **Remove Association**.

## Users

- Change user settings
- Add or remove users

### Change user settings in a delivery group

The name of this page appears as either **User Settings** or **Basic Settings**.

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **User Settings** (or **Basic Settings**) page, change any of the settings in the following table.
4. Select **Apply** to apply any changes you made and keep the window open. Or, select **OK** to apply changes and close the window.

---

Setting	Description
Description	The text that Citrix Workspace (or StoreFront) uses and that users see.
Enable Delivery Group	Whether the delivery group is enabled.
Time zone	



Setting	Description
Enable Secure ICA	Secures communications to and from machines in the delivery group using SecureICA, which encrypts the ICA protocol. The default level is 128-bit. The level can be changed using the SDK. Citrix recommends using more encryption methods such as TLS encryption when traversing public networks. Also, SecureICA does not check data integrity.

### Add or remove users in a delivery group

For detailed information about users, see [Users](#).

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **Users** page:
  - To add users, select **Add**, and then specify the users you want to add.
  - To remove users, select one or more users and then select **Remove**.
  - Select or clear the check box to allow access by unauthenticated users.
4. Select **Apply** to apply any changes you made and keep the window open. Or, select **OK** to apply changes and close the window.

### Import or export user lists

For delivery groups containing physical single-session OS machines, you can import user information from a .csv file after you create the delivery group. You can also export user information to a .csv file. The .csv file can contain data from a previous product version.

The first line in the .csv file must contain comma-separated column headings (in any order), which can include: `ADComputerAccount`, `AssignedUser`, `VirtualMachine`, and `HostId`. Subsequent lines in the file contain comma-separated data. The `ADComputerAccount` entries can be common names, IP addresses, distinguished names, or domain and computer name pairs.

To import or export user information:

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **Machine Allocation** page, select **Import** list or **Export** list, and then browse to the file location.

4. Select **Apply** to apply any changes you made and keep the window open. Or, select **OK** to apply changes and close the window.

## Machines

- Change assignments of machines to users
- Change the maximum number of machines per user
- Update a machine
- Add, change, or remove a tag restriction for a desktop
- Remove a machine
- Restrict access to machines
- Prevent users from connecting to a machine (maintenance mode)
- Shut down and restart machines
- Create and manage restart schedules for machines
- Load manage machines
- Manage Autoscale

In addition to the features described in this article, see [Autoscale](#) for information about proactively power managing machines.

### Change assignments of machines to users in a delivery group

You can change the assignments of single-session OS machines provisioned with MCS. You cannot change assignments for multi-session OS machines or machines provisioned with Citrix Provisioning.

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **Desktops** or **Desktop Assignment Rules** page (the page title depends on the type of machine catalog the delivery group uses), specify the new users.
4. Select **Apply** to apply any changes you made and keep the window open. Or, select **OK** to apply changes and close the window.

### Change the maximum number of machines per user in a delivery group

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **Desktop Assignment Rules** page, set the maximum desktops per user value.
4. Select **Apply** to apply any changes you made and keep the window open. Or, select **OK** to apply changes and close the window.

### Update a machine in a delivery group

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **View Machines** in the action bar.
3. Select a machine and then select **Update Machines** in the action bar.

To choose a different image, select **Master image** and then select a snapshot.

To apply changes and notify machine users, select **Rollout notification to end-users**. Then specify:

- When to update the image: now or on the next restart
- The restart distribution time (the total time to begin updating all machines in the group)
- Whether users are notified of the restart
- The message users will receive

### Add, change, or remove a tag restriction for a desktop

Adding, changing, and removing tag restrictions can have unanticipated effects on which desktops are considered for launch. Review the considerations and cautions in [Tags](#).

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **Desktops** page, select the desktop and select **Edit**.
4. To add a tag restriction, select **Restrict launches to machines with the tag** and then select the tag.
5. To change or remove a tag restriction, either:
  - Select a different tag.
  - Remove the tag restriction by clearing **Restrict launches to machines with this tag**.
6. Select **Apply** to apply any changes you made and keep the window open. Or, select **OK** to apply changes and close the window.

### Remove a machine from a delivery group

Removing a machine deletes it from a delivery group. It does not delete it from the machine catalog that the delivery group uses. Therefore, that machine is available for assignment to another delivery group.

Machines must be shut down before they can be removed. To temporarily stop users from connecting to a machine while you are removing it, put the machine into maintenance mode before shutting it down.

Machines might contain personal data, so use caution before allocating the machine to another user. Consider reimaging the machine.

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **View Machines** in the action bar.
3. Ensure that the machine is shut down.
4. Select the machine and then select **Remove from Delivery Group** in the action bar.

You can also remove a machine from a delivery group through the [connection](#) the machine uses.

### Restrict access to machines in a delivery group

Any changes you make to restrict access to machines in a delivery group supersede previous settings, regardless of the method you use. You can:

- **Restrict access for administrators using delegated administration scopes:** You can create and assign a scope that permits administrators to access all applications, and another scope that provides access to only certain applications. For details, see [Delegated administration](#).
- **Restrict access for users through SmartAccess policy expressions:** Use policy expressions to filter user connections made through Citrix Gateway.
  1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
  2. Select a group and then select **Edit Delivery Group** in the action bar.
  3. On the **Access Policy** page, select **Connections through NetScaler Gateway**.
  4. To choose a subset of those connections, select **Connections meeting any of the following filters**. Then define the Citrix Gateway site, and add, edit, or remove the SmartAccess policy expressions for the allowed user access scenarios. For details, see the Citrix Gateway documentation.
  5. Select **Apply** to apply any changes you made and keep the window open. Or, select **OK** to apply changes and close the window.
- **Restrict access for users through exclusion filters:** Use exclusion filters on access policies that you set in the SDK. Access policies are applied to delivery groups to refine connections. For example, you can restrict machine access to a subset of users, and you can specify allowed user devices. Exclusion filters further refine access policies. For example, for security, you can deny access to a subset of users or devices. By default, exclusion filters are disabled.

For example, for a teaching lab on a corporate network subnet, to prevent access from that lab to a particular delivery group, regardless of who is using the machines in the lab, use the command: `Set-BrokerAccessPolicy -Name VPDesktops_Direct -ExcludedClientIPFilterEnabled $True -`

You can use the asterisk (\*) wildcard to match all tags that start with the same policy expression. For example, if you add the tag `VPDesktops_Direct` to one machine and `VPDesktops_Test` to another, setting the tag in the `Set-BrokerAccessPolicy` script to `VPDesktops_*` applies the filter to both machines.

If you are connected using a web browser or with the Citrix Workspace app user experience feature enabled in the store, you cannot use a client name exclusion filter.

### **Prevent users from connecting to a machine (maintenance mode) in a delivery group**

When you need to temporarily stop new connections to machines, you can turn on maintenance mode for one or all machines in a delivery group. You might do this before applying patches or using management tools.

- When a multi-session OS machine is in maintenance mode, users can connect to existing sessions, but cannot start new sessions.
- When a single-session OS machine (or a PC using Remote PC Access) is in maintenance mode, users cannot connect or reconnect. Current connections remain connected until they disconnect or log off.

To turn maintenance mode on or off:

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group.
3. To turn on maintenance mode for all machines in the delivery group, select **Turn On Maintenance Mode** in the action bar.

To turn on maintenance mode for one machine, select **View Machines** in the action bar. Select a machine, and then select **Turn On Maintenance Mode** in the action bar.

4. To turn maintenance mode off for one or all machines in a delivery group, follow the previous instructions, but select **Turn Off Maintenance Mode** in the action bar.

Windows Remote Desktop Connection (RDC) settings also affect whether a multi-session OS machine is in maintenance mode. Maintenance mode is on when any of the following occur:

- Maintenance mode is set to on, as described earlier.
- RDC is set to **Don't allow connections to this computer**.
- RDC is not set to **Don't allow connections to this computer** and the Remote Host Configuration User Logon Mode setting is either **Allow reconnections, but prevent new logons** or **Allow reconnections, but prevent new logons until the server is restarted**.

You can also turn maintenance mode on or off for:

- A connection, which affects the machines using that connection.
- A machine catalog, which affects the machines in that catalog.

### **Shut down and restart machines in a delivery group**

This procedure is not supported for Remote PC Access machines.

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **View Machines** in the action bar.
3. Select the machine and then select one of the following entries in the action bar (some options might not be available, depending on the machine state):
  - **Force shut down:** Forcibly powers off the machine and refreshes the list of machines.
  - **Restart:** Requests the operating system to shut down and then start the machine again. If the operating system cannot comply, the machine remains in its current state.
  - **Force restart:** Forcibly shuts down the operating system and then restarts the machine.
  - **Suspend:** Pauses the machine without shutting it down, and refreshes the list of machines.
  - **Shut down:** Requests the operating system to shut down.

For non-force actions, if the machine does not shut down within 10 minutes, it is powered off. If Windows attempts to install updates during the shutdown, there is a risk that the machine will be powered off before the updates finish.

Citrix recommends that you prevent single-session OS machine users from selecting **Shut down** within a session. See the Microsoft policy documentation for details.

You can also shut down and restart machines on a [connection](#).

### **Create and manage restart schedules for machines in a delivery group**

A restart schedule specifies when machines in a delivery group are periodically restarted. You can create one or more schedules for a delivery group. A schedule can affect either:

- All the machines in the group.
- One or more (but not all) machines in the group. The machines are identified by a tag that you apply to the machine. This is called a tag restriction, because the tag restricts an action to only items (in this case, machines) that have the tag.

For example, let's say all of your machines are in one delivery group. You want every machine restarted once every week, and you want the machines used by the accounting team restarted daily. To accomplish this, set up one schedule for all machines, and another schedule for only the machines in accounting.

A schedule includes the day and time the restart begins, and the duration. The duration is either "start all affected machines at the same time" or an interval it will likely take to restart all affected machines.

You can enable or disable a schedule. Disabling a schedule can be helpful when testing, during special intervals, or when preparing schedules before you need them.

You cannot use schedules for automated power-on or shutdown from the management console, only to restart.

## Schedule overlap

Multiple schedules can overlap. In the example above, both schedules affect the accounting machines. Those machines might be restarted twice on Sunday. The scheduling code is designed to avoid restarting the same machine more often than intended, but it cannot be guaranteed.

- If the schedules coincide precisely in start and duration times, it is more likely that the machines will be restarted only once.
- The more the schedules differ in start and duration times, it's more likely that multiple restarts will occur.
- The number of machines affected by a schedule also affects the chance of an overlap. In the example, the weekly schedule that affects all machines might initiate restarts significantly faster than the daily schedule for accounting machines, depending on the duration specified for each.

For an in-depth look at restart schedules, see [Reboot schedule internals](#).

## View restart schedules

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. Select the **Restart Schedule** page.

The **Restart Schedule** page contains the following information for each configured schedule:

- Schedule name.
- Tag restriction used, if any.
- How often the machine restarts occur.
- Whether machine users receive a notification.
- Whether the schedule is enabled. Disabling a schedule can be helpful when testing, during special intervals, or when preparing schedules before you need them.

## Add (apply) tags

When you configure a restart schedule that uses a tag restriction, ensure that the tag has been added (applied) to the machines that the schedule affects. In the example above, each of the machines used by the accounting team has a tag applied. For details, see [Tags](#).

Although you can apply more than one tag to a machine, a restart schedule can specify only one tag.

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select the group containing the machines to be controlled by the schedule.
3. Select **View Machines** and then select the machines you want to add a tag to.
4. Select **Manage Tags** in the action bar.

5. If the tag exists, enable the check box next to the tag name. If the tag does not exist, select **Create** and then specify the name for the tag. After the tag is created, enable the check box next to the newly created tag name.
6. Select **Save** in the **Manage Tags** dialog.

### Create a restart schedule

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **Restart Schedule** page, select **Add**.
4. On the **Add Restart Schedule** page:
  - To enable the schedule, select the check box. To disable the schedule, clear the check box.
  - Type a schedule name and description.
  - For **Restart frequency**, select how often the restart occurs: daily, weekdays, weekend days, or a specific day each week.
  - For **Begin restart at**, specify the time of day to begin the restart in a 24-hour clock format.
  - For **Restart duration**, choose whether to:
    - Restart all machines at the same time.
    - Begin restarting all affected machines within a certain interval. An internal algorithm determines when each machine is restarted during that interval.
    - Restart all machines after draining all sessions. When the restart time is reached, machines are put into the drain state and restarted when all sessions are logged off.
  - In **Send notification to users**, choose whether to display a notification message on the affected machines before a restart begins. By default, no message is displayed.
  - If you choose to display a message 15 minutes before the restart begins, you can choose (in **Notification frequency**) to repeat the message every five minutes after the initial message. By default, the message is not repeated.
  - Enter the notification title and text. There is no default text.

If you want the message to include the number of minutes before restart, include the variable **%m%**. For example: “Warning: Your computer is automatically restarted in %m% minutes.” The value decrements by five minutes in each repeated message. Unless you chose to restart all machines at the same time, the message displays on each machine at the appropriate time before the restart, calculated by the internal algorithm.
  - Advanced settings:
    - **Restrict to tag**. Lets you apply a tag restriction.



- **Include machines in maintenance mode.** Lets you include machines that are in maintenance mode in this restart schedule. This setting is available only in the *web-based console*. To use PowerShell instead, see Scheduled restarts for machines in maintenance mode.
5. Select **Apply** to apply the changes you made and keep the window open. Or, select **OK** to apply changes and close the window.

### Immediately run a restart schedule

**Note:**

This feature is available only in the *web-based console*.

A restart schedule specifies when machines in a delivery group restart regularly. You can also run a restart schedule immediately to restart the machines in that schedule.

To run a restart schedule immediately, follow these steps:

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select the applicable delivery group and then select **Edit Delivery Group** in the action bar. The **Edit Delivery Group** wizard appears.
3. Select the **Restart Schedule** page.
4. Select a schedule that you want to run and then select **Run schedule now**.

**Note:**

- You cannot run a schedule immediately if it is configured with the **Restart all machines after draining sessions** setting.
- You can apply **Run schedule now** only to one schedule at a time.
- After you edit a schedule, **Run schedule now** becomes unavailable. Select **Apply** to make it available.

### Edit, remove, enable, or disable a restart schedule

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **Restart Schedule** page, select the check box for a schedule.
  - To edit a schedule, select **Edit**. Update the schedule configuration, using the guidance in Create a restart schedule.
  - To enable or disable a schedule, select **Edit**. Select or clear the **Enable restart schedule** check box.
  - To remove a schedule, select **Remove**. Confirm the removal. Removing a schedule does not affect any tags applied to machines in the affected machines.

## Scheduled restarts delayed due to database outage

### Note:

This feature is available only in PowerShell.

If a site database outage occurs before a scheduled restart begins for machines (VDAs) in a delivery group, the restarts begin when the outage ends. This can have unintended results.

For example, let's say you've scheduled a delivery group's restarts to occur during off-production hours (beginning at 3 am). A site database outage occurs one hour before a scheduled restart begins (2 am). The outage lasts six hours (until 8 am). The restart schedule begins when the connection between the Delivery Controller and the site database is restored. The VDA restarts now begin five hours after their original schedule. This might result in VDAs restarting during production hours.

To help avoid this situation, you can use the `MaxOvertimeStartMins` parameter for the `New-BrokerRebootScheduleV2` and `Set-BrokerRebootScheduleV2` cmdlets. The value specifies the maximum number of minutes beyond the scheduled start time that a restart schedule can begin.

- If the database connection is restored within that time (scheduled time + `MaxOvertimeStartMins`), the VDA restarts begin.
- If the database connection is not restored within that time, the VDA restarts do not begin.
- If this parameter is omitted or has a zero value, the scheduled restart begins when the connection to the database is restored, regardless of the outage duration.

For more information, see the cmdlet help. This feature is available only in PowerShell.

## Scheduled restarts for machines in maintenance mode

To indicate whether a restart schedule affects machines that are in maintenance mode, use the `IgnoreMaintenanceMode` option with the `BrokerRebootScheduleV2` cmdlets.

For example, the following cmdlet creates a schedule that restarts both machines that are and machines that are not in maintenance mode.

```
New-BrokerRebootScheduleV2 rebootSchedule1 -DesktopGroupName <myDesktopGroup> -IgnoreMaintenanceMode $true
```

The following cmdlet modifies an existing restart schedule.

```
Set-BrokerRebootScheduleV2 rebootSchedule1 -IgnoreMaintenanceMode $true
```

For more information, see the cmdlet help.

## Load manage machines in delivery groups

You can load manage multi-session OS machines only.

Load management measures the server load and determines which server to select under the current environment conditions. This selection is based on:

- **Server maintenance mode status:** A multi-session OS machine is considered for load balancing only when maintenance mode is off.
- **Server load index:** Determines how likely a server delivering multi-session OS machines is to receive connections. The index is a combination of load evaluators: the number of sessions and the settings for performance metrics such as CPU, disk, and memory use. Load evaluators are specified in load management policy settings.

A server load index of 10000 indicates that the server is fully loaded. If no other servers are available, users might receive a message that the desktop or application is currently unavailable when they launch a session.

You can monitor the load index in Director (Monitor), a Full Configuration management interface search, and the SDK.

In console displays, to display the **Server Load Index** column (which is hidden by default), select a machine, right-click a column heading, and then select **Select Column**. In the **Machine category**, select **Load Index**.

In the SDK, use the `Get-BrokerMachine` cmdlet. For details, see [CTX202150](#).

- **Concurrent logon tolerance policy setting:** The maximum number of concurrent requests to log on to the server. (This setting is equivalent to load throttling in XenApp 6.x versions.)

When all servers are at or higher than the concurrent logon tolerance setting, the next logon request is assigned to the server with the lowest pending logons. If more than one server meets these criteria, the server with the lowest load index is selected.

## Manage Autoscale

### Important:

This feature is available only in the web-based console.

By default, Autoscale is disabled for delivery groups. To manage Autoscale for a delivery group (if applicable), follow these steps:

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Manage Autoscale** in the action bar. The **Manage Autoscale** window appears.
3. Configure settings as needed. For information about Autoscale settings, see [Autoscale](#).
4. Select **Apply** to apply any changes you made and to keep the window open. Or, select **Save** to apply changes and to close the window.

## Sessions

- Log off or disconnect a session, or send a message to users
- Configure session prelaunch and session linger
- Control session reconnection when disconnected from machine in maintenance mode

### Log off or disconnect a session, or send a message to delivery group users

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **View Machines** in the action bar.
3. To log a user off a session, select the session or desktop and then select **Log off** in the action bar. The session closes and the machine becomes available to other users, unless it is allocated to a specific user.
4. To disconnect a session, select the session or desktop and then select **Disconnect** in the action bar. Applications continue to run and the machine remains allocated to that user. The user can reconnect to the same machine.
5. To send a message to users, select the session, machine, or user and then select **Send message** in the action bar. Enter the message.

### Configure session prelaunch and session linger in a delivery group

These features are supported only on multi-session OS machines.

The session prelaunch and session linger features help specified users access applications quickly, by:

- Starting sessions before they are requested (session prelaunch)
- Keeping application sessions active after a user closes all applications (session linger)

By default, session prelaunch and session linger are not used. A session starts (launches) when a user starts an application, and remains active until the last open application in the session closes.

Considerations:

- The delivery group must support applications, and the machines must be running a VDA for multi-session OS, minimum version 7.6.
- These features are supported only when using Citrix Workspace app for Windows, and also require more Citrix Workspace app configuration. For instructions, search for session prelaunch in the product documentation for your Citrix Workspace app for Windows version.
- Citrix Workspace app for HTML5 is not supported.
- When using session prelaunch, if a user's machine is put into suspend or hibernate mode, prelaunch does not work (regardless of session prelaunch settings). Users can lock their machines/sessions. However, if a user logs off from Citrix Workspace app, the session is ended and prelaunch no longer applies.

- When using session prelaunch, physical client machines cannot use the suspend or hibernate power management functions. Client machine users can lock their sessions but should not log off.
- Prelaunched and lingering sessions consume a concurrent license, but only when connected. If using a user/device license, the license lasts 90 days. Unused prelaunched and lingering sessions disconnect after 15 minutes by default. This value can be configured in PowerShell ([New/Set-BrokerSessionPreLaunch](#) cmdlet).
- Careful planning and monitoring of your users' activity patterns are essential to tailoring these features to complement each other. Optimal configuration balances the benefits of earlier application availability for users against the cost of keeping licenses in use and resources allocated.
- You can also configure session prelaunch for a scheduled time of day in Citrix Workspace app.

### How long unused prelaunched and lingering sessions remain active

There are several ways to specify how long an unused session remains active if the user does not start an application: a configured timeout and server load thresholds. You can configure all of them. The event that occurs first causes the unused session to end.

- **Timeout:** A configured timeout specifies the number of minutes, hours, or days an unused prelaunched or lingering session remains active. If you configure too short a timeout, prelaunched sessions end before they provide the user benefit of quicker application access. If you configure too long a timeout, incoming user connections might be denied because the server doesn't have enough resources.

You can enable this timeout from the SDK only ([New/Set-BrokerSessionPreLaunch](#) cmdlet), not from the management console. If you disable the timeout, it does not appear in the console display for that delivery group or in the **Edit Delivery Group** pages.

- **Thresholds:** Automatically ending prelaunched and lingering sessions based on server load ensures that sessions remain open as long as possible, assuming that server resources are available. Unused prelaunched and lingering sessions do not cause denied connections because they are ended automatically when resources are needed for new user sessions.

You can configure two thresholds: the average percentage load of all servers in the delivery group, and the maximum percentage load of a single server in the group. When a threshold is exceeded, the sessions that have been in the prelaunch or lingering state for the longest time are ended. Sessions are ended one-by-one at minute intervals until the load falls below the threshold. While the threshold is exceeded, no new prelaunch sessions are started.

Servers with VDAs that have not registered with a Controller and servers in maintenance mode are considered fully loaded. An unplanned outage causes prelaunch and lingering sessions to end automatically to free capacity.

## To enable session prelaunch

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **Application Prelaunch** page, enable session prelaunch by choosing when sessions launch:
  - When a user starts an application. This is the default setting. Session prelaunch is disabled.
  - When any user in the delivery group logs on to Citrix Workspace app for Windows.
  - When anyone in a list of users and user groups logs on to Citrix Workspace app for Windows. Be sure to also specify users or user groups if you choose this option.

The screenshot shows the 'Edit Delivery Group' configuration page. On the left is a navigation pane with 'Application Prelaunch' selected. The main area is titled 'Prelaunch Sessions for Applications'. It includes a description: 'With prelaunch, sessions launch when users log on to Receiver, so applications are available sooner.' Below this, it asks 'When do you want sessions to launch?' with three radio button options: 'Launch when users start an application (no prelaunch)', 'Prelaunch when any user in the Delivery Group logs on to Receiver for Windows' (which is selected), and 'Prelaunch when any of the following users log on to Receiver for Windows:'. There is an empty list box with 'Add...' and 'Remove' buttons. Below that, it asks 'If no application is started, when do you want prelaunched sessions to end?' with three options: 'After a specified time:' (set to 2 hours), 'When average load on all machines exceeds (%):' (checked, set to 80%), and 'When load on any machine exceeds (%):' (checked, set to 85%). At the bottom are 'OK', 'Cancel', and 'Apply' buttons.

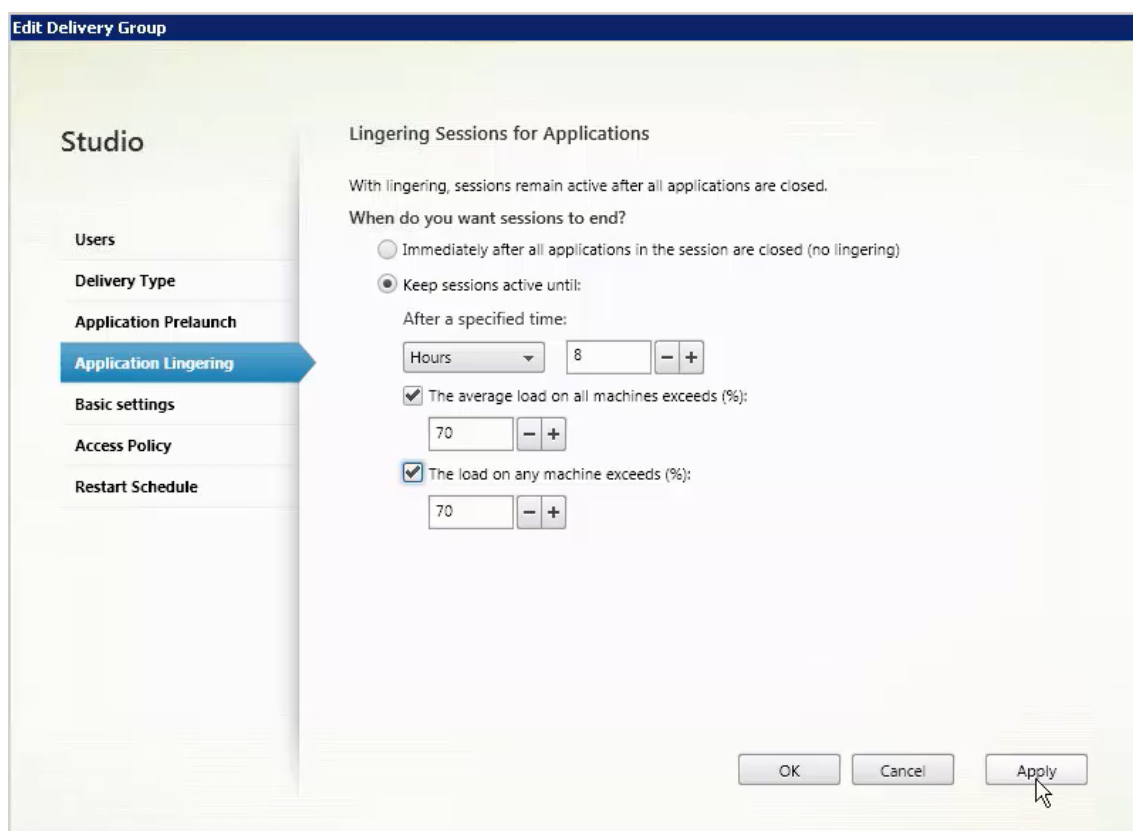
4. A prelaunched session is replaced with a regular session when the user starts an application. If the user does not start an application (the prelaunched session is unused), the following settings affect how long that session remains active.
  - When a specified time interval elapses. You can change the time interval (1–99 days, 1–2376 hours, or 1–142,560 minutes).
  - When the average load on all machines in the delivery group exceeds a specified percentage (1–99%).

- When the load on any machine in the delivery group exceeds a specified percentage (1–99%).

Recap: A prelaunched session remains active until one of the following events occurs: a user starts an application, the specified time elapses, or a specified load threshold is exceeded.

### To enable session linger

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select a group and then select **Edit Delivery Group** in the action bar.
3. On the **Application Linging** page, enable session linger by selecting **Keep sessions active until**.



4. Several settings affect how long a lingering session remains active if the user does not start another application.
  - When a specified time interval elapses. You can change the time interval: 1–99 days, 1–2376 hours, or 1–142,560 minutes.
  - When the average load on all machines in the delivery group exceeds a specified percentage: 1–99%.
  - When the load on any machine in the delivery group exceeds a specified percentage: 1–99%.

Recap: A lingering session remains active until one of the following events occurs: a user starts an application, the specified time elapses, or a specified load threshold is exceeded.

### Control session reconnection when disconnected from machine in maintenance mode

**Note:**

This feature is available only in PowerShell.

You can control whether sessions that are disconnected on machines in maintenance mode are allowed to reconnect to machines in the delivery group.

Before late May 2021, reconnection was not allowed for single-session pooled desktop sessions that had disconnected from machines in maintenance mode. Reconnection was always allowed for static single-session desktops and multi-session machines. Now, you can configure a delivery group to allow or prohibit reconnections (regardless of VDA type) after disconnection from a machine in maintenance mode.

When creating or editing a delivery group (`New-BrokerDesktopGroup`, `Set-BrokerDesktopGroup`), use the `-AllowReconnectInMaintenanceMode <boolean>` parameter to allow or prohibit reconnections for machines that were disconnected from a machine in maintenance mode.

- When set to true, sessions can reconnect to machines in the group.
- When set to false, sessions cannot reconnect to machines in the group.

Default values:

- Single-session: Disabled
- Multi-session: Enabled

### Troubleshoot

- VDAs that are not registered with a Delivery Controller are not considered when launching brokered sessions. This results in underutilization of otherwise available resources. There are various reasons a VDA might not be registered, many of which an administrator can troubleshoot. The details display provides troubleshooting information in the catalog creation wizard, and after you add a catalog to a delivery group.

After you create a delivery group, the details pane for a delivery group indicates the number of machines that should be registered but are not. For example, one or more machines are powered on and not in maintenance mode, but are not currently registered with a Controller. When viewing a “not registered, but should be” machine, review the **Troubleshoot** tab in the details pane for possible causes and recommended corrective actions.

For messages about functional level, see [VDA versions and functional levels](#).

For information about VDA registration troubleshooting, see [CTX136668](#).



- In the display for a delivery group, the **Installed VDA version** in the details pane might differ from the actual version installed on the machines. The machine's Windows Programs and Features display shows the actual VDA version.
- For machines with **Power State Unknown** status, see [CTX131267](#) for guidance.

## Create application groups

July 7, 2021

### Introduction

Application groups let you manage collections of applications. You can create application groups for applications shared across different delivery groups or used by a subset of users within delivery groups. Application groups are optional. They offer an alternative to adding the same applications to multiple delivery groups. Delivery groups can be associated with more than one application group, and an application group can be associated with more than one delivery group.


Using application groups can provide application management and resource control advantages over using more delivery groups:

- The logical grouping of applications and their settings lets you manage those applications as a single unit. For example, you don't have to add (publish) the same application to individual delivery groups one at a time.
- Session sharing between application groups can conserve resource consumption. In other cases, disabling session sharing between application groups may be beneficial.
- You can use the tag restriction feature to publish applications from an application group, considering only a subset of the machines in selected delivery groups. With tag restrictions, you can use your existing machines for more than one publishing task, saving the costs associated with deploying and managing additional machines. A tag restriction can be thought of as subdividing (or partitioning) the machines in a delivery group. Using an application group or desktops with a tag restriction can be helpful when isolating and troubleshooting a subset of machines in a delivery group.

### Example configurations

#### Example 1

The following graphic shows a deployment that includes application groups:




In this configuration, applications are added to the application groups, not the delivery groups. The delivery groups specify which machines will be used. (Although not shown, the machines are in machine catalogs.)

Application group 1 is associated with delivery group 1. The applications in application group 1 can be accessed by the users specified in application group 1, as long as they are also in the user list for delivery group 1. This follows the guidance that the user list for an application group should be a subset (a restriction) of the user lists for the associated delivery groups. The settings in application group 1 (such as application session sharing between application groups, associated delivery groups) apply to applications and users in that group. The settings in delivery group 1 (such as anonymous user support) apply to users in application groups 1 and 2, because those application groups have been associated with that delivery group.

Application group 2 is associated with two delivery groups: 1 and 2. Each of those delivery groups can be assigned a priority in application group 2, which indicates the order in which the delivery groups will be checked when an application is launched. delivery groups with equal priority are load balanced. The applications in application group 2 can be accessed by the users specified in application group 2, as long as they are also in the user lists for delivery group 1 and delivery group 2.

## Example 2

This simple layout uses tag restrictions to limit which machines will be considered for certain desktop and application launches. The site has one shared delivery group, one published desktop, and one application group configured with two applications.



Tags have been added to each of the three machines (VDA 101-103).

The application group was created with the “Orange” tag restriction, so each of its applications (Calculator and Notepad) can be launched only on machines in that delivery group that have the tag “Orange”: VDA 102 and 103.

For more comprehensive examples and guidance for using tag restrictions in application groups (and for desktops), see [Tags](#).

## Guidance and considerations

Citrix recommends adding applications to either application groups or delivery groups, but not both. Otherwise, the additional complexity of having applications in two group types can make it more difficult to manage.

By default, an application group is enabled. After you create an application group, you can edit the group to change this setting. See [Manage application groups](#).

By default, application session sharing between application groups is enabled. See [Session sharing between application groups](#).

Citrix recommends upgrading your delivery groups to the current version. This requires:

1. Upgrading VDAs on the machines used in the delivery group.
2. Upgrading the machine catalogs containing those machines
3. Upgrading the delivery group.

For details, see [Manage delivery groups](#).

To use application groups, your core components must be minimum version 7.9.

Creating application groups requires the delegated administration permission of the Delivery Group Administrator built-in role. See [Delegated administration](#) for details.

This article refers to “associating” an application with more than one application group to differentiate that action from adding a new instance of that application from an available source. Similarly, delivery groups are associated with application groups (and vice versa), rather than being additions or components of one another.

### **Session sharing with application groups**

When application session sharing is enabled, all applications launch in the same application session. This saves the costs associated with launching additional application sessions, and allows the use of application features that involve the clipboard, such as copy-paste operations. However, in some situations you may wish to turn off session sharing.

When you use application groups you can configure application session sharing in the following three ways which extend the standard session sharing behavior available when you are using only delivery groups:

- Session sharing enabled between application groups.
- Session sharing enabled only between applications in the same application group.
- Session sharing disabled.

### **Session sharing between application groups**

You can enable application session sharing between application groups, or you can disable it to limit application session sharing only to applications in the same application group.

- **An example when enabling session sharing between application groups is helpful:**

Application group 1 contains Microsoft Office applications such as Word and Excel. Application group 2 contains other applications such as Notepad and Calculator, and both application groups are attached to the same delivery group. A user who has access to both application groups starts an application session by launching Word, and then launches Notepad. If the user’s existing session running Word is suitable for running Notepad then Notepad is started within the existing session. If Notepad cannot be run from the existing session—for example if the tag restriction excludes the machine that the session is running on—then a new session on a suitable machine is created rather than using session sharing.

- **An example when disabling session sharing between application groups is helpful:**

You have a set of applications that do not interoperate well with other applications that are installed on the same machines, such as two different versions of the same software suite or

two different versions of the same web browser. You prefer not to allow a user to launch both versions in the same session.

You create an application group for each version of the software suite, and add the applications for each version of the software suite to the corresponding application group. If session sharing between groups is disabled for each of those application groups, a user specified in those groups can run applications of the same version in the same session, and can still run other applications at the same time, but not in the same session. If the user launches one of the different-versioned applications (that are in a different application group), or launches any application that is not contained in an application group, then that application is launched in a new session.

This session sharing between application groups feature is not a security sandboxing feature. It is not foolproof, and it cannot prevent users from launching applications into their sessions through other means (for example, through Windows Explorer).

If a machine is at capacity, new sessions are not started on it. New applications are started in existing sessions on the machine as needed using session sharing (providing that this complies with the session sharing restrictions described here).

You can only make prelaunched sessions available to application groups which have application session sharing allowed. (Sessions which use the session linger feature are available to all application groups.) These features must be enabled and configured in each of the delivery groups associated with the application group. You cannot configure them in the application groups.

By default, application session sharing between application groups is enabled when you create an application group. You cannot change this when you create the group. After you create an application group, you can edit the group to change this setting. See [Manage application groups](#).

### **Disable session sharing within an application group**

You can prevent application session sharing between applications which are in the same application group.

- **An example when disabling session sharing within application groups is helpful:**

You want your users to access multiple simultaneous full screen sessions of an application on separate monitors.

You create an application group and add the applications to it. If session sharing is prohibited between applications in that application group, when a user specified in it starts one application after another they launch in separate sessions, and the user can move each to a separate monitor.

By default, application session sharing is enabled when you create an application group. You cannot

change this when you create the group. After you create an application group, you can edit the group to change this setting. See [Manage application groups](#).

## Create an application group

To create an application group:

1. From **Manage > Full Configuration**, select **Applications** in the left pane, and then select the **Application Groups** tab.
2. Select **Create Application Group**.
3. The group creation wizard launches with an **Introduction** page, which you can remove from future launches of this wizard.
4. The wizard guides you through the pages described below. When you are done with each page, select **Next** until you reach the **Summary** page.

### Step 1. Delivery groups

The **Delivery Groups** page lists all delivery groups, with the number of machines each group contains.

- The **Compatible Delivery Groups** list contains delivery groups you can select. Compatible delivery groups contain random (not permanently or statically assigned) server or desktop OS machines.
- The **Incompatible Delivery Groups** list contains delivery groups you cannot select. Each entry explains why it is not compatible, such as containing static assigned machines.

An application group can be associated with delivery groups containing shared (not private) machines that can deliver applications.

You can also select delivery groups containing shared machines that deliver only desktops, if both of the following conditions are met:

- The delivery group contains shared machines and was created with a XenDesktop version earlier than 7.9.
- You have Edit Delivery Group permission.

The delivery group type is automatically converted to “desktops and applications” when the group creation wizard is committed.

Although you can create an application group that has no associated delivery groups (perhaps to organize applications or to serve as storage for applications not currently used) the application group cannot be used to deliver applications until it specifies at least one delivery group. Additionally, you cannot add applications to the application group from the **From Start** menu source if there are no delivery groups specified.

The delivery groups you select specify the machines that will be used to deliver applications. Select the check boxes next to the delivery groups you want to associate with the application group.

To add a tag restriction, select **Restrict launches to machines with the tag** and then select the tag from the dropdown.

## Step 2. Users

Specify who can use the applications in the application group. You can either allow all users and user groups in the delivery groups you selected on the previous page, or select specific users and user groups from those delivery groups. If you restrict use to users you specify, then only the users specified in the delivery group and the application group can access the applications in this application group. Essentially, the user list in the application group provides a filter on the user lists in the delivery groups.

Enabling or disabling application use by unauthenticated users is available only in delivery groups, not in application groups.

For information about where user lists are specified in a deployment, see [Where user lists are specified](#).

## Step 3. Applications

Good to know:

- By default, new applications you add are placed in a folder named **Applications**. You can specify a different folder. If you try to add an application and one with the same name already exists in that folder, you are prompted to rename the application you are adding. If you agree with the suggested unique name, the application is added with that new name. Otherwise, you must rename it yourself before it can be added. For details, see [Manage application folders](#).
- You can change an application's properties (settings) when you add it, or later. See [Change application properties](#). If you publish two applications with the same name to the same users, change the **Application name (for user)** property in Full Configuration management interface. Otherwise, users will see duplicate names in the Citrix Workspace app.
- When you add an application to more than one application group, a visibility issue can occur if you do not have sufficient permission to view the application in all of those groups. In such cases, either consult an administrator with greater permissions or have your scope extended to include all the groups to which the application was added.

Select the **Add** dropdown to display the application sources.

- **From Start menu:** Applications that are discovered on a machine in the selected delivery groups. When you select this source, a new page launches with a list of discovered applications. Select the check boxes of applications to add, and then select **OK**.

This source cannot be selected if you selected any of the following:

- Application groups that have no associated delivery groups.
  - Application groups with associated delivery groups that contain no machines.
  - A delivery group containing no machines.
- **Manually defined:** Applications located in the site or elsewhere in your network. When you select this source, a new page launches where you type the path to the executable, working directory, optional command line arguments, and display names for administrators and users. After entering this information, select **OK**.
  - **Existing:** Applications previously added to the site. When you select this source, a new page launches with a list of discovered applications. Select the check boxes of applications to add and then select **OK**. This source cannot be selected if the site has no applications.
  - **App-V:** Applications in App-V packages. When you select this source, a new page launches where you select **App-V server** or **Application Library**. From the resulting display, select the check boxes of applications to add, and then select **OK**. For more information, see [App-V](#). This source cannot be selected (or might not appear) if App-V is not configured for the site.

**Note:**

On VDA version 2003 and later, publishing App-V packages from HTTP URLs is not supported. You cannot select those applications from the list.

As noted, certain entries in the **Add** dropdown will not be selectable if there is no valid source of that type. Sources that are incompatible are not listed at all (for example, you cannot add application groups to application groups, so that source is not listed when you create an application group).

#### **Step 4. Scopes**

This page appears only if you have previously created a custom scope. By default, the **All** scope is selected. For more information, see [Delegated administration](#).

#### **Step 5. Summary**

Enter a name for the application group. You can also (optionally) enter a description.

Review the summary information and then select **Finish**.

## **Manage application groups**

July 7, 2021



## Introduction

This article describes how to manage the application groups you [created](#).

See [Applications](#) for information about managing applications in application groups or delivery groups, including how to:

- Add or remove applications in an application group.
- Change application group associations.

Managing application groups requires the delegated administration permissions of the Delivery Group Administrator built-in role. For details, see [Delegated administration](#).

## Enable or disable an application group

When an application group is enabled, it can deliver the applications that have been added to it. Disabling an application group disables each application in that group. However, if those applications are also associated with other enabled application groups, they can be delivered from those other groups. Similarly, if the application was explicitly added to delivery groups associated with the application group (in addition to being added to the application group), disabling the application group does not affect the applications in those delivery groups.

An application group is enabled when you create it. You cannot change this when you create the group.

1. From **Manage > Full Configuration**, select **Applications** in the left pane, and then select the **Application Groups** tab.
2. Select an application group and then select **Edit Application Group** in the action bar.
3. On the **Settings** page, select or clear the **Enable Application Group** check box.
4. Select **Apply** to apply any changes you made and keep the window open, or select **OK** to apply changes and close the window.

## Enable or disable application session sharing between application groups

Session sharing between application groups is enabled when you create an application group. You cannot change this when you create the group. For more information, see [Session sharing with application groups](#).

1. From **Manage > Full Configuration**, select **Applications** in the left pane, and then select the **Application Groups** tab.
2. Select an application group and then select **Edit Application Group** in the action bar.
3. On the **Settings** page, select or clear the **Enable application session sharing between Application Groups** check box.
4. Select **Apply** to apply any changes you made and keep the window open, or select **OK** to apply changes and close the window.

## Disable application session sharing within an application group

Session sharing between applications in the same application group is enabled by default when you create an application group. If you disable application session sharing between application groups, session sharing between applications in the same application group remains enabled.

You can use the PowerShell SDK to configure application groups with application session sharing disabled between the applications they contain. In some circumstances this can be desirable. For example, you might want users to start non-seamless applications in full-size application windows on separate monitors.

When you disable application session sharing within an application group, each application in that group launches in a new application session. If a suitable disconnected session is available which is running the same application, it is reconnected. For example, if you launch Notepad, and there is a disconnected session with Notepad running, that session is reconnected instead of creating a new one. If multiple suitable disconnected sessions are available, one of the sessions is chosen to reconnect to, in a random but deterministic manner. If the situation reoccurs in the same circumstances, the same session is chosen, but the session is not necessarily predictable otherwise.

You can use the PowerShell SDK to either disable application session sharing for all applications in an existing application group, or to create an application group with application session sharing disabled.

### PowerShell cmdlet examples

To disable session sharing, use the Broker PowerShell cmdlets `New-BrokerApplicationGroup` or `Set-BrokerApplicationGroup` with the parameter `SessionSharingEnabled` set to `False` and the parameter `SingleAppPerSession` set to `True`.

- For example, to create an application group with application session sharing disabled for all applications in the group:

```
New-BrokerApplicationGroup AppGr1 -SessionSharingEnabled $False -  
SingleAppPerSession $True
```

- For example, to disable application session sharing between all applications in an existing application group:

```
Set-BrokerApplicationGroup AppGR1 -SessionSharingEnabled $False -  
SingleAppPerSession $True
```

### Considerations

- To enable the `SingleAppPerSession` property you must set the `SessionSharingEnabled` property to `False`. The two properties must not be enabled at the same time. The

`SessionSharingEnabled` parameter refers to sharing sessions between application groups.

- Application session sharing works only for applications that are associated with application groups but are not associated with delivery groups. All applications that are associated directly with a delivery group share sessions by default.
- If an application is assigned to multiple application groups, make sure that the groups do not have conflicting settings. For example, one group with the option set to True, and another group's option set to False results in unpredictable behavior.

### **Rename an application group**

1. From **Manage > Full Configuration**, select **Applications** in the left pane, and then select the **Application Groups** tab.
2. Select an application group and then select **Rename Application Group** in the action bar.
3. Specify the new unique name and then select **OK**.

### **Add, remove, or change the priority of delivery group associations with an application group**

An application group can be associated with delivery groups containing shared (not private) machines that can deliver applications.

You can also select delivery groups containing shared machines that deliver only desktops, if both of the following conditions are met:

- The delivery group contains shared machines and was created with a version earlier than 7.9.
- You have Edit Delivery Group permission.

The delivery group type is automatically converted to “desktops and applications” when the **Edit Application Group** dialog is committed.

1. From **Manage > Full Configuration**, select **Applications** in the left pane, and then select the **Application Groups** tab.
2. Select an application group and then select **Edit Application Group** in the action bar.
3. Select the **Delivery Groups** page.
4. To add delivery groups, select **Add**. Select the check boxes of available delivery groups. (Incompatible delivery groups cannot be selected.) When you finish your selections, select **OK**.
5. To remove delivery groups, select the check boxes of the groups you want to remove and then select **Remove**. Confirm the deletion when prompted.
6. To change the priority of delivery groups, select the check box of the delivery group and then select **Edit Priority**. Enter the priority (0 = highest) and then select **OK**.
7. Select **Apply** to apply any changes you made and keep the window open, or select **OK** to apply changes and close the window.

## Add, change, or remove a tag restriction in an application group

Adding, changing, and removing tag restrictions can have unanticipated effects on which machines are considered for application launch. Review the considerations and cautions in [Tags](#).

1. From **Manage > Full Configuration**, select **Applications** in the left pane, and then select the **Application Groups** tab.
2. Select an application group and then select **Edit Application Group** in the action bar.
3. Select the **Delivery Groups** page.
4. To add a tag restriction, select **Restrict launches to machines with the tag** and then select the tag from the menu.
5. To change or remove a tag restriction, either select a different tag from the menu or remove the tag restriction by clearing **Restrict launches to machines with this tag**.
6. Select **Apply** to apply any changes you made and keep the window open, or select **OK** to apply changes and close the window.

## Add or remove users in an application group

For detailed information about users, see [Create application groups](#).

1. From **Manage > Full Configuration**, select **Applications** in the left pane, and then select the **Application Groups** tab.
2. Select an application group and then select **Edit Application Group** in the action bar.
3. Select the **Users** page. Indicate whether you want to allow all users in the associated delivery groups to use applications in the application group, or only specific users and groups. To add users, select **Add**, and then specify the users you want to add. To remove users, select one or more users and then select **Remove**.
4. Select **Apply** to apply any changes you made and keep the window open, or select **OK** to apply changes and close the window.

## Add, change, or remove an application icon in an application group

### Important:

This feature is available only in the web-based console.

Perform the following steps to add, change, or remove an application icon.

1. In the navigation pane, select **Applications**.
2. On the **All Applications** tab, select an application and then select **Properties**.  
To make changes at an application group level, navigate to the **Application Groups** tab, select an application in a group, and then select **Properties**.
3. Select the **Delivery** page and then select **Change**. The **Select Icon** window appears.

4. In the **Select Icon** window, do either of the following:
  - To add an icon, select **Add** and then browse to the icon.
  - To remove an icon, select it and then select **Remove**.
  - To change an icon, select it for the application.

**Important:**

- You cannot add an icon whose size is greater than 200 KB.
- You can add only .icon files.
- You cannot remove built-in icons.
- You cannot remove an icon of an application that is in use.

5. Select **OK** to apply changes and close the window.

## Change scopes in an application group

You can change a scope only if you have created a scope (you cannot edit the All scope). For more information, see [Delegated administration](#).

1. From **Manage > Full Configuration**, select **Applications** in the left pane, and then select the **Application Groups** tab.
2. Select an application group in the middle pane and then select **Edit Application Group** in the action bar.
3. Select the **Scopes** page. Select or clear the check box next to the scopes you want to change.
4. Select **Apply** to apply any changes you made and keep the window open, or select **OK** to apply changes and close the window.

## Delete an application group

An application must be associated with at least one delivery group or application group. If deleting an application group will result in one or more applications no longer belonging to a group, you are warned that deleting that group will also delete those applications. You can then confirm or cancel the deletion.

Deleting an application does not delete it from its original source. However, if you want to make it available again, you must add it again.

1. From **Manage > Full Configuration**, select **Applications** in the left pane, and then select the **Application Groups** tab.
2. Select an application group and then select **Delete Group** in the action bar.
3. Confirm the deletion when prompted.

## Remote PC Access

September 8, 2021

### Note:

This article describes how to configure Remote PC Access using the Full Configuration interface. If you're using the Quick Deploy interface, follow the guidance in [Remote PC Access in Quick Deploy](#).

Remote PC Access is a feature of Citrix Virtual Apps and Desktops that enables organizations to easily allow their employees to access corporate resources remotely in a secure manner. The Citrix platform makes this secure access possible by giving users access to their physical office PCs. If users can access their office PCs, they can access all the applications, data, and resources they need to do their work. Remote PC Access eliminates the need to introduce and provide other tools to accommodate teleworking. For example, virtual desktops or applications and their associated infrastructure.

Remote PC Access uses the same Citrix Virtual Apps and Desktops components that deliver virtual desktops and applications. As a result, the requirements and process of deploying and configuring Remote PC Access are the same as those required for deploying Citrix Virtual Apps and Desktops for the delivery of virtual resources. This uniformity provides a consistent and unified administrative experience. Users receive the best user experience by using Citrix HDX to deliver their office PC session.

The feature consists of a machine catalog of type **Remote PC Access** that provides the following functionality:

- Ability to add machines by specifying OUs. This ability facilitates the addition of PCs in bulk.
- Ability to add machines by using CSV files. This ability facilitates the addition of PCs in bulk in scenarios with OU structure restrictions.
- Automatic user assignment based on the user that logs into the office Windows PC. We support single user and multiple users assignments. By default, automatic assignment is restricted to a single user. To automatically assign multiple users to the next unassigned machine, navigate to **Full Configuration > Settings** and turn on the **Enable automatic assignment of multiple users for Remote PC Access** setting.

Citrix Virtual Apps and Desktops can accommodate more use cases for physical PCs by using other types of machine catalogs. These use cases include:

- Physical Linux PCs
- Pooled physical PCs (that is, randomly assigned, not dedicated)

### Notes:

For details on the supported OS versions, see the system requirements for the VDA for [single-session OS](#) and [Linux VDA](#).

For on-premises deployments, Remote PC Access is valid only for Citrix Virtual Apps and Desktops Advanced or Premium licenses. Sessions consume licenses in the same way as other Citrix Virtual Desktops sessions. For Citrix Cloud, Remote PC Access is valid for the Citrix Virtual Apps and Desktops Service and Workspace Premium Plus.

### Considerations

While all the technical requirements and considerations that apply to Citrix Virtual Apps and Desktops in general also apply to Remote PC Access, some might be more relevant or exclusive to the physical PC use case.

### Deployment considerations

While planning the deployment of Remote PC Access, make a few general decisions.

- You can add Remote PC Access to an existing Citrix Virtual Apps and Desktops deployment. Before choosing this option, consider the following:
  - Are the current Delivery Controllers or Cloud Connectors appropriately sized to support the additional load associated with the Remote PC Access VDAs?
  - Are the on-premises site databases and database servers appropriately sized to support the additional load associated with the Remote PC Access VDAs?
  - Will the existing VDAs and the new Remote PC Access VDAs exceed the number of maximum supported VDAs per site?
- You must deploy the VDA to office PCs through an automated process. The following are two of options available:
  - Electronic Software Distribution (ESD) tools such as SCCM: [Install VDAs using SCCM](#).
  - Deployment scripts: [Install VDAs using scripts](#).
- Review the [Remote PC Access security considerations](#).

### Machine catalog considerations

The type of machine catalog required depends on the use case:

- Remote PC Access
  - Windows dedicated PCs
  - Windows dedicated multi-user PCs. This use case applies to physical office PCs that multiple users can access remotely in different shifts.
- Single-session OS
  - Static - Dedicated Linux PCs
  - Random - Pooled Windows and Linux PCs

Once you identify the type of machine catalog, consider the following:

- A machine can be assigned to only one machine catalog at a time.
- To facilitate delegated administration, consider creating machine catalogs based on geographic location, department, or any other grouping that eases delegating administration of each catalog to the appropriate administrators.
- When choosing the OUs in which the machine accounts reside, select lower-level OUs for greater granularity. If such granularity is not required, you can choose higher-level OUs. For example, in the case of Bank/Officers/Tellers, select **Tellers** for greater granularity. Otherwise, you can select **Officers** or **Bank** based on the requirement.
- Moving or deleting OUs after being assigned to a Remote PC Access machine catalog affects VDA associations and causes issues with future assignments. Therefore, make sure to plan accordingly so that OU assignment updates for machine catalogs are accounted for in the Active Directory change plan.
- You can choose OUs to add machines to the machine catalog in bulk. In some scenarios, doing that is not easy because of OU structure restrictions. Instead, you can add machines in bulk by using CSV files. That feature gives you more flexibility to bulk add machines. You can add only machines (for use with user auto-assignments) or add machines along with user assignments.
- Integrated Wake on LAN is available only with the **Remote PC Access** type machine catalog.

### Linux VDA considerations

These considerations are specific to the Linux VDA:

- Use the Linux VDA on physical machines only in non-3D mode. Due to limitations on NVIDIA's driver, the local screen of the PC cannot be blacked out and displays the activities of the session when HDX 3D mode is enabled. Showing this screen is a security risk.
- Use machine catalogs of type single-session OS for physical Linux machines.
- The integrated Wake on LAN functionality is not available for Linux machines.

### Technical requirements and considerations

This section contains the technical requirements and considerations for physical PCs.

- The following are not supported:
  - KVM switches or other components that can disconnect a session.
  - Hybrid PCs, including All-in-One and NVIDIA Optimus laptops and PCs.
  - Dual boot machines.
- Connect the keyboard and mouse directly to the PC. Connecting to the monitor or other components that can be turned off or disconnected, can make these peripherals unavailable. If you must connect the input devices to components such as monitors, do not turn those components off.



- The PCs must be joined to an Active Directory Domain Services domain.
- Secure Boot is supported on Windows 10 only.
- The PC must have an active network connection. A wired connection is preferred for greater reliability and bandwidth.
- If using Wi-Fi, do the following:
  1. Set the power settings to leave the wireless adapter turned on.
  2. Configure the wireless adapter and network profile to allow automatic connection to the wireless network before the user logs on. Otherwise, the VDA does not register until the user logs on. The PC isn't available for remote access until a user has logged on.
  3. Ensure that the Delivery Controllers or Cloud Connectors can be reached from the Wi-Fi network.
- You can use Remote PC Access on laptop computers. Ensure the laptop is connected to a power source instead of running on the battery. Configure the laptop power options to match the options of a desktop PC. For example:
  1. Disable the hibernate feature.
  2. Disable the sleep feature.
  3. Set the close lid action to **Do Nothing**.
  4. Set the "press the power button" action to **Shut Down**.
  5. Disable video card and NIC energy-saving features.
- Remote PC Access is supported on Surface Pro devices with Windows 10. Follow the same guidelines for laptops mentioned previously.
- If using a docking station, you can undock and redock laptops. When you undock the laptop, the VDA reregisters with the Delivery Controllers or Cloud Connectors over Wi-Fi. However, when you redock the laptop, the VDA doesn't switch to use the wired connection unless you disconnect the wireless adapter. Some devices provide built-in functionality to disconnect the wireless adapter upon establishing a wired connection. The other devices require custom solutions or third-party utilities to disconnect the wireless adapter. Review the Wi-Fi considerations mentioned previously.

Do the following to enable docking and undocking for Remote PC Access devices:

1. In the **Start** menu, select **Settings > System > Power & Sleep**, and set **Sleep** to **Never**.
  2. Under the **Device Manager > Network adapters > Ethernet adapter** go to **Power Management** and clear **Allow the computer to turn off this device to save power**. Ensure that **Allow this device to wake the computer** is checked.
- Multiple users with access to the same office PC see the same icon in Citrix Workspace. When a user logs on to Citrix Workspace, that resource appears as unavailable if already in use by another user.

- Install the Citrix Workspace app on each client device (for example, a home PC) that accesses the office PC.

## Configuration sequence

This section contains an overview of how to configure Remote PC Access when using the **Remote PC Access** type machine catalog. For information on how to create other types of machine catalogs, see the [Create machine catalogs](#).

1. On-premises site only - To use the integrated Wake on LAN feature, configure the prerequisites outlined in [Wake on LAN](#).
2. If a new Citrix Virtual Apps and Desktops site was created for Remote PC Access:
  - a) Select the **Remote PC Access** site type.
  - b) On the **Power Management** page, choose to enable or disable power management for the default Remote PC Access machine catalog. You can change this setting later by editing the machine catalog properties. For details on configuring Wake on LAN, see [Wake on LAN](#).
  - c) Complete the information on the **Users** and **Machine Accounts** pages.

Completing these steps creates a machine catalog named **Remote PC Access Machines** and a delivery group named **Remote PC Access Desktops**.

3. If adding to an existing Citrix Virtual Apps and Desktops site:
  - a) Create a machine catalog of type **Remote PC Access** (Operating System page of the wizard). For details on how to create a machine catalog, see [Create machine catalogs](#). Make sure to assign the correct OU so that the target PCs are made available for use with Remote PC Access.
  - b) Create a delivery group to provide users access to the PCs in the machine catalog. For details on how to create a delivery group, see [Create delivery groups](#). Make sure to assign the delivery group to an Active Directory group that contains the users that require access to their PCs.
4. Deploy the VDA to the office PCs.
  - We recommend using the single-session OS core VDA installer ([VDAWorkstationCoreSetup.exe](#)).
  - You can also use the single-session full VDA installer ([VDAWorkstationSetup.exe](#)) with the `/remotepc` option, which achieves the same outcome as using the core VDA installer.
  - Consider enabling Windows Remote Assistance to allow help desk teams to provide remote support through Citrix Director. To do so, use the `/enable_remote_assistance` option. For details, see [Install using the command line](#).
  - To be able to see logon duration information in Director, you must use the single-session full VDA installer and include the **Citrix User Profile Management WMI Plugin** compo-

ment. Include this component by using the `/includeadditional` option. For details, see [Install using the command line](#).

- For information about deploying the VDA using SCCM, see [Install VDAs using SCCM](#).
- For information about deploying the VDA through deployment scripts, see [Install VDAs using scripts](#).

After you successfully complete steps 2–4, users are automatically assigned to their own machines when they log in locally on the PCs.

5. Instruct users to download and install Citrix Workspace app on each client device that they use to access the office PC remotely. Citrix Workspace app is available from the Citrix download site or the application stores for supported mobile devices.

## Features managed through the registry

### Caution:

Editing the registry incorrectly can cause serious problems that might require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

### Sleep mode (minimum version 7.16)

To allow a Remote PC Access machine to go into a sleep state, add this registry setting on the VDA, and then restart the machine. After the restart, the operating system power saving settings are respected. The machine goes into sleep mode after the preconfigured idle timer passes. After the machine wakes up, it reregisters with the Delivery Controller.

`HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\PortICA`

- Name: DisableRemotePCSleepPreventer
- Type: DWORD
- Data: 1

### Session management

By default, a remote user's session is automatically disconnected when a local user initiates a session on that machine (by pressing CTRL+ATL+DEL). To prevent this automatic action, add the following registry entry on the office PC, and then restart the machine.

`HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\PortICA\RemotePC`

- Name: SasNotification

- Type: DWORD
- Data: 1

By default, the remote user has preference over the local user when the connection message is not acknowledged within the timeout period. To configure the behavior, use this setting:

`HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\PortICA\RemotePC`

- Name: RpcaMode
- Type: DWORD
- Data:
  - 1 - The remote user always has preference if he or she does not respond to the messaging UI in the specified timeout period. This behavior is the default if this setting is not configured.
  - 2 - The local user has preference.

The timeout for enforcing the Remote PC Access mode is 30 seconds by default. You can configure this timeout, but do not set it lower than 30 seconds. To configure the timeout, use this registry setting:

`HKLM\SOFTWARE\Citrix\PortICA\RemotePC`

- Name: RpcaTimeout
- Type: DWORD
- Data: number of seconds for timeout in decimal values

When a user wants to forcibly get the console access: The local user can press Ctrl+Alt+Del twice in a gap of 10 seconds to get local control over a remote session and force a disconnect event.

After the registry change and machine restart, if a local user presses Ctrl+Alt+Del to log on to that PC while it is in use by a remote user, the remote user receives a prompt. The prompt asks whether to allow or deny the local user's connection. Allowing the connection disconnects the remote user's session.

## Wake on LAN

Remote PC Access supports Wake on LAN, which gives users the ability to turn on physical PCs remotely. This feature enables users to keep their office PCs turned off when not in use to save energy costs. It also enables remote access when a machine has been turned off inadvertently.

With the Wake on LAN feature, the magic packets are sent directly from the VDA running on the PC to the subnet in which the PC resides when instructed by the delivery controller. This allows the feature to work without dependencies on additional infrastructure components or third-party solutions for delivery of magic packets.

The Wake on LAN feature differs from the legacy SCCM-based Wake on LAN feature. SCCM-integrated Wake on LAN is an alternative Wake on LAN option for Remote PC Access that is only available with on-premises Citrix Virtual Apps and Desktops. For information on the SCCM-based Wake on LAN, see [Wake on LAN – SCCM-integrated](#).

## System requirements

The following are the system requirements for using the Wake on LAN feature:

- Control plane:
  - Citrix Virtual Apps and Desktops Service
  - Citrix Virtual Apps and Desktops 2009 or later
- Physical PCs:
  - VDA version 2009 or later
  - Windows 10. For supportability details, see the [VDA system requirements](#).
  - Wake on LAN enabled in BIOS/UEFI
  - Wake on LAN enabled in network adapter's properties within Windows configuration

## Configure Wake on LAN

To configure Wake on LAN, you can use the Full Configuration management interface or PowerShell.

### Important:

- You can use only the *web-based console* to configure Wake on LAN for Remote PC Access.
- The legacy console displays Remote PC Access catalogs configured with Wake on LAN that you added using the web-based console or PowerShell. However, the legacy console does not support adding machines to those catalogs. Use the web-based console instead.

## Configure Wake on LAN in the Full Configuration interface

### Important:

To configure Wake on LAN in the Full Configuration interface, verify that you have a Wake on LAN host connection created. This feature is not available to existing connections whose connection type is not **Remote PC Wake on LAN**.

A general workflow to configure Wake on LAN in the Full Configuration interface is as follows:

1. Create a Wake on LAN host connection to your resource location. (Select **Remote PC Wake on LAN** as the connection type.)
2. Create a Remote PC Access machine catalog.
3. Select the Wake on LAN connection you created for the catalog.

To select the Wake on LAN connection, navigate to the **Machine Catalog Setup > Operating System** page. If you decide not to choose a Remote PC Access power management connection, select **Do not use power management**. If needed, you can enable the Wake on LAN connection later by editing the catalog.

## Configure Wake on LAN through PowerShell

A general workflow to configure Wake on LAN through PowerShell is as follows:

1. Unless a Remote PC Access machine catalog exists, create one.
2. Unless a Wake on LAN host connection exists, create one.
3. Retrieve the unique identifier of the Wake on LAN host connection you created in Step 2.
4. Associate the Wake on LAN host connection with the machine catalog you created in Step 1.

To create the Wake on LAN host connection:

```
1 # Load Citrix SnapIns
2 Add-PSSnapIn -Name "*citrix*"
3
4 # Provide the name of the Wake on LAN host connection
5 [string]$connectionName = "Remote PC Access Wake on LAN"
6
7 # Create the hypervisor connection
8 $hypHc = New-Item -Path xdhyp:\Connections `
9             -Name $connectionName `
10            -HypervisorAddress "N/A" `
11            -UserName "woluser" `
12            -Password "wolpwd" `
13            -ConnectionType Custom `
14            -PluginId VdaWOLMachineManagerFactory `
15            -CustomProperties "<CustomProperties></CustomProperties>" `
16            -Persist
17
18 $bhc = New-BrokerHypervisorConnection -HypHypervisorConnectionUid
19             $hypHc.HypervisorConnectionUid
20
21 # Wait for the connection to be ready before trying to use it
22 while (-not $bhc.IsReady)
23 {
24     Start-Sleep -s 5
25     $bhc = Get-BrokerHypervisorConnection -HypHypervisorConnectionUid
26             $hypHc.HypervisorConnectionUid
27 }
28 <!--NeedCopy-->
```

When the host connection is ready, run the following commands to retrieve the host connection's unique identifier:

```
1 $bhc = Get-BrokerHypervisorConnection -Name "<WoL Connection Name>"
2 $hypUid = $bhc.Uid
3 <!--NeedCopy-->
```

After you retrieve the connection's unique identifier, run the following commands to associate the connection with the Remote PC Access machine catalog:

```
1 Get-BrokerCatalog -Name "<Catalog Name>" | Set-BrokerCatalog -
   RemotePCHypervisorConnectionUid $hypUid
2 <!--NeedCopy-->
```

### Design considerations

When you are planning to use Wake on LAN with Remote PC Access, consider the following:

- Multiple machine catalogs can use the same Wake on LAN host connection.
- For a PC to wake up another PC, both PCs must be in the same subnet and use the same Wake on LAN host connection. It does not matter if the PCs are in the same or different machine catalogs.
- Host connections are assigned to specific zones. If your deployment contains more than one zone, you need a Wake on LAN host connection in each zone. The same applies to machine catalogs.
- Magic packets are broadcasted using the global broadcast address 255.255.255.255. Ensure that the address is not blocked.
- There must be at least one PC turned on in the subnet - for every Wake on LAN connection - to be able to wake up machines in that subnet.

### Operational considerations

The following are considerations for using the Wake on LAN feature:

- The VDA must register at least once before the PC can be woken up using the integrated Wake on LAN feature.
- Wake on LAN can only be used to wake up PCs. It does not support other power actions, such as restart or shut down.
- After the Wake on LAN connection is created, it is visible in the Full Configuration interface. However, editing its properties within the Full Configuration interface is currently not supported.
- Magic packets are sent in one of the two ways:
  1. When a user tries to launch a session to their PC and the VDA is unregistered
  2. When an administrator manually sends a power on command from the Full Configuration interface or PowerShell

- Because the Delivery Controller is unaware of a PC's power state, the Full Configuration interface displays **Not Supported** under power state. The delivery controller uses the VDA registration state to determine whether a PC is on or off.

## Troubleshoot

### Monitor blanking not working

If the Windows PC's local monitor is not blank while there is an active HDX session (the local monitor displays what's happening in the session) it is likely due to issues with the GPU vendor's driver. To resolve the issue, give the Citrix Indirect Display driver (IDD) higher priority than the graphic card's vendor driver by setting the following registry value:

`HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\Graphics\AdapterMerits`

- Name: CitrixIDD
- Type: DWORD
- Data: 3

For more details about display adapter priorities and monitor creation, see the Knowledge Center article [CTX237608](#).

### Session disconnects when you select Ctrl+Alt+Del on the machine that has session management notification enabled

The session management notification controlled by the **SasNotification** registry value only works when Remote PC Access mode is enabled on the VDA. If the physical PC has the Hyper-V role or any virtualization-based security features enabled, the PC reports as a virtual machine. If the VDA detects that it is running on a virtual machine, it automatically disables Remote PC Access mode. To enable Remote PC Access mode, add the following registry value:

`HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\PortICA`

- Name: ForceEnableRemotePC
- Type: DWORD
- Data: 1

Restart the PC for the setting to take effect.

### Diagnostic information

Diagnostic information about Remote PC Access is written to the Windows Application Event log. Informational messages are not throttled. Error messages are throttled by discarding duplicate messages.

- 3300 (informational): Machine added to catalog



- 3301 (informational): Machine added to delivery group
- 3302 (informational): Machine assigned to user
- 3303 (error): Exception

### Power management

If power management for Remote PC Access is enabled, subnet-directed broadcasts might fail to start machines that are on a different subnet from the Controller. If you need power management across subnets using subnet-directed broadcasts, and AMT support is not available, try the Wake-up proxy or Unicast method. Ensure those settings are enabled in the advanced properties for the power management connection.

### The active remote session records the local touchscreen input

When the VDA enables Remote PC Access mode, the machine ignores the local touchscreen input during an active session. If the physical PC has the Hyper-V role or any virtualization-based security features enabled, the PC reports as a virtual machine. If the VDA detects that it is running on a virtual machine, it automatically disables Remote PC Access mode. To enable Remote PC Access mode, add the following registry setting:

`HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\PortICA`

- Name: ForceEnableRemotePC
- Type: DWORD
- Data: 1

Restart the PC for the setting to take effect.

### More resources

The following are other resources for Remote PC Access:

- Solution design guidance: [Remote PC Access Design Decisions](#).
- Examples of Remote PC Access architectures: [Reference Architecture for Citrix Remote PC Access Solution](#).

### Remove components

April 9, 2020

To remove components that you installed (such as VDAs), Citrix recommends using the Windows feature for removing or changing programs. Alternatively, you can remove components using the command line, or a script.

When you remove components, prerequisites are not removed, and firewall settings are not changed.

When you remove a VDA, the machine restarts automatically after the removal, by default.

## Remove components using the Windows feature for removing or changing programs

From the Windows feature for removing or changing programs:

- To remove a VDA, select **Citrix Virtual Delivery Agent <version>**, then right-click and select **Uninstall**. The installer launches and you can select the components to be removed.
- To remove the Universal Print Server, select **Citrix Universal Print Server**, then right-click and select **Uninstall**.

## Remove a VDA using the command line

Run the command that was used to install the VDA: `VDAServerSetup.exe`, `VDAWorkstationSetup.exe`, or `VDAWorkstationCoreSetup.exe`. See [Install using the command line](#) for syntax descriptions.

- To remove only the VDA or only the Citrix Workspace app, use the `/remove` and `/components` options.
- To remove the VDA and Citrix Workspace app, use the `/removeall` option.

For example, the following command removes the VDA and Citrix Workspace app from a multi-session OS machine.

```
VDAServerSetup.exe /removeall
```

For example, the following command removes the VDA but not Citrix Workspace app for Windows (if it is installed) from a single-session OS machine.

```
VDAWorkstationSetup.exe /remove /components vda
```

You can also remove a VDA using a script provided by Citrix. See [Remove VDAs using the script](#).

## User personalization layer

June 3, 2021

The user personalization layer feature for Citrix Virtual Apps and Desktops extends the capabilities of non-persistent machine catalogs. User personalization layers preserve users' data and locally installed applications across sessions. Powered by Citrix App Layering, this feature replaces Personal vDisk (PvD).

Like PvD, the user personalization layer feature supports Citrix Provisioning and Machine Creation Services (MCS) in a non-persistent machine catalog. You install the feature components alongside the Virtual Delivery Agent within the master image.

A VHD file stores any applications that the user installs locally. The VHD, which is mounted on the image, acts as the user's own user layer virtual hard drive.

This document includes instructions for deploying and configuring the user personalization layer feature. It describes the requirements for successful deployment, limitations, and known issues.

To use the User personalization layer feature, you must first deploy it using the steps detailed in the article. Until then, the feature is not available for you to use.

## Application support

Aside from the following exceptions, all applications that a user installs locally on the desktop are supported in the user personalization layer.

### Exceptions

The following applications are the exception and are not supported on the user personalization layer:

- Enterprise applications, such as MS Office and Visual Studio.
- Applications that modify the network stack or hardware. Example: a VPN client.
- Applications that have boot level drivers. Example: a virus scanner.
- Applications with drivers that use the driver store. Example: a printer driver.

#### Note:

You can make printers available using Windows Group Policy Objects (GPOs).

Do *not* allow users to install any unsupported applications locally. Rather, install these applications directly on the master image.

### Applications that require a local user or administrator account

When a user installs an application locally, the app goes into their user layer. If the user then adds or edits a local user or group, the changes do not persist beyond the session.

**Important:**

Add any required local user or group in the master image.

## Requirements

The user personalization layer feature requires the following components:

- Citrix Virtual Apps and Desktops 7 1909 or later
- Virtual Delivery Agent (VDA), version 1912
- Citrix Provisioning, version 1909 or later
- Windows File Share (Server Message Block protocol, SMB)

You can deploy the user personalization layer feature on the following Windows versions when the OS is deployed as single session. Support is limited to a single user on a single session.

- Windows 10 Enterprise x64, version 1607 or later
- Windows 10 Multi-Session\*
- Windows Server 2016\*
- Windows Server 2019\*

\*For Citrix Virtual Apps and Desktops 7, Azure Files with user personalization layers is supported on Windows Server 2019, Windows Server 2016v, and Windows 10 client. Desktop VDAs running Windows 10 and single-user server VDAs running Windows server 2016 or 2019 are also supported.

When using a Server OS, UPL is supported only as a Server VDI deployment For details, see the [Server VDI](#) article.

If you installed the preview version of the user personalization layer feature, uninstall the software and reboot the master image before installing this release.

## Set up your file share

The user personalization layer feature requires Windows Server Message Block (SMB) storage. To create a Windows file share, follow the usual steps for the Windows operating system that you are on.

For details about using Azure Files with Azure-based catalogs, see [Set up Azure Files storage for User personalization layers](#).

## Recommendations

Follow the recommendations in this section for a successful user personalization layer deployment.

### **Profile Management solution**

User personalization layer stores all changes the user makes for a single machine catalog image. To add enhanced capabilities such as roaming profile data across multiple catalog images, Citrix recommends also using Profile Management. Refer to the [Profile Management documentation](#) for more details.

When using Profile Management with the user personalization layer feature, clear deletion of the user's information on logoff. You can clear deletion using a Group Policy Object (GPO) or the policy on the Delivery Controller (DDC).

For details about available Profile Management policies, see [Profile Management policy descriptions and defaults](#).

### **Microsoft System Center Configuration Manager (SCCM)**

If you are using SCCM with the user personalization layer feature, follow the Microsoft guidelines for preparing your image in a VDI environment. Refer to this [Microsoft TechNet article](#) for more information.

### **Maximum user layer size**

We recommend at least 10 GB as the user layer size.

**Note:**

During installation, the value zero (0) results in the default user layer size of 10 GB.

### **A quota set in Windows can override the maximum user layer size**

You can override Studio's maximum user layer size by defining a quota for the user layer file share. The user layer size is set to a maximum of the quota size.

To set a hard quota on the user layer size, use either of Microsoft's quota tools:

- File Server Resource Manager (FSRM)
- Quota Manager

The quota must be set on the user layer directory named Users.

**Note:**

Increasing or decreasing the quota only impacts new user layers. It does not change the maximum size of existing user layers. Existing user layers remain unchanged when the quota is updated.

## Deploy a user personalization layer

When deploying the user personalization feature, you define the policies within Studio. You then assign the policies to the delivery group bound to the machine catalog, where the feature is deployed.

If you leave the master image with no user personalization layer configuration, the services remain idle and do not interfere with authoring activities.

If you set the policies in the master image, the services attempt to run and mount a user layer within the master image. The master image would exhibit unexpected behaviors and instability.

To deploy the user personalization layer feature, complete the following steps in this order:

- Step 1: Verify availability of a Citrix Virtual Apps and Desktops environment.
- Step 2: Prepare your master image.
- Step 3: Create a machine catalog.
- Step 4: Create a delivery group.
- Step 5: Create delivery group custom policies.


### Step 1: Verify that the Citrix Virtual Apps and Desktops environment is available

Be sure that your Citrix Virtual Apps and Desktops environment is available to use with this new feature. For setup details, see [Install and configure Citrix Virtual Apps and Desktops](#).


### Step 2: Prepare your master image

To prepare your master image:

1. Locate the master image. Install your organization's enterprise applications and any other apps your users generally find useful.
2. If you are deploying Server VDI, follow the steps in the [Server VDI](#) article. Be sure to include the optional component, **User personalization layer**. For details, see the [Command-line options to install a VDA](#).
3. If you are using Windows 10, install Virtual Delivery Agent (VDA) 1912. If an older version of the VDA is already installed, uninstall the old version first. When installing the new version, be sure to select and install the optional component, **Citrix User Personalization Layer**, as follows:
  - a) Click the tile, **Virtual Delivery Agent for Windows Desktop OS**:




- a) **Environment:** Select either Create a master MCS image or Create a master image using Citrix Provisioning or third-party provisioning tools.



- a) **Core Components:** Click **Next**.
- b) **Additional Components:** Check **Citrix User Personalization Layer**.





a) Click through the remaining installation screens, configuring the VDA as needed, and click Install. The image reboots one or more times during installation.

4. Leave **Windows updates** disabled. The user personalization layer installer disables Windows updates on the image. Leave the updates disabled.


The image is ready for you to upload into Studio.

### Step 3: Create a machine catalog


In Studio, follow the steps to create a machine catalog. Use the following options during catalog creation:

1. Select **Operating System** and set it to **Single session OS**.
2. Select **Machine Management** and set it to **Machines that are power managed**. For example, virtual machines or blade PCs.
3. Select **Desktop Experience** and set it to either **pooled-random** or **pooled-static** catalog type, as in the following examples:

- **Pooled-random:**



- **Pooled-static:** If you select pooled-static, configure desktops to discard all changes and clear virtual desktops when the user logs off, as shown in the following screenshot:



**Note:**

User personalization layer does not support pooled-static catalogs configured to use Citrix Personal vDisk or assigned as dedicated virtual machines.

4. If you are using MCS, select **Master Image** and the snapshot for the image created in the previous section.
5. Configure the remaining catalog properties as needed for your environment.

**Step 4: Create a delivery group**

Create and configure a **delivery group**, including machines from the machine catalog you created. For details, see the [Create Delivery Groups](#).

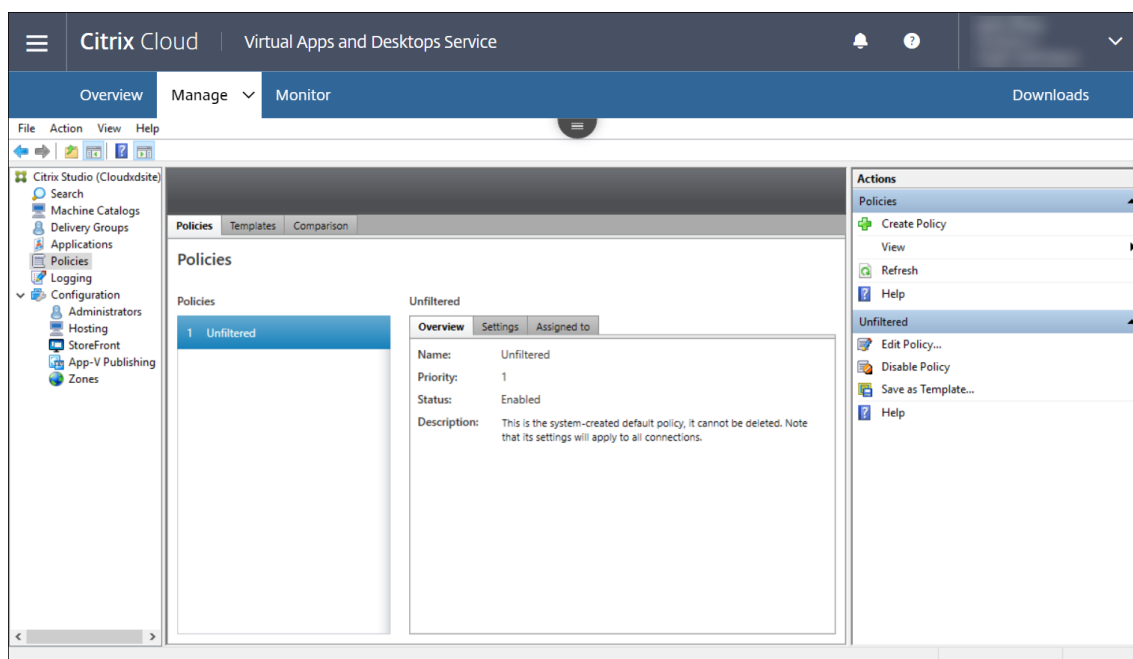
**Step 5: Create delivery group custom policies**

To enable mounting of user layers within the Virtual Delivery Agents, you use configuration parameters to specify:

- Where on the network to access the user layers.
- How large to permit the user layer disks to grow.

To define the parameters as custom Citrix policies in Studio and assign them to your delivery group.

1. In Studio, select Policies in the navigation pane:




2. Select Create Policy in the Actions pane. The Create Policy window appears.
3. Type 'user layer' into the search field. The following two policies appear in the list of available policies:
  - User Layer Repository Path
  - User Layer Size GB


**Note:**

Changing the User Layer Size in the policy does not change the size of existing layers.


4. Click **Select** next to User Layer Repository Path. The Edit Setting window appears.




5. Enter a path in the format `\\server name or address\folder name` in the Value field, Click **OK**:



6. Optional: Click **Select** next to User Layer Size in GB:



7. The Edit Settings window appears.
  8. Optional: Change the default value of '0' to the maximum size (in GB) that the user layer can grow. Click OK.
- Note:**
- If you keep the default value, the maximum user layer size is 10 GB.
9. Click Next to configure Users and Machines. Click the Delivery Group Assign link highlighted in this image:



10. In the Delivery Group menu, select the delivery group created in the previous section. Click OK.



**Assign Policy**

**Delivery Group**  
**Applies to:** Virtual Delivery Agent: 5.6, 7.0 Server OS, 7.0 Desktop OS, 7.1 Server OS, 7.1 Desktop OS, 7.5 Server OS, 7.5 Desktop OS, 7.6 Server OS, 7.6 Desktop OS, 7.7 Server OS, 7.7 Desktop OS, 7.8 Server OS, 7.8 Desktop OS, 7.9 Server OS, 7.9 Desktop OS, 7.11 Server OS, 7.11 Desktop OS, 7.12 Server OS, 7.12 Desktop OS, 7.13 Server OS, 7.13 Desktop OS, 7.14 Server OS, 7.14 Desktop OS, 7.15 Server OS, 7.15 Desktop OS, 7.16 Server OS, 7.16 Desktop OS, 7.17 Server OS, 7.17 Desktop OS, 7.18 Server OS, 7.18 Desktop OS, 1808 Server OS, 1808 Desktop OS, 1811 Server OS, 1811 Desktop OS, 1903 Server OS, 1903 Desktop OS, 1906 Server OS, 1906 Desktop OS, 1909 Server OS, 1909 Desktop OS


Apply policy based on the delivery group membership of the desktop running the session.

**Delivery Group elements:**

Mode	Controller	Delivery Group	
Allow		Win10 - UPL	+ -
<input checked="" type="checkbox"/> Enable			

OK Cancel

11. Enter a name for the policy. Click the check box to enable the policy, and click Finish.



### Configure security settings on the user layer folder

As a domain administrator, you can specify more than one storage location for your user layers. Create a `\Users` subfolder for each storage location (including the default location). Secure each location using the following settings.

Setting name	Value	Apply to
Creator Owner	Modify	Subfolders and Files only
Owner Rights	Modify	Subfolders and Files only
Users or group:	Create Folder, Append Data, Traverse Folder, Execute File, List Folders, Read Data, Read Attributes	Selected Folder Only
System	Full Control	Selected Folder, Subfolders, and Files

Setting name	Value	Apply to
Domain Admins and selected Admin group	Full Control	Selected Folder, Subfolders, and Files

## User layer messages

When a user is unable to access their user layer, they receive one of these notification messages.

- **User Layer In Use**

We were unable to attach your user layer because it is in use. Any changes you make to application settings or data will not be saved. Be sure to save any work to a shared network location.<!--NeedCopy-->

- **User Layer Unavailable**

We were unable to attach your user layer. Any changes you make to application settings or data will not be saved. Be sure to save any work to a shared network location.<!--NeedCopy-->

- **System not reset after user sign-out**

This system was not shut down properly. Please log off immediately and contact your system administrator.<!--NeedCopy-->

## Log files to use when troubleshooting

The log file, `ulayersvc.log`, contains the output of the user personalization layer software where changes are logged.

```
1 C:\ProgramData\Unidesk\Logs\ulayersvc.log
2 <!--NeedCopy-->
```

## Limitations

Keep the following limitations in mind when installing and using the user personalization layer feature.

- Do *not* configure the user personalization layer feature with persistent machine catalogs.
- Do *not* use Session hosts.
- Do *not* update the machine catalog with an image running a new OS install (even the same version of Windows 10). Best practice is to apply updates to the OS within the same master image used when creating the machine catalog.

- Do *not* use boot-time drivers, or any other early boot personalization for user-installed apps.
- Do *not* migrate PvD data to the user personalization layer feature.
- Do *not* migrate existing user layers from the full App Layering product to the user personalization layer feature.
- Do *not* change the user layer SMB path to access user layers created using a different master OS image.
- Do *not* enable Secure Boot within User personalization layer virtual machines, as it is not currently supported.
- When a user logs out of a session and then logs in again, the new session runs on a different machine in the pool. In a VDI environment, Microsoft Software Center lists an application as **Installed** on the first machine, but shows it as **Unavailable** on the second machine.

To find out the true status of the application, instruct the user to select the application in Software Center and click **Install**. SCCM then updates the status to the true value.

- Software Center occasionally stops immediately after launching within a VDA that has the user personalization layer feature enabled. To avoid this issue, follow Microsoft's recommendations for [Implementing SCCM in a XenDesktop VDI environment](#). Also, make sure that the `ccmexec` service is running before you start the Software Center.
- In Group Policies (Computer Settings), User layer settings override settings applied to the master image. Therefore, changes you make in Computer Settings using a GPO are not always present for the user on the next session login.

To get around this issue, create a User Logon Script that issues the command:

```
gpupdate /force
```

For example, one customer set the following command to run at each user login:

```
gpupdate /Target:Computer /force
```

For best results, apply changes to Computer Settings directly on the user layer, after the user has logged in.

## Upgrade

June 16, 2021

## Introduction

Citrix maintains all the Citrix Virtual Apps and Desktops service components in your deployment, except VDAs.

Before beginning a VDA upgrade:

- Review this entire article, so you know what to expect.
- Review the [Lifecycle policy](#) for the service.

To upgrade a VDA, download a VDA installer and run it on the machine or image. You can use the installer's graphical or command-line interface. For guidance, see:

- [VDA installers](#)
- [Install VDAs using the graphical interface](#)
- [Install VDAs using the command line](#)

If the VDA was originally installed using `VDAWorkstationCoreSetup.exe`:

- You retain that configuration if you upgrade the VDA with the latest version of the same installer.
- If you run `VDAWorkstationSetup.exe` on that machine, you can enable the features that are not supported in `VDAWorkstationCoreSetup.exe`. Keep in mind that some of those features might be enabled by default in the `VDAWorkstationSetup.exe` installer. You can also install Citrix Workspace app.

When upgrading a VDA to version 7.17 or a later supported version, a machine restart occurs during the upgrade process. This restart cannot be avoided. The upgrade resumes automatically after the restart (unless you specify `/noresume` on the command line).

After you upgrade VDAs, [update the images and catalogs](#) that use that VDA.

## If the VDA has Personal vDisk installed

If the Personal vDisk (PvD) component was ever installed on a VDA, that VDA cannot be upgraded to version 1912 LTSR or later until you remove that component.

This instruction applies even if you never used PvD. Here's how the PvD component might have been installed in earlier versions:

- In the VDA installer's graphical interface, PvD was an option on the **Additional Components** page. The 7.15 LTSR and earlier 7.x releases enabled this option by default. So, if you accepted the defaults (or explicitly enabled the option in any release), PvD was installed.
- On the command line, the `/baseimage` option installed PvD. If you specified this option, or used a script that contained this option, PvD was installed.

## What to do

If the VDA installer does not detect the PvD component in the currently installed VDA, the upgrade proceeds as usual.

If the installer detects the PvD component in the currently installed VDA:

- **Graphical interface:** The upgrade pauses. A message asks if you want the unsupported component removed automatically. When you click **OK**, the component is removed automatically and the upgrade proceeds.
- **CLI:** The command fails if the installer detects the PvD component. To avoid command failure, include the following option in the command: `/remove_pvd_ack`.

If you want to continue using PvD on your Windows 7 or Windows 10 (1607 and earlier, without updates) machines, VDA 7.15 LTSR is the latest supported version.

## Earlier operating systems

The [System requirements](#) article lists the supported Windows operating systems for current release VDAs.

- For LTSR VDAs, see the system requirements article for your LTSR version.
- For Linux VDAs, see the [Linux Virtual Delivery Agent](#) documentation.

For Windows machines with OSs that are no longer supported for installation of the latest VDA (such as Windows 7 and Windows Server 2008 R2), you have several options.

- Windows 7 is supported in an Azure Virtual Desktop environment.
- For non-WVD environments:
  - Reimage the machine to a supported Windows version, and then install the new VDA.
  - If reimaging the machine is not an option but you want to upgrade the OS, uninstall the VDA before upgrading the OS. Otherwise, the VDA will be in an unsupported state. After upgrading the OS, install the new VDA.
  - To continue using machines with an OS that is no longer supported, XenApp and XenDesktop 7.15 LTSR is the most current supported VDA version for Windows 7 and Windows Server 2008 R2.
    - \* If the machine has version 7.15 LTSR installed (and you try to install a newer version), a message informs you that you're using the latest supported version.
    - \* If the machine has a version earlier than 7.15 LTSR installed, a message guides you to [CTX139030](#) for information. You can download 7.15 LTSR VDAs from the Citrix website.

## Migrate to cloud

July 8, 2021

If you have a Citrix Virtual Apps and Desktops on-premises configuration and you want to move it to a Citrix Virtual Apps and Desktops Service deployment, or you have a Citrix Virtual Apps and Desktops Service deployment that you want to move to another region, you want to back up and restore your configuration, or you have reached your resource limits, you can now migrate all or part of your configuration using the Automated Configuration tool.

The following 2-minute video provides a quick tour of Automated Configuration.



For more information on Automated Configuration, see [Proof of Concept: Automated Configuration Tool](#) on Tech Zone.

For a deeper look into moving your deployment and readying your configuration for migration, see [Deployment Guide: Migrating Citrix Virtual Apps and Desktops from on-premises to Citrix Cloud](#) on Tech Zone.

### Getting started with Automated Configuration

Automated Configuration for Citrix Virtual Apps and Desktops is a tool that allows you to automate migrating your Citrix Virtual Apps and Desktops Service deployment from:

- An on-premises site

- Another cloud site
- A cloud site backup

Automated Configuration exports your configuration information into a collection of .yaml files that can then be optionally edited (for staging your migration) and imported into your Citrix Virtual Apps and Desktops Service deployment.

This section provides common details needed to perform the functions of Automated Configuration.

### Download Automated Configuration

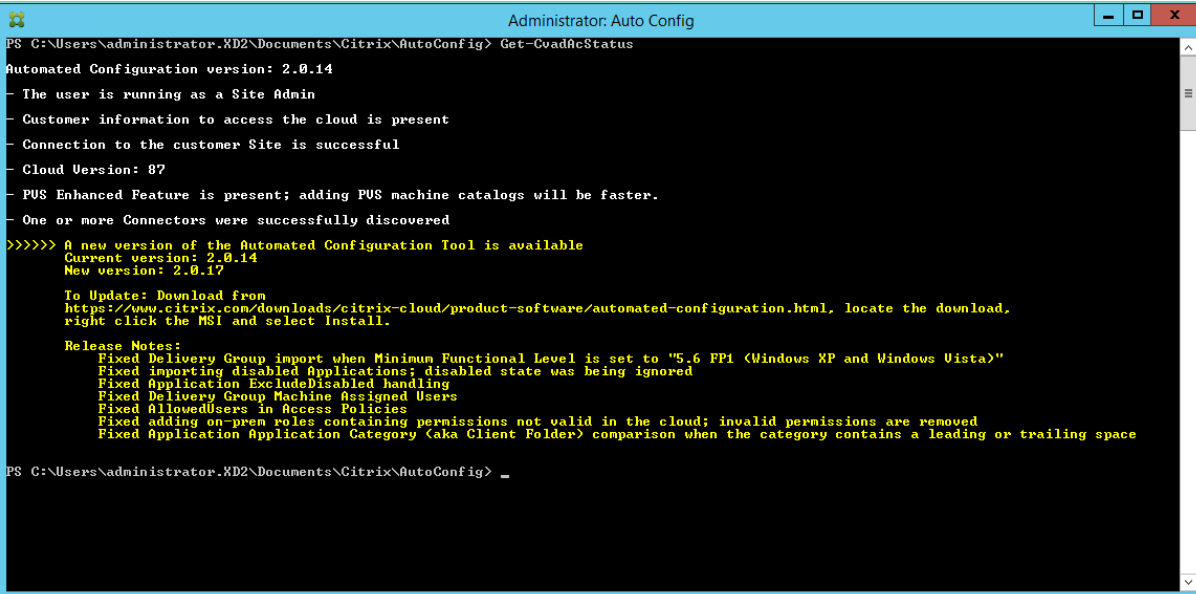
Download and install the Automated Configuration tool from [Citrix Downloads](#).

#### Important:

To prevent errors in functionality, you should always use the latest available version of Automated Configuration.

### Upgrading Automated Configuration

When running cmdlets that access the cloud in Automated Configuration, the tool alerts you when there is a newer version available for download.



```
Administrator: Auto Config
PS C:\Users\administrator.XD2\Documents\Citrix\AutoConfig> Get-CvadAcStatus
Automated Configuration version: 2.0.14
- The user is running as a Site Admin
- Customer information to access the cloud is present
- Connection to the customer Site is successful
- Cloud Version: 87
- PUS Enhanced Feature is present; adding PUS machine catalogs will be faster.
- One or more Connectors were successfully discovered
>>>>> A new version of the Automated Configuration Tool is available
Current version: 2.0.14
New version: 2.0.17

To Update: Download from
https://www.citrix.com/downloads/citrix-cloud/product-software/automated-configuration.html, locate the download,
right click the MSI and select Install.

Release Notes:
Fixed Delivery Group import when Minimum Functional Level is set to "5.6 FPI (Windows XP and Windows Vista)"
Fixed importing disabled Applications; disabled state was being ignored
Fixed Application ExcludeDisabled handling
Fixed Delivery Group Machine Assigned Users
Fixed AllowedUsers in Access Policies
Fixed adding on-prem roles containing permissions not valid in the cloud; invalid permissions are removed
Fixed Application Application Category (aka Client Folder) comparison when the category contains a leading or trailing space

PS C:\Users\administrator.XD2\Documents\Citrix\AutoConfig> _
```

You can make sure you have the latest version by following the steps below:

1. Double-click the **Auto Config** icon. A PowerShell window appears.
2. Run the following command to check your version number.

```
Get-CvadAcStatus
```



3. Check your tool version against the version listed in the alert or at [Citrix Downloads](#). The latest version of the tool is located there.
4. Download and install the latest version of the tool. You do **not** need to uninstall the old version to upgrade Automated Configuration.

**Note:**

The alert appears every time you run a cmdlet that accesses the cloud. For more information on cmdlets, see Automated Configuration tool cmdlets.

### Prerequisites for migrating your configuration

Users of Automated Configuration must be proficient with the following:

- Citrix Virtual Apps and Desktops setup and administration
- The YAML standard
- PowerShell

For **exporting** your configuration from Citrix Virtual Apps and Desktops:

- Citrix Virtual Apps and Desktops: current release and its immediate predecessor or Citrix Virtual Apps and Desktops, XenApp and XenDesktop LTSRs: all versions
- An on-premises Delivery Controller and at least a single on-premises VDA
- A machine with .NET Framework 4.7.2 or later
- A machine with the Citrix PowerShell SDK. This is automatically installed on the Delivery Controller. (To run on machine other than the on-premises Delivery Controller, the machine must be domain joined and Citrix Studio must be installed, as Studio installs the correct PowerShell snap-ins. The Studio installer can be found on the Citrix Virtual Apps and Desktops [installation media](#).)

For **importing** your configuration into Citrix Virtual Apps and Desktops Service:

- A machine with access to Citrix Cloud. This does not have to be a Delivery Controller or a domain joined machine.
- Citrix Virtual Apps and Desktops Service provisioned.
- An active resource location with Connector installed and domain-joined to the same domain as the on-premises setup.
- Connectivity to sites accessing Citrix Cloud must be allowed and available. For more information, see [Virtual Apps and Desktops service service connectivity](#).

### Supported objects for migration

Automated Configuration supports the following components.

- Tags

- Delegated Admin
  - Scopes
  - Roles
- Host Connections
  - A Single Resource Pool
  - Admin Scopes
- Machine Catalogs
  - Admin Scopes
  - Machines
  - Remote PC Access, Physical, Pooled, Provisioned (except MCS), Assigned
- Delivery Groups
  - Access Policy
  - Admin Scope Association
  - Application Access Policy
  - Assignment Policy
  - Entitlement/Desktop Policy
  - Power Schedules
  - Session Lingering
  - Session Prelaunch
  - Reboot Schedules
  - Tags
- Application Groups
  - Admin Scope Association
  - Delivery Groups
  - Users and Groups
- Applications
  - Application Folders
  - Icons
  - Applications
  - Broker Configured FTAs
  - Tags
- Group Policies

### **Component dependency**

When creating an on-premises site, the creation steps must be followed in a specific order, due to the dependencies on previous items. Automated Configuration works on the same principles.

The import process is designed to accurately perform updates, only perform needed updates, and verify that all updates have been correctly made.

Components must be imported in an order that creates dependencies in the order they are needed. The following list identifies selectable components in the [Export](#), [Import](#), [Merge](#), [New](#), [Sync](#), and [Restore](#) commands. They are listed in their dependency order. Importing out of dependency order might result in failures and cause operations to fail.

1. Tags
2. Admin Roles and Scopes (includes Scopes and Roles, does not include administrators)
3. Host Connections
4. Machine Catalogs
5. Delivery Groups (includes StoreFront)
6. Application Groups
7. Applications (includes Application Folders)
8. Group Policies

Components must be imported in an order that creates dependencies in the order they are needed. Automated Configuration correctly orders components when importing, merging, or restoring multiple components in a single cmdlet execution.

When importing single components with multiple commands, it is necessary to order the components so dependencies are imported in the required order.

Individual imports can then be made for updates after a complete site configuration has been imported successfully.

### Dependency details

The components and their dependencies are listed here. A component's dependencies must be in place before it can be imported or merged. If a dependency is missing, it can cause the import or merge command to fail. The **Fixups** section of the log file shows missing dependencies if an import or merge fails.

1. Tags
  - No pre-dependencies
2. Delegated Admin
  - No pre-dependencies
3. Host Connections
  - Security Information in CvadAcSecurity.yml
4. Machine Catalogs
  - Machines present in Active Directory
  - Host Connections
  - Tags
5. Delivery Groups
  - Machines present in Active Directory

- Users present in Active Directory
  - Machine Catalogs
  - Tags
6. Application Groups
- Delivery Groups
  - Tags
7. Applications
- Delivery Groups
  - Application Groups
  - Tags
8. Group Policies
- Delivery Groups
  - Tags

### **Known limitations**

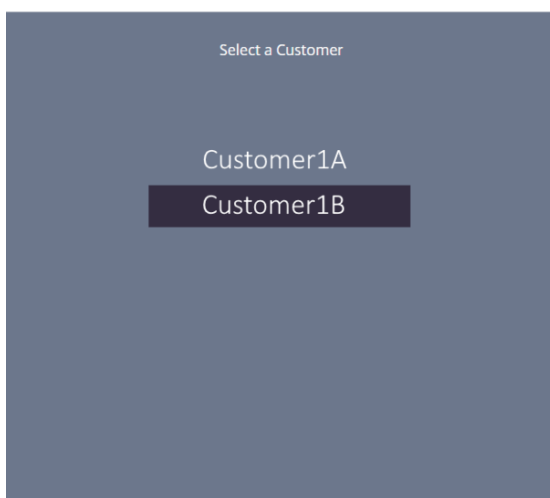
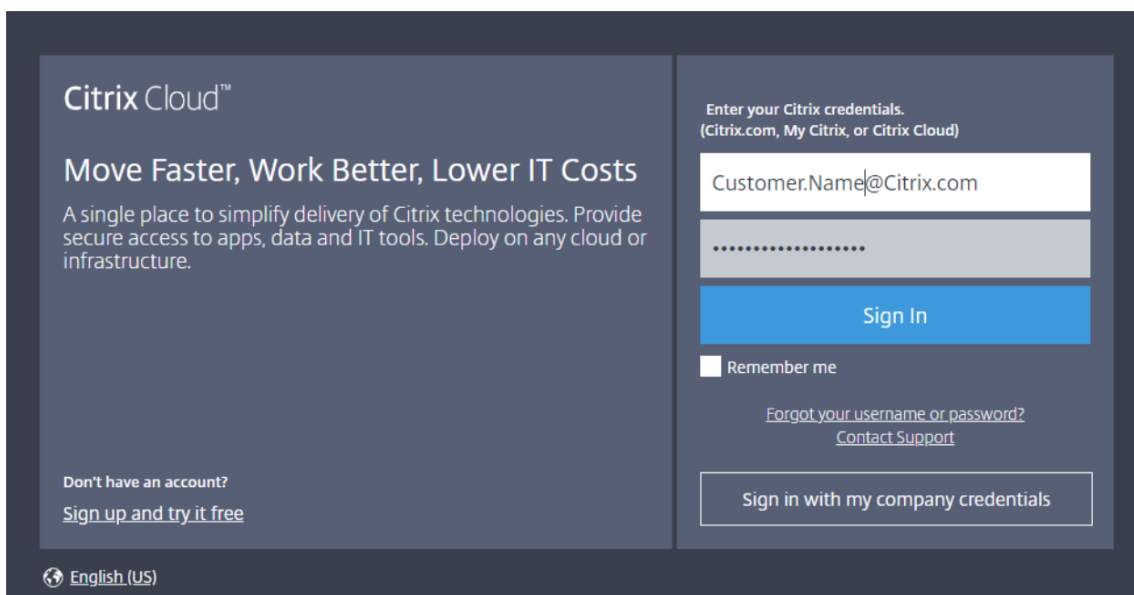
- Machine catalogs provisioned through Machine Creation Services are not currently supported. You can still import other objects from your site. For more information on MCS, see [Importing other objects when you have Machine Creation Services machine catalogs](#).
- Icons are not applied to machines or desktops.

### **Generating the customer ID, client ID, and secret key**

The following steps allow you to retrieve the customer ID and create the client ID and secret key that are required to import your configuration to Citrix Cloud. All cmdlets accessing the cloud require these values.

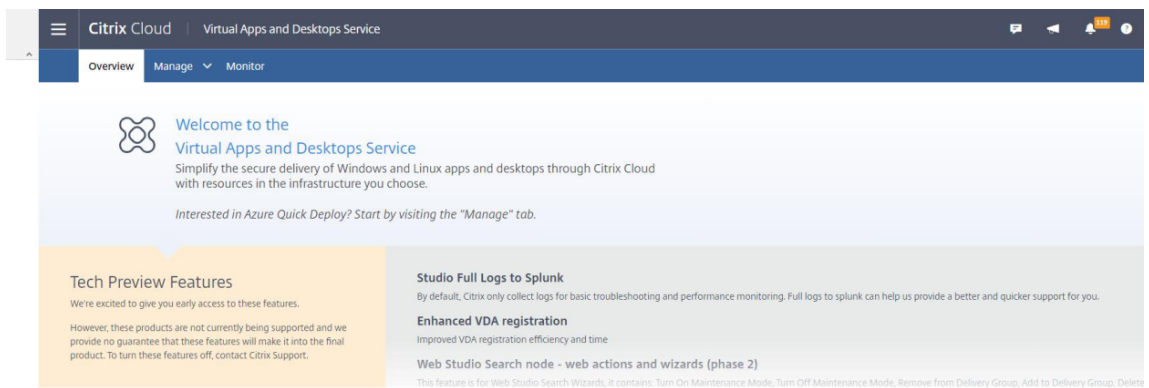
To retrieve the **Customer ID**:

1. Sign into your Citrix Cloud account and select the customer.



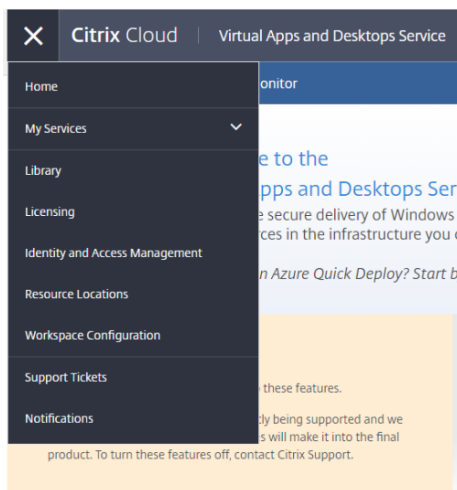
2. Click the hamburger menu, then select **Identity and Access Management** in the drop-down menu.

# Citrix Virtual Apps and Desktops service




## Get Started with your Virtual Apps and Desktops Service

- 1** Connect to infrastructure  
Install two Cloud Connectors to enable communications between the resource location and Citrix Cloud.
- 2** Register resources  
Install a Virtual Delivery Agent. VDAs manage connections between VMs that deliver apps and desktops and the user device.
- 3** Create collection of resources  
Create a machine catalog of VMs containing apps and desktops to be delivered.
- 4** Assign users  
Create a delivery group to specify who can use the apps and desktops.
- 5** Launch apps and desktops  
Users can now access the virtualized apps and desktops through the [Workspace URL](#).



3. The **Customer ID** is located on the **Identity and Access Management** page.

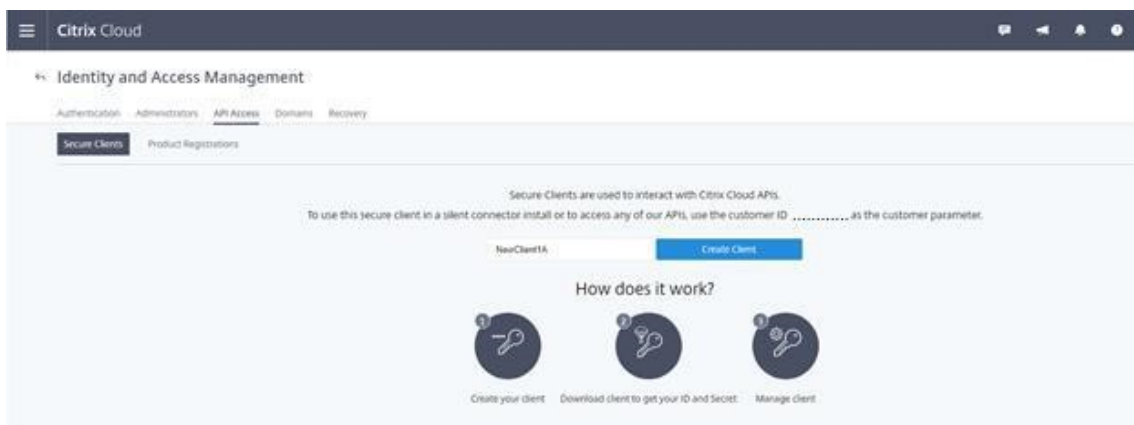


To retrieve the **Client ID** and **Secret Key**:

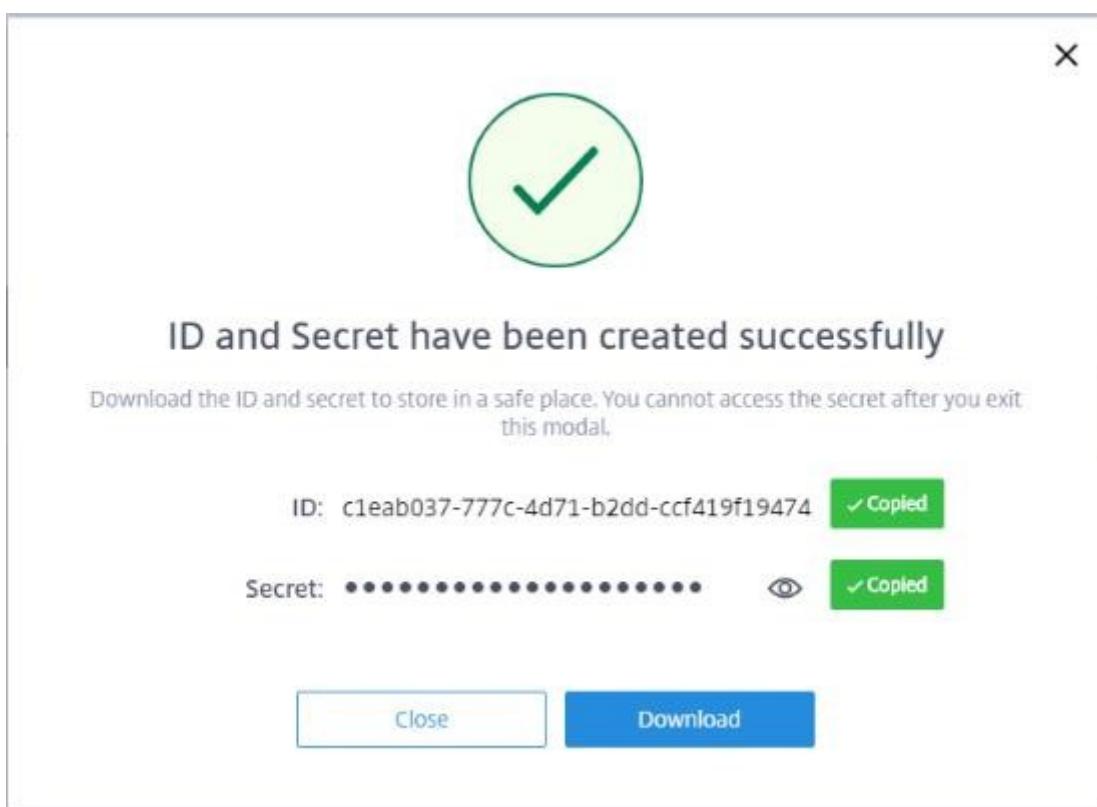
1. On the **Identity and Access Management** page, click the **API Access** tab.



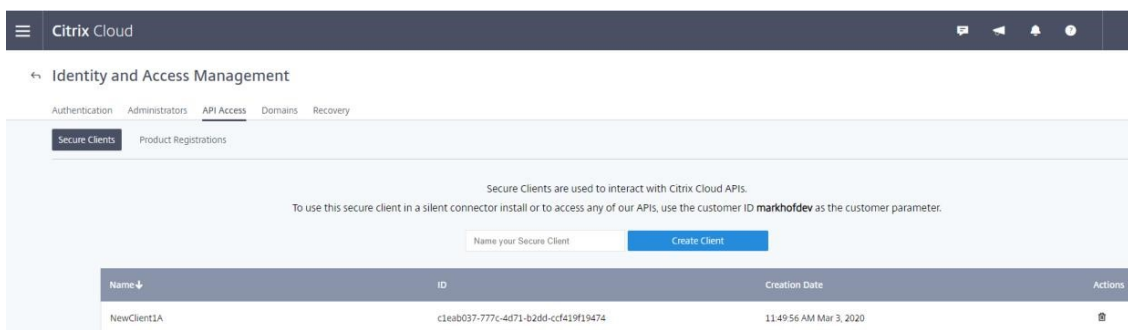
2. Enter a name in the box. This name is used to differentiate between multiple client IDs and secret keys. Click **Create Client** to create the client ID and the secret key.



3. The following dialog appears after you successfully create the client ID and the secret key. Be sure to copy both values to a secure location or download the .csv file containing this information.



4. The client ID and the secret key are successfully created.



Place these values in a secure location and share only with trusted company members who need access to the tool or access the cloud Rest APIs. The client ID and secret key do not expire. If they are compromised, immediately remove them by using the **Trash** icon and create new ones.

**Note:**

The secret key cannot be retrieved if it is lost or forgotten; a new client ID and secret key must be created.



## Populating customer info file

Using the CustomerInfo.yml file eliminates the need to provide customer information parameters with each cmdlet's execution. Any of the customer information can be overridden by using cmdlet parameters.

Create the CustomerInfo.yml file by using the `New-CvadAcCustomerInfoFile` cmdlet. `New-CvadAcCustomerInfoFile` has the following required parameters.

- `CustomerId` – customer's ID.
- `ClientId` – customer's client ID created on Citrix Cloud.
- `Secret` – customer's secret created on Citrix Cloud.

```
New-CvadAcCustomerInfoFile -CustomerId markhof123 -ClientId 6813EEA6-46CC-4F8A-BC71-539F2DAC5984 -Secret TwBLaaaaaaaaaaaaaaaaaw==
```

Update the CustomerInfo.yml file by using the `Set-CvadAcCustomerInfoFile` cmdlet.

**Note:**

The cmdlet only changes the Client ID.

```
Set-CvadAcCustomerInfoFile -ClientId C80487EE-7113-49F8-85DD-2CFE30CC398E
```

The following is a sample CustomerInfo.yml file.

```
1 ---
2 # Created/Updated on 2020/01/29 16:46:47
3 CustomerId: "markhof123"
4 CustomerId: "markhof123"
5 ClientId: "6713FEA6-46CC-4F8A-BC71-539F2DDK5384"
6 Secret: "TwBLaaabbbbaaaaaaaaaaw=="
7 LogFileName: "CitrixLog"
8 Environment: Production
9 AltRootUrl: ""
10 StopOnError: False
11 AlternateFolder: ""
12 Locale: "en-us"
13 Editor: "C:\Program Files\Notepad++\notepad++.exe"
14 Confirm: True
15 DisplayLog: True
16 <!--NeedCopy-->
```

## Populating zone mapping file

An on-premises zone is the equivalent of the cloud resource location. Unlike other site components, you cannot import the on-premises zone to the cloud automatically. Instead, it must be manually

mapped using the ZoneMapping.yml file. Import failures can occur if the zone name is not associated with an existing resource location name.

For on-premises sites having only one zone and cloud sites only one resource location, the Automated Configuration tool makes the correct association, eliminating the need to manually manage the ZoneMapping.yml file.

For on-premises sites having multiple zones or cloud sites having multiple resource locations, the ZoneMapping.yml file must be manually updated to reflect the correct mapping of on-premises zones to cloud resource locations. This must be done prior to attempting any import operation to the cloud.

The ZoneMapping.yml file is located in `%HOMEPATH%\Documents\Citrix\AutoConfig`. The content of the .yml file is a dictionary with the zone name as the key and the resource location name as the value.

As an example, an on-premises Citrix Virtual Apps and Desktops site with a primary zone called “Zone-1” and a secondary zone called “Zone-2” is migrated to a Citrix Virtual Apps and Desktops cloud deployment with two newly created cloud resource locations called “Cloud-RL-1” and “Cloud-RL-2”. In this instance, the ZoneMapping.yml would be configured as follows:

```
1 ---
2 Zone-1: Cloud-RL-1
3
4 Zone-2: Cloud-RL-2
5 <!--NeedCopy-->
```

**Note:**

A space must be between the colon and resource location name. If spaces are used in the zone or resource location name, enclose the name with quotes.

## Migrating from on-premises to cloud

Automated Configuration allows you to automate moving your on-premises configuration to a cloud site.

### Migration high level view

Migration is the process of exporting data from a source followed by importing that same data into a target. This is what Automated Configuration does. For customers migrating to the cloud from an on-premises environment, the source is their on-premises site and the target their cloud site. In its simplest form, the first-time-use migration process is the following steps, with example cmdlets.

1. Download and install the Automated Configuration tool.
2. Export the on-premises data (`Export-CvadAcToFile`).
3. Obtain the needed security information to access your cloud site.

4. Create the customer info file containing the security information used by the Automated Configuration tool (`New-CvadAcCustomerInfoFile`).
5. Import the on-premises data into the cloud (`Import-CvadAcToSite`, `Merge-CvadAcToSite`, `New-CvadAcToSite`).

Successive migrations are even simpler.

1. Export the on-premises data (`Export-CvadAcToFile`).
2. Import the on-premises data into the cloud (`Import-CvadAcToSite`, `Merge-CvadAcToSite`, `New-CvadAcToSite`).

The above can even be reduced to one step.

1. Export and import (`Sync-CvadAcSiteToSite`).

### Exporting your Citrix Virtual Apps and Desktops on-premises configuration

#### Important:

- You must have your `CustomerInfo.yml` file with your customer ID, client ID, and the secret key information included. For more information on how to retrieve your customer ID, client ID and secret key, see [Generating the customer ID, client ID, and secret key](#). For information on how to add this information to the `CustomerInfo.yml` file, see [Populating customer info file](#).
- The `ZoneMapping.yml` file must include information that maps your on-premises zone to Resource Locations in the cloud. For more information on how to map your zones, see [Populating zone mapping file](#).
- If you have host connections, you must input the corresponding info in the `CvadAcSecurity.yml` file.

1. Install Automated Configuration.
2. Double-click the **Auto Config** icon. A PowerShell window appears.
3. Run the following command to export all components.

```
Export-CvadAcToFile
```

After you run any cmdlet for the first the time, an export folder with the `.yml` configuration files and logs is created. The folder is at `%HOMEPATH%\Documents\Citrix\AutoConfig`. Each successive export creates a subfolder. The parent folder `%HOMEPATH%\Documents\Citrix\AutoConfig` always contains the exported files from the most recent export.

#### Note:

If Automated Configuration is not installed on the Delivery Controller, run `import-module Citrix.AutoConfig.Commands` before using the tool through PowerShell. This is not needed

if you open Automated Configuration using the **Auto Config** icon.

If you encounter any errors or exceptions, see the **Fixups** section in the log file.

## Importing your configuration to Citrix Virtual Apps and Desktops Service

### Important:

- You must have your CustomerInfo.yml file with your customer ID, client ID, and the secret key information included. For more information on how to retrieve your customer ID, client ID and secret key, see [Generating the customer ID, client ID, and secret key](#). For information on how to add this information to the CustomerInfo.yml file, see [Populating customer info file](#).
- The ZoneMapping.yml file must include information that maps your on-premises zone to resource locations in the cloud. For more information on how to map your zones, see [Populating zone mapping file](#).
- If you have host connections, you must input the corresponding info in the CvadAcSecurity.yml file.

## Running an import

1. Double-click the **Auto Config** icon. A PowerShell window appears.
2. Run the following command to import all components.

```
Merge-CvadAcToSite
```

Verify the expected state with the new current state. Various import options control whether the import results are identical or a subset of the on-premises site. See [Derived state results by command](#) for an explanation of the options and resultant state.

After you run the cmdlet, an export folder with the .yml configuration files and logs is created. The folder is at `%HOMEPATH%\Documents\Citrix\AutoConfig`.

If you encounter any errors or exceptions, see the **Fixups** section in the log file.

### Note:

If Automated Configuration is not installed on the Delivery Controller, run `import-module Citrix.AutoConfig.Commands` before using the tool through PowerShell. This is not needed if you open Automated Configuration using the **Auto Config** icon.

To revert to your original Citrix Virtual Apps and Desktops Service configuration, see [Backing up your Citrix Virtual Apps and Desktops Service configuration](#).

## Import operation in detail

The import process is designed to accurately perform updates, only perform needed updates and verify that all updates have been correctly made. The steps followed in all import operations follow.

1. Read the exported .yml file (expected state).
2. Read the cloud (current state).
3. Back up the pre-import cloud state to .yml files (pre-backup can be restored if necessary).
4. Evaluate the differences between the expected and current state. This determines which updates to make.
5. Make the updates.
6. Reread the cloud (new current state).
7. Back up the post-import cloud state to .yml files (post-backup can be restored if necessary).
8. Compare the new current state with the expected state.
9. Report the results of the comparison.

## Activating sites

### Note:

This feature is available only on versions 2.0 and later. Check your version by using `Get-CvadAcStatus` within Automated Configuration.

The Delivery Controller in both on-premises and cloud sites control resources such as brokering desktops, applications, and rebooting machines. Problems occur when a common set of resources is controlled by two or more sites. Such a situation can occur when migrating from an on-premises site to a cloud site. It is possible for both the on-premises and cloud Delivery Controllers to manage the same set of resources. Such dual management can lead to resources becoming unavailable and unmanageable, and can be difficult to diagnose.

Site activation allows you to control where the active site is controlled.

Site activation is managed using the delivery group maintenance mode. Delivery groups are placed in maintenance mode when the site is inactive. Maintenance mode is removed from delivery groups for sites that are active.

- `Set-CvadAdSiteActiveStateCloud`
- `Set-CvadAdSiteActiveStateOnPrem`

All cmdlets support the `IncludeByName` and `ExcludeByName` filtering.

## Import and transferring control to the cloud

The following is a high level description on how to import and transfer control from the on-prem site to the cloud site.

1. Export and import the on-premises site to the cloud. Make sure the `-SiteActive` parameter is not present on any of the import cmdlets. The on-premises site is active and the cloud site inactive. By default, cloud site delivery groups are in maintenance mode.
2. Verify the cloud content and configuration.
3. During off hours, set the on-premises site to inactive. The `-SiteActive` parameter must be absent. All on-premises site delivery groups are in maintenance mode.
  - `Set-CvadAcSiteActiveStateOnPrem`
4. Set the cloud site to active. The `-SiteActive` parameter must be present. No cloud site delivery groups are in maintenance mode.
  - `Set-CvadAcSiteActiveStateCloud -SiteActive`
5. Verify that the cloud site is active and the on-premises site is inactive.

### Transferring control back to the on-premises site

To transfer control from the cloud site to the on-premises site:

1. During off hours, set the cloud site to inactive. All cloud site delivery groups are in maintenance mode.
  - `Set-CvadAcSiteActiveStateCloud`
2. Set the on-premises site to active. No on-prem site delivery groups are in maintenance mode.
  - `Set-CvadAcSiteActiveStateOnPrem -SiteActive`

### Migrating from cloud to cloud

Automated Configuration allows you to automate moving your cloud configuration to another cloud site or allowing you to restore your own cloud site.

Moving your Citrix Virtual Apps and Desktops Service configuration can be beneficial for many reasons:

- Migrating from one region to another
- Syncing your site from test or stage to production
- Reaching resource limits

### Backing up your Citrix Virtual Apps and Desktops Service configuration

**Note:**

Before you begin, follow the **Import** steps in Prerequisites for migrating your configuration to migrate your configuration from one cloud to another.

**Important:**

- You must have your CustomerInfo.yml file with your customer ID, client ID, and the secret key information included. For more information on how to retrieve your customer ID, client ID and secret key, see [Generating the customer ID, client ID, and secret key](#). For information on how to add this information to the CustomerInfo.yml file, see [Populating customer info file](#).
- The ZoneMapping.yml file must include information that maps your resource locations in the cloud. For more information on how to map your zones, see [Populating zone mapping file](#).
- If you have host connections, you must input the corresponding info in the CvadAcSecurity.yml file.

1. Install Automated Configuration.

**Note:**

For cloud-to-cloud migration, Automated Configuration can be installed on a machine having access to the internet that the administrator has direct access to.

2. Double-click the **Auto Config** icon. A PowerShell window appears.

3. Run the following command to do a backup.

```
Backup-CvadAcToFile
```

After you run any cmdlet for the first time, an export folder with the .yml configuration files and logs is created. The folder is at `%HOMEPATH%\Documents\Citrix\AutoConfig`.

If you encounter any errors or exceptions, see the **Fixups** section in the log file.

## Restoring your configuration to Citrix Virtual Apps and Desktops Service

**Note:**

This section follows [Backing up your Citrix Virtual Apps and Desktops Service configuration and the steps documented there](#). Follow those steps before beginning the restore.

### Running a restore

1. Double-click the **Auto Config** icon. A PowerShell window appears.

2. Run the following command to do a restore.

```
Restore-CvadAcToSite -RestoreFolder <folder path of the backup files>
```

Verify the expected state with the new current state.

After you run the cmdlet, an export folder with the .yml configuration files and logs is created. The folder is at `%HOMEPATH%\Documents\Citrix\AutoConfig`.

If you encounter any errors or exceptions, see the **Fixups** section in the log file.

## Activating sites

Site activation allows you to control where the active site is controlled. For more information on activating sites, see [Activating sites](#).

## Backup and restore

The backup and restore process protects you from unintentional cloud site configuration changes or corruption. While Automated Configuration makes backups each time a change is made, this backup reflects the state of the cloud site configuration before the changes. Protecting yourself requires that you periodically back up your cloud site configuration and saving it in a safe place. If an undesirable change or corruption takes place, the backup can be used to fix the change or corruption at either a granular or full site configuration level.

### Backup

To back up, run the backup cmdlet: `Backup-CvadAcToFile`

While individual components can be selected when backing up, Citrix suggests backing up all components. Backups are placed in a uniquely named folder under the root folder: `%HOMEPATH%\Documents\Citrix\AutoConfig\Backup_yyyy_mm_dd_hh_mm_ss`

### Restore

Restore can be done from any backup to restore either one or more component members, one or more components or the full cloud site configuration.

### Restoring Component Members

Restoring one or more component members makes use of the `IncludeByName` feature. The `Restore` cmdlet is invoked with the `RestoreFrom` parameter along with the selected single component and the inclusion list.

To restore two group policies from a backup, follow this example:

```
Restore-CvadAcToSite -RestoreFrom %HOMEPATH%\Documents\Citrix\AutoConfig/  
Backup_yyyy_mm_dd_hh_mm_ss  
  
-GroupPolicies $true -IncludeByName Policy1,Policy2
```

### Restoring Entire Components

Restoring one component involves selecting one or more component parameters.



To restore the entire delivery group and machine catalog components, follow this example:

```
Restore-CvadAcToSite -RestoreFrom %HOMEPATH%\Documents\Citrix\AutoConfig/  
Backup_yyyy_mm_dd_hh_mm_ss  
  
-DeliveryGroups $true -MachineCatalogs $true
```

### Restoring the entire cloud site configuration

Restoring the full cloud site configuration means selecting all components to restore.

To restore the entire cloud site configuration, follow this example:

```
Restore-CvadAcToSite -RestoreFrom %HOMEPATH%\Documents\Citrix\AutoConfig/  
Backup_yyyy_mm_dd_hh_mm_ss
```

### Merging multiple sites into a single site

**Note:**

This feature is available only on versions 2.0 and later. Check your version by using `Get-CvadAcStatus` within Automated Configuration.

Multi-site support provides a method to merge multiple on-premises sites into a single cloud site.

Multi-site support adds unique prefixes and suffixes to component names on a per on-premises site basis, ensuring name uniqueness after multiple on-premises sites are merged to a single cloud site.

Prefixes and suffixes can be assigned for each of the following components on a per-on-premises-site-basis.

- AdminScope
- AdminRole
- ApplicationAdmin
- ApplicationFolder
- ApplicationGroup
- ApplicationUser
- DeliveryGroup
- GroupPolicy
- HostConnection
- MachineCatalog
- StoreFront
- Tag

Application folders support prefixing, suffixing, and rerooting. Rerooting adds an extra top level folder to an application's existing folder structure.

### Prefixing and suffixing rules

1. Prefixes and suffixes cannot contain any of the following special characters: \ , / ; : ## . \* ? = < > | ( ) "' { } [ ]
2. Prefixes and suffixes can contain trailing spaces but not leading spaces.
3. Prefixes and suffixes must be double quoted to contain trailing spaces.
4. Prefixes and suffixes are applied at the time of import, merge, and add. The source .yml files are never modified.
5. The prefix and suffix process automatically prefixes or suffixes dependent component names when applicable. For example, if machine catalog names are prefixed with “East,” delivery groups referencing them are also prefixed with “East.”
6. If a component name already begins with the prefix or suffix, no prefix or suffix is added. Component names cannot contain double identical prefixes or suffixes.
7. Prefixes and suffixes can be individually used or used in combination.
8. Use of a prefix or a suffix on a component is optional.

#### Note:

The Full Configuration interface displays components in alphabetical order.

### Group by site

Use prefixing to visually group components from a single site. Each site is listed in its own group with prefixing alphabetically controlling the ordering of different site groups.

### Group by name

Use suffixing to visually group like-named components from multiple sites. Like-named components from different sites visually alternate.

### SitePrefixes.yml file

Site prefixing begins with the SiteMerging.yml file that contains the site prefix and suffix mapping for one or more on-premises sites. You can manage the SiteMerging.yml file manually, or by using the available cmdlets listed at the Site merging cmdlets section.

### Exporting, importing, merging, and adding

Merging cannot begin until you have exported an on-premises site. To export an on-premises site, see Migrating from on-premises to cloud.

### Central export target folder

The methods described in this section place multiple site exports into a central file share location. The SiteMerging.yml file, CustomerInfo.yml file, and all export files reside in that file share location, allowing you to do the import from one location independent of the on-premises sites.

Cloud accessing operations never reference the on-premises sites or Active Directory, therefore allowing you to do cloud-accessing operations from anywhere.

### Direct file share

The export, import, merge, and new/add operations provide a parameter to target or source a folder other than the default folder, %HOMEPATH%\Documents\Citrix\AutoConfig. The following examples use a central file share located at \\share.central.net that the admin already has access to, having provided credentials as needed.

To target the export to a site-specific folder, use the `-TargetFolder` parameter:

From the East DDC:

```
mkdir \\share.central.net\AutoConfig\SiteEast
```

```
Export-CvadAcToFile -TargetFolder \\share.central.net\AutoConfig\SiteEast
```

From the West DDC:

```
mkdir \\share.central.net\AutoConfig\SiteWest
```

```
Export-CvadAcToFile -TargetFolder \\share.central.net\AutoConfig\SiteWest
```

After the exports are complete, create the CustomerInfo.yml and SiteMerging.yml files and place them in \\share.central.net\AutoConfig.

#### Note:

Do not use the `SiteRootFolder` parameter when creating the SitePrefixes.yml when using this direct file share reference method.

To import, merge, or add from the direct file share, you must decide from which machine you want to do the cloud accessing operation. Options include:

- One of the on-premises DDCs where the tool is already installed.
- The machine hosting the file share.
- A different machine.

Automated Configuration must be installed on the machine accessing the cloud. Neither the on-premises PowerShell SDK, DDC, nor Active Directory are used, so the cloud accessing execution requirements are simpler than the export requirements.

To merge the East DDC to the cloud:

```
Merge-CvadAcToSite -SiteName East -SourceFolder \\share.central.net\AutoConfig
\SiteEast -CustomerInfoFileSpec \\share.central.net\AutoConfig\CustomerInfo
.yml
```

To merge the West DDC to the cloud:

```
Merge-CvadAcToSite -SiteName West -SourceFolder \\share.central.net\AutoConfig
\SiteWest -CustomerInfoFileSpec \\share.central.net\AutoConfig\CustomerInfo
.yml
```

The following is a sample SitePrefixes.yml file used in the previous example.

```
1 East:
2   SiteRootFolder: "" # Important: leave this empty
3   AdminScopePrefix: "East_"
4   AdminRolePrefix: "East_"
5   ApplicationAdminPrefix: "East_"
6   ApplicationFolderPrefix: "" # Note that a new parent root folder is
   used instead
7   ApplicationFolderRoot: "East"
8   ApplicationGroupPrefix: "East_"
9   ApplicationUserPrefix: "East_"
10  DeliveryGroupPrefix: "East_"
11  GroupPolicyPrefix: "East_"
12  HostConnectionPrefix: "East_"
13  MachineCatalogPrefix: "East_"
14  StoreFrontPrefix: "East_"
15  TagPrefix: "East_"
16  AdminScopeSuffix: "_east"
17  AdminRoleSuffix: "_east"
18  ApplicationAdminSuffix: "_east"
19  ApplicationFolderSuffix: "_east"
20  ApplicationGroupSuffix: "_east"
21  ApplicationUserSuffix: "_east"
22  DeliveryGroupSuffix: "_east"
23  GroupPolicySuffix: "_east"
24  HostConnectionSuffix: "_east"
25  MachineCatalogSuffix: "_east"
26  StoreFrontSuffix: "_east"
27  TagSuffix: "_east"
28 West:
29   SiteRootFolder: "" # Important: leave this empty
30   AdminScopePrefix: "Western "
31   AdminRolePrefix: "Western "
32   ApplicationAdminPrefix: "Western "
33   ApplicationFolderPrefix: "" # Note that a new parent root folder is
```

```

        used instead
34   ApplicationFolderRoot: "Western"
35   ApplicationGroupPrefix: "Western "
36   ApplicationUserPrefix: "Western "
37   DeliveryGroupPrefix: "Western "
38   GroupPolicyPrefix: "Western "
39   HostConnectionPrefix: "Western "
40   MachineCatalogPrefix: "Western "
41   StoreFrontPrefix: "Western "
42   TagPrefix: "Western "
43   AdminScopeSuffix: ""
44   AdminRoleSuffix: ""
45   ApplicationAdminSuffix: ""
46   ApplicationFolderSuffix: ""
47   ApplicationGroupSuffix: ""
48   ApplicationUserSuffix: ""
49   DeliveryGroupSuffix: ""
50   GroupPolicySuffix: ""
51   HostConnectionSuffix: ""
52   MachineCatalogSuffix: ""
53   StoreFrontSuffix: ""
54   TagSuffix: ""
55
56 <!--NeedCopy-->

```

### File share reference using SiteMerging.yml

This method uses the `SiteRootFolder` member of the site's prefixes set. While more involved than the direct file share method, this method reduces the odds of targeting the wrong folder when exporting, importing, merging, or adding.

First, set the `SiteRootFolder` for each site in the `SiteMerging.yml` file. You must do this on the shared location.

```

Set-CvadAcSitePrefixes -SiteName East -SiteRootFolder \\share.central.net\
AutoConfig\SiteEast -SitePrefixesFolder \\share.central.net\AutoConfig
Set-CvadAcSitePrefixes -SiteName West -SiteRootFolder SiteWest -SitePrefixesFolder
\\share.central.net\AutoConfig

```

In this example, East is a fully qualified folder specification and West is a relative folder specification.

To target the export to a site-specific folder using the `SiteMerging.yml` file:

From the East DDC:

```
mkdir \\share.central.net\AutoConfig\SiteEast
```

```
Export-CvadaCToFile -SiteName East -CustomerInfoFileSpec \\share.central.net\AutoConfig\CustomerInfo.yml
```

From the West DDC:

```
mkdir \\share.central.net\AutoConfig\SiteWest
```

```
Export-CvadaCToFile -SiteName West -CustomerInfoFileSpec \\share.central.net\AutoConfig\CustomerInfo.yml
```

The export cmdlet uses the CustomerInfo.yml folder location to locate the SiteMerging.yml file. In the case of East, the `SiteRootFolder` is fully qualified. It is used as-is. In the case of West, the `SiteRootFolder` is not fully qualified. It is combined with the CustomerInfo.yml folder location to retrieve a fully qualified folder location for West.

To merge the East DDC to the cloud:

```
Merge-CvadaCToSite -SiteName East -CustomerInfoFileSpec \\share.central.net\AutoConfig\CustomerInfo.yml
```

To merge the West DDC to the cloud:

```
Merge-CvadaCToSite -SiteName West -CustomerInfoFileSpec \\share.central.net\AutoConfig\CustomerInfo.yml
```

The following is a sample SitePrefixes.yml file used in the previous example.

```
1 East:
2   SiteRootFolder: "\\share.central.net\AutoConfig\SiteEast"
3   AdminScopePrefix: "East_"
4   AdminRolePrefix: "East_"
5   ApplicationAdminPrefix: "East_"
6   ApplicationFolderPrefix: "" # Note that a new parent root folder is
   used instead
7   ApplicationFolderRoot: "East"
8   ApplicationGroupPrefix: "East_"
9   ApplicationUserPrefix: "East_"
10  DeliveryGroupPrefix: "East_"
11  GroupPolicyPrefix: "East_"
12  HostConnectionPrefix: "East_"
13  MachineCatalogPrefix: "East_"
14  StoreFrontPrefix: "East_"
15  TagPrefix: "East_"
16  AdminScopeSuffix: "_east"
17  AdminRoleSuffix: "_east"
18  ApplicationAdminSuffix: "_east"
19  ApplicationFolderSuffix: "_east"
20  ApplicationGroupSuffix: "_east"
```

```
21 ApplicationUserSuffix: "_east"
22 DeliveryGroupSuffix: "_east"
23 GroupPolicySuffix: "_east"
24 HostConnectionSuffix: "_east"
25 MachineCatalogSuffix: "_east"
26 StoreFrontSuffix: "_east"
27 TagSuffix: "_east"
28 West:
29 SiteRootFolder: "\\share.central.net\AutoConfig\SiteWest"
30 AdminScopePrefix: "Western "
31 AdminRolePrefix: "Western "
32 ApplicationAdminPrefix: "Western "
33 ApplicationFolderPrefix: "" # Note that a new parent root folder is
    used instead
34 ApplicationFolderRoot: "Western"
35 ApplicationGroupPrefix: "Western "
36 ApplicationUserPrefix: "Western "
37 DeliveryGroupPrefix: "Western "
38 GroupPolicyPrefix: "Western "
39 HostConnectionPrefix: "Western "
40 MachineCatalogPrefix: "Western "
41 StoreFrontPrefix: "Western "
42 TagPrefix: "Western "
43 AdminScopeSuffix: ""
44 AdminRoleSuffix: ""
45 ApplicationAdminSuffix: ""
46 ApplicationFolderSuffix: ""
47 ApplicationGroupSuffix: ""
48 ApplicationUserSuffix: ""
49 DeliveryGroupSuffix: ""
50 GroupPolicySuffix: ""
51 HostConnectionSuffix: ""
52 MachineCatalogSuffix: ""
53 StoreFrontSuffix: ""
54 TagSuffix: ""
55
56 <!--NeedCopy-->
```

If a central file share method is not used and the import, merge, or add is done from the individual DDCs, then create and replicate the SiteMerging.yml file on each DDC being migrated into the cloud. The default location is %HOMEPATH%\Documents\Citrix\AutoConfig. You must specify the - SiteName parameter to select the correct site prefixes.

## Merging the sites

Citrix recommends performing the cloud operations in steps and to do a complete review of each result before doing the next cloud operation. For example, if merging three sites to a single cloud site:

1. Merge the initial site to the cloud using the appropriate `SiteName` value.
2. Review the results in Studio (on-premises or web.)
3. If the results are incorrect, determine the issue and its cause, correct it, and then rerun the merge. If necessary, remove the cloud components and start from scratch by using `Remove-CvadAcFromSite` for the selected component and members. If the results are correct, continue.
4. If the initial merge is correct, merge the second site to the single cloud site.
5. Repeat steps 2 and 3.
6. If the second merge is correct, merge the third site to the single cloud site.
7. Repeat steps 2 and 3.
8. Review the resources from the user's perspective and verify that the view is in the desired state.

## Remove a component using the site prefix

You can selectively remove single site components by using the prefix on the `-IncludeByName` parameter of the `Remove-CvadAcFromSite` cmdlet. In the following example, the West DDC delivery groups are not correct. To remove the delivery groups for just the West site:

```
Remove-CvadAcFromSite -DeliveryGroups -IncludeByName "Western *"
```

To remove all West components, run the following cmdlets in order.

```
Remove-CvadAcFromSite -GroupPolicies -IncludeByName "Western *"
```

```
Remove-CvadAcFromSite -Applications -IncludeByName "Western *"
```

```
Remove-CvadAcFromSite - ApplicationGroups -IncludeByName "Western *"
```

```
Remove-CvadAcFromSite -DeliveryGroups -IncludeByName "Western *"
```

```
Remove-CvadAcFromSite -MachineCatalogs -IncludeByName "Western *"
```

```
Remove-CvadAcFromSite -HostConnections -IncludeByName "Western *"
```

```
Remove-CvadAcFromSite -Tags -IncludeByName "Western *"
```

To remove group policies of the East components, use the suffix:

```
Remove-CvadAcFromSite -GroupPolicies -IncludeByName "*_east"
```

## Automated Configuration tool cmdlets

- Site management cmdlets



- Customer info file cmdlets
- Support and troubleshooting cmdlets
- Site activation cmdlets
- Site merging cmdlets
- Component selection parameters
- Filtering by object names
- Migration parameters
- Cloud-accessing parameters
- Log display parameters
- Derived state results by command
- Cmdlet return values
- Help

### Site management cmdlets

- [Export-CvadAcToFile](#)

Exports configuration from your on-premises setup. This is the default export operation for Automated Configuration. No modifications are made to the on-premises site configuration. Exported files are placed in the directory `%HOMEPATH%\Documents\Citrix\AutoConfig` in a uniquely named **Export** subfolder. The folder `%HOMEPATH%\Documents\Citrix\AutoConfig` always contains the latest exported on-premises site configuration.

Parameters:

- See Component selection parameters
- See Filtering by object names
- [TargetFolder](#) - specifies the export destination folder.
- [Locale](#) - specifies the language of human-readable text that can be exported.
- [Quiet](#) - suppress logging to the console (available only on version 2.0 or later).
- [AdminAddress](#) - specifies the Delivery Controller's DNS or IP address when the export is not being run on the Delivery Controller.
- [CheckUserAndMachines](#) - verifies if users and machines are in Active Directory. Users and machines that are not in Active Directory might result in import failures.

Returns:

- See Cmdlet return values

- [Import-CvadAcToSite](#)

Imports all the on-premises files to the cloud. This command ensures that the cloud end state is identical to the on-premises state. This option deletes any changes that exist in the cloud. Imported site configuration files are sourced from `%HOMEPATH%\Documents\Citrix\AutoConfig`.

**Use with caution.**

Parameters:

- See Component selection parameters
- See Filtering by object names
- See Cloud-accessing parameters
- `SourceFolder` - identifies a substitute root folder for `%HOMEPATH%\Documents\Citrix\AutoConfig`.
- `Locale` - specifies the language of human-readable text that can be exported.
- `Quiet` - suppress logging to the console (available only on version 2.0 or later).
- `DisplayLog` - displays the log file at the completion of the cmdlet. Set to `$false` to suppress the log display.
- `Merge` - When set to `$true`, only adds components to the cloud site. Components are not removed. Set to `$false` to remove components.
- `AddOnly` - When set to `$true`, adds only new components, does not update or delete existing components. Set to `$false` to allow updates and deletions. `Merge` is ignored when this parameter is `$true`.

Returns:

- See Cmdlet return values

- [Merge-CvadAcToSite](#)

Merges the on-premises files to the cloud, but does not delete any components in the cloud. This preserves changes already made in the cloud. If a component exists in Citrix Cloud with the same name, this command can modify that component. This is the default import operation for Automated Configuration. Merged site configuration files are sourced from `%HOMEPATH%\Documents\Citrix\AutoConfig`.

Parameters:

- See Component selection parameters
- See Filtering by object names
- See Cloud-accessing parameters
- `SourceFolder` - identifies a substitute root folder for `%HOMEPATH%\Documents\Citrix\AutoConfig`.
- `Locale` - specifies the language of human-readable text that can be exported.
- `Quiet` - suppress logging to the console (available only on version 2.0 or later).
- `DisplayLog` - displays the log file at the completion of the cmdlet. Set to `$false` to suppress the log display.

- `AddOnly` – When set to `$true`, adds only new components, does not update or delete existing components. Set to `$false` to allow updates and deletions. `Merge` is ignored when this parameter is `$true`.

Returns:

- See Cmdlet return values

- `Restore-CvadAcToSite`

Restores the cloud site to the previous configuration. Imported files are sourced from the folder specified in the `-RestoreFolder` parameter. This can be used for reverting to your previous configuration or for backing up and restoring your cloud site.

Parameters:

- See Component selection parameters
- See Filtering by object names
- See Cloud-accessing parameters
- `RestoreFolder` - identifies the folder containing the .yml files to restore to the cloud site. This must be a fully qualified folder specification.
- `Locale` – specifies the language of human-readable text that can be exported.
- `Quiet` - suppress logging to the console (available only on version 2.0 or later).
- `DisplayLog` – displays the log file at the completion of the cmdlet. Set to `$false` to suppress the log display.
- `Merge` – When set to `$true`, only adds components to the cloud site. Components are not removed. Set to `$false` to remove components.
- `AddOnly` – When set to `$true`, adds only new components, does not update or delete existing components. Set to `$false` to allow updates and deletions. `Merge` is ignored when this parameter is `$true`.

Returns:

- See Cmdlet return values

- `New-CvadAcToSite`

Imports on-premises site configuration to the cloud but only adds new components. Existing cloud site components are neither updated nor deleted. Use this command if your existing cloud site components must remain unchanged.

Parameters:

- See Component selection parameters
- See Filtering by object names
- See Cloud-accessing parameters
- `SourceFolder` - identifies a substitute root folder for `%HOMEPATH%\Documents\Citrix\AutoConfig`.

- `Locale` – specifies the language of human-readable text that can be exported.
- `Quiet` - suppress logging to the console (available only on version 2.0 or later).
- `DisplayLog` – displays the log file at the completion of the cmdlet. Set to `$false` to suppress the log display.

Returns:

- See Cmdlet return values

- `Sync-CvadAcToSite`

Sync performs both an export and import in one step.

Parameters:

- See Component selection parameters
- See Filtering by object names
- See Cloud-accessing parameters
- `SourceTargetFolder` - specifies the export/import destination folder.
- `Locale` – specifies the language of human-readable text that can be exported.
- `AdminAddress` - specifies the Delivery Controller's DNS or IP address when the export is not being run on the Delivery Controller.
- `Quiet` - suppress logging to the console (available only on version 2.0 or later).
- `DisplayLog` – displays the log file at the completion of the cmdlet. Set to `$false` to suppress the log display.
- `Merge` – When set to `$true`, only adds components to the cloud site. Components are not removed. Set to `$false` to remove components.
- `AddOnly` – When set to `$true`, adds only new components, does not update or delete existing components. Set to `$false` to allow updates and deletions. `Merge` is ignored when this parameter is `$true`.

Returns:

- See Cmdlet return values

- `Backup-CvadAcToFile`

Exports your cloud configuration to .yaml files. This backup can be used in a backup and restore process to restore lost components.

Parameters:

- See Component selection parameters
- See Cloud-accessing parameters
- `TargetFolder` - specifies the export destination folder.
- `Locale` – specifies the language of human-readable text that can be exported.
- `Quiet` - suppress logging to the console (available only on version 2.0 or later).

- `DisplayLog` – displays the log file at the completion of the cmdlet. Set to `$false` to suppress the log display.

Returns:

- See Cmdlet return values

- `Compare-CvadAcToSite`

Compares the on-premises .yaml files with the cloud configuration, producing a report of changes that are made by an `Import`, `Merge`, or `Restore` cmdlet.

Parameters:

- See Component selection parameters
- See Filtering by object names
- See Cloud-accessing parameters
- `SourceFolder` - identifies a substitute root folder for `%HOMEPATH%\Documents\Citrix\AutoConfig`.
- `Locale` – specifies the language of human-readable text that can be exported.
- `Quiet` - suppress logging to the console (available only on version 2.0 or later).
- `DisplayLog` – displays the log file at the completion of the cmdlet. Set to `$false` to suppress the log display.
- `Merge` – When set to `$true`, only adds components to the cloud site. Components are not removed. Set to `$false` to remove components.
- `AddOnly` – When set to `$true`, adds only new components, does not update or delete existing components. Set to `$false` to allow updates and deletions. `Merge` is ignored when this parameter is `$true`.

Returns:

- See Cmdlet return values

- `Remove-CvadAcFromSite`

Can reset the entire site or remove member items from a component (For example, removing one machine catalog from the list of catalogs). This can be used when coupled with the `IncludeByName` parameter to selectively remove specific members.

Parameters:

- See Component selection parameters
- See Filtering by object names
- See Cloud-accessing parameters
- `Quiet` - suppress logging to the console (available only on version 2.0 or later).
- `DisplayLog` – displays the log file at the completion of the cmdlet. Set to `$false` to suppress the log display.

Returns:

- See Cmdlet return values

### Customer info file cmdlets

- `New-CvadAcCustomerInfoFile`

Create a customer info file. By default, the customer info file is located at `%HOMEPATH%\Documents\Citrix\AutoConfig`.

Parameters:

- `CustomerId` – customer’s ID (required).
- `ClientId` – customer’s client ID created on Citrix Cloud (required).
- `Secret` – customer’s secret key created on Citrix Cloud (required).
- `Environment` – Production or ProductionGov environment.
- `LogFileName` – Change the log file prefix from CitrixLog to something else.
- `StopOnError` – Stops the operation upon first error.
- `AlternateRootFolder` – Use the specified folder as the root folder instead of `%HOMEPATH%\Documents\Citrix\AutoConfig`.
- `Locale` – use the specified local instead of the locale derived from the system the tool is run on.
- `Editor` – use the specified editor to display the log at the completion of each cmdlet. Notepad.exe is the default editor. This parameter must include the fully qualified file specification to the editor and the editor must take the log file spec as its only parameter.

Returns:

- See Cmdlet return values

Example:

```
New-CvadAcCustomerInfoFile -CustomerId markhof123 -ClientId 6813EEA6-46CC-4F8A-BC71-539F2DAC5984 -Secret TwBLaaaaaaaaaaaaaaaaaaw==
```

- `Set-CvadAcCustomerInfoFile`

Update an existing customer info file. Only cmdlet specified parameters are changed, all unspecified parameter values in the CustomerInfo.yml file are unchanged.

Parameters:

- `CustomerId` – customer’s ID (required).
- `ClientId` – customer’s client ID created on Citrix Cloud (required).
- `Secret` – customer’s secret key created on Citrix Cloud (required).
- `Environment` – Production or ProductionGov environment.
- `LogFileName` – Change the log file prefix from CitrixLog to something else.
- `StopOnError` – Stops the operation upon first error.

- **AlternateRootFolder** – Use the specified folder as the root folder instead of `%HOMEPATH%\Documents\Citrix\AutoConfig`.
- **Locale** – use the specified local instead of the locale derived from the system the tool is run on.
- **Editor** – use the specified editor to display the log at the completion of each cmdlet. Notepad.exe is the default editor. This parameter must include the fully qualified file specification to the editor and the editor must take the log file spec as its only parameter.

Returns:

- See Cmdlet return values

### Support and troubleshooting cmdlets

- **New-CvadAcTemplateToFile**

Creates a template file for selected components, allowing you to manually create an import file.

Parameters:

- See Component selection parameters
- **TargetFolder** - specifies the export destination folder.

Returns:

- See Cmdlet return values

- **Show-CvadAcDocument** - Displays this document in the default browser.

Parameters:

- None.

Returns:

- Display this webpage in the default web browser.

- **Test-CvadAcConnectionWithSite**

Test the connection with the cloud site to verify that the communication connection is working. This cmdlet uses the cloud accessing parameters or the CustomerInfo.yml file to specify the customer connection information.

Parameters:

- See Cloud-accessing parameters
- **SiteId** - identifies the site to test for a valid connection.
- **Quiet** - suppress logging to the console (available only on version 2.0 or later).

Returns:

- Test results are displayed on the command line.

- **Find-CvadAcConnector**

Locates existing connectors and determines their running state. This cmdlet uses information from the CustomerInfo.yml file or the customer ID parameter to locate the customer's connectors.

Parameters:

- **CustomerInfoFileSpec** - The file specification pointing to a customer information file to override the default location and name. This parameter is ignored when the **CustomerId** parameter is provided.
- **CustomerId** - the customer's ID, this parameter overrides the same value in the CustomerInfo.yml file.

Returns:

- Results are shown on the command line.

- **New-CvadAcZipInfoForSupport**

**Note:**

The **New-CvadAcZipInfoForSupport** cmdlet is available only on versions 2.0 and later. Check your version by using **Get-CvadAcStatus** within Automated Configuration.

Zips all log and .yml files in a single zip file to send to Citrix for support. Customer sensitive information (CustomerInfo.yml and CvadAcSecurity.yml) is not included in the zip. The Icon.yml file is also excluded due to its size. The zip file is placed in *%HOMEPATH%\Documents\Citrix\AutoConfig* and named *CvadAcSupport\_YYYY\_MM\_DD\_HH\_MM\_SS.zip*, based on the date and timestamp. This zip file can also act as a backup.

Parameters:

- **AlternateFolder** - specifies an alternate folder to save the zip file to.
- **Quiet** - suppress logging to the console (available only on version 2.0 or later).

Returns:

- Zip file with zip file name and location is displayed on the command prompt.

- **Get-CvadAcCustomerSites**

Returns the list of all the customer sites. This cmdlet uses the cloud accessing parameters or the CustomerInfo.yml file to specify the customer connection information.

Parameters:

- See Cloud-accessing parameters

Returns:



- Displays a list of found customer site IDs.

- [Get-CvadAcStatus](#)

Returns information about the tool and the environment it is running in, including the tool version, whether the user has site administrator privileges for exporting, whether customer information (Customer ID, Client ID, and Secret) is present, whether the site can be communicated with, whether the Provisioning Services enhanced feature is present, and whether one or more connectors are present.

Parameters:

- See Cloud-accessing parameters
- [SiteId](#) - identifies the site to connect to (optional).
- [AdminAddress](#) - is the DNS or IP address of the on-premises Delivery Controller used to verify the admins access level. This is required if the tool is not being run on a Delivery Controller.

Returns:

- Displays the results for each item.

### Site activation cmdlets

**Note:**

This feature is available only on versions 2.0 and later. Check your version by using [Get-CvadAcStatus](#) within Automated Configuration.

- [Set-CvadAcSiteActiveStateOnPrem](#)

Sets the on-premises site state to either active or inactive.

Parameters:

- See Cloud-accessing parameters
- [SiteActive](#) - when present, sets the on-premises site to active removing the maintenance mode from all delivery groups. When this parameter is not present, maintenance mode is set on all delivery groups.
- [IncludeByName](#) - a list specifying the names of delivery groups to include when setting the site active state to active. The '\*' and '?' wildcards are supported in names.
- [ExcludeByName](#) - a list specifying the names of delivery groups to exclude when setting the site active state to active. The '\*' and '?' wildcards are supported in names.
- [Quiet](#) - suppress logging to the console
- [DisplayLog](#) - displays the log file at the completion of the cmdlet. Set to `$false` to suppress the log display.

Returns:

- See Cmdlet return values

- [Set-CvadAcSiteActiveStateCloud](#)

Sets the cloud site state to either active or inactive.

Parameters:

- See Cloud-accessing parameters
- [SiteActive](#) - when present, sets the on-premises site to active removing the maintenance mode from all delivery groups. When this parameter is not present, maintenance mode is set on all delivery groups.
- [IncludeByName](#) - a list specifying the names of delivery groups to include when setting the site active state to active. The '\*' and '?' wildcards are supported in names.
- [ExcludeByName](#) - a list specifying the names of delivery groups to exclude when setting the site active state to active. The '\*' and '?' wildcards are supported in names.
- [Quiet](#) - suppress logging to the console
- [DisplayLog](#) - displays the log file at the completion of the cmdlet. Set to `$false` to suppress the log display.

Returns:

- See Cmdlet return values

### Site merging cmdlets

**Note:**

This feature is available only on versions 2.0 and later. Check your version by using [Get-CvadAcStatus](#) within Automated Configuration.

For more information on site merging and usage of these cmdlets, see [Merging multiple sites into a single site](#).

- [New-CvadAcSiteMergingInfo](#) - Creates a site merging prefix/suffix info set. It is not necessary to know all prefixes or suffixes at the beginning. They can be updated with [Set-CvadAcSiteMergingInfo](#) or by manually editing the `SiteMerging.yml` file.
- [Set-CvadAcSiteMergingInfo](#) - Updates an existing site merging prefix/suffix info set.
- [Remove-CvadAcSiteMergingInfo](#) - Removes an existing site merging prefix/suffix info set.

### Site merging parameters

- [SiteName](#) - the name used to identify the set of prefixes/suffixes for a specific site. It can match the name of the actual site but does not need to. `SiteName` is a required parameter. **Note:** This is the only supported parameter for [Remove-CvadAcSiteMergingInfo](#).
- [Quiet](#) - suppress logging to the console (available only on version 2.0 or later).

- `AdminScopedPrefix` – the prefix to apply to administrator scopes.
- `ApplicationPrefix` – the prefix to apply to applications.
- `ApplicationFolderPrefix` – the prefix to apply to application folders; `ApplicationFolderPrefix` can be combined with `ApplicationFolderRoot`.
- `ApplicationFolderRoot` – the new root folder to application folders. This creates an extra folder hierarchy. `ApplicationFolderRoot` can be combined with `ApplicationFolderPrefix`.
- `ApplicationGroupPrefix` – the prefix for application groups.
- `ApplicationUserPrefix` – the prefix to apply to the application name the user sees.
- `ApplicationAdminPrefix` – the prefix to apply to the application name the administrator sees.
- `DeliveryGroupPrefix` – the prefix to apply to delivery groups.
- `GroupPolicyPrefix` – the prefix to apply to policy names.
- `HostConnectionPrefix` – the prefix to apply to host connections.
- `MachineCatalogPrefix` – the prefix to apply to machine catalogs.
- `StoreFrontPrefix` – the prefix to apply to StoreFront names.
- `TagPrefix` – the prefix to apply to tags.
- `AdminScopedSuffix` – the suffix to apply to administrator scopes.
- `ApplicationSuffix` – the suffix to apply to applications.
- `ApplicationFolderSuffix` – the suffix to apply to application folders; `ApplicationFolderSuffix` can be combined with `ApplicationFolderRoot`.
- `ApplicationGroupSuffix` – the suffix for application groups.
- `ApplicationUserSuffix` – the suffix to apply to the application name the user sees.
- `ApplicationAdminSuffix` – the suffix to apply to the application name the administrator sees.
- `DeliveryGroupSuffix` – the suffix to apply to delivery groups.
- `GroupPolicySuffix` – the suffix to apply to policy names.
- `HostConnectionSuffix` – the suffix to apply to host connections.
- `MachineCatalogSuffix` – the suffix to apply to machine catalogs.
- `StoreFrontSuffix` – the suffix to apply to StoreFront names.
- `TagSuffix` – the suffix to apply to tags.
- `SiteRootFolder` – the fully qualified folder name to use for exports and imports; this can be a local folder or a file share.

### Component selection parameters

The following components can be specified with cmdlets supporting them. For tool versions including 1.x, follow each parameter with `$true`. For tool versions 2.x or later, specify the parameter only. Selecting `All` causes the other parameters to be ignored. The `All` option is automatically selected when no component parameters are specified.

- All
- Tags
- AdminRolesScopes
- MachineCatalogs
- DeliveryGroups
- ApplicationGroups
- Applications
- GroupPolicies

## Filtering by object names

### Include and Exclude by name

The `IncludeByName` and `ExcludeByName` parameters enable including and excluding component members in cmdlets by name. Only one component (for example, delivery groups) can be chosen at a time in any of the supported cmdlets. If a component member is in both areas, exclude overrides any other parameter and an entry is made in the log fixup list identifying the component and member name that was excluded.

`IncludeByName` and `ExcludeByName` take a list of component member names. Any name can contain one or more wildcards. Two types of wildcards are supported. The list of component member names must be enclosed in single-quotes when any member name contains special characters.

- \* Matches any number of characters
- ? Matches a single character

`IncludeByName` and `ExcludeByName` can also take a file containing a list of members where each member can be explicit or contain wildcards. Each line in the file can contain one member. Leading and trailing spaces are trimmed from the member name. The file name must be preceded by the @ sign and be surrounded by single quotes (a PowerShell requirement so the @ is not reinterpreted). Multiple files can be listed in addition to being mixed with member names.

One example of merging all delivery groups whose names begin with `DgSite1` and contain `Home2` would be written:

```
Merge-CvadAcToSite -DeliveryGroups $true -IncludeByName 'DgSite1*,*Home2*
```

### By Delivery Group Name

`ByDeliveryGroupName` filters by the delivery group name for applications and application groups. This parameter is always an inclusion list identifying members to include based on their delivery group association.

`ByDeliveryGroupName` takes a list of delivery group names. Any name can contain one or more wildcards. Two types of wildcards are supported.

- \* matches any number of characters
- ? matches a single character

The following example merges all applications that reference all delivery group names beginning with `EastDg`.

```
Merge-CvadAcToSite -Applications -ByDeliveryGroupName EastDg*
```

### Exclude Disabled

`ExcludeDisabled` filters out from import operations all applications and application groups that are disabled. `ExcludeDisabled` defaults to **false**, meaning all applications and application groups are imported regardless of their enabled state.

### By Machine Name

#### Note:

This feature is available only on versions 2.0 and later. Check your version by using `Get-CvadAcStatus` within Automated Configuration.

`ByMachineName` filters by the machine name for machine catalogs and delivery groups. This parameter is always an inclusion list identifying members to include based on their machine name association.

`ByMachineName` takes a list of machine names where any name can contain one or more wildcards. Two types of wildcards are supported.

- \* matches any number of characters
- ? matches a single character

When exporting or importing and using `ByMachineName` and a machine name filter results in no machines in the machine catalog or delivery group, the machine catalog or delivery group is excluded from the export or import.

#### Note:

Use of `ByMachineName` in any import type cmdlet results in `AddMachinesOnly` being set to `$true`.

### Add Machines Only

`AddMachinesOnly`, when set to `$true`, instructs the import operation to add machines only to the machine catalog or delivery group. Machines are not removed, allowing for incremental additive operations.

`AddMachineOnly` defaults to false meaning machines are removed if they are not present in the machine catalog or delivery group .yaml file. `AddMachinesOnly` is set to `$true` when `ByMachineName` is used but can be overridden by setting `AddMachinesOnly` to false.

### Migration parameters

Cmdlets modifying the cloud site configuration (`Import`, `Restore`, `Merge`, `New`, and `Sync`) support the following extra parameters to provide further flexibility.

- `CheckMode` – Performs the import operation but makes **no** changes. All expected changes are reported before the import completes. You can use this command to test your import before doing the actual import.
- `BackupFirst` – Backs up the cloud contents to .yaml files before modifying the cloud configuration. This is enabled by default.
- `Confirm` – When true, prompts users to confirm that they want to make changes to the cloud site configuration. The `Remove` cmdlet shows a prompt due to its destructive nature. Set to false if no prompt is desired, such as running inside automated scripts. `Confirm` defaults to true.

#### Note:

The `SecurityFileFolder`, `SiteName`, and `SiteActive` cmdlets are available only on versions 2.0 and later. Check your version by using `Get-CvadAcStatus` within Automated Configuration.

- `SecurityFileFolder` – This is the fully qualified folder containing the `CvadSecurity.yaml` file which might point to a local folder or a network share folder that is under authentication control. The tool does not prompt for credentials; access to the controlled resource must be obtained before running the tool.
- `SiteName` – Specifies the site merging prefix and suffix set to use when importing. See [Merging multiple sites into a single site](#) for more information.
- `SiteActive` – Specifies whether the imported site is active or inactive. By default, this parameter is set to `$false` meaning the imported site is inactive.

### Cloud-accessing parameters

All cmdlets accessing the cloud support the following extra parameters.

#### Note:

The `CustomerId`, `ClientId`, and `Secret` can be placed in the `CustomerInfo.yaml` file or specified with the cmdlet using the preceding parameters. When they are specified in both places, the cmdlet parameters take precedence.

- **CustomerId** – The customer ID used in the Rest APIs and is required to access all Rest APIs. Your customer ID is found in Citrix Cloud.
- **ClientId** – The clientID created on the Citrix Cloud Identity and Access Management website. This is required to obtain the bearer token needed for authentication for all Rest APIs.
- **Secret** – The secret key created on the Citrix Cloud Identity and Access Management website. This is required to obtain the bearer token needed for authentication for all Rest APIs.
- **CustomerInfoFileSpec** – The file specification pointing to a customer information file to override the default location and name.

### Log display parameters

The **Export**, **Import**, **Sync**, **Restore**, **Backup**, **Compare**, and **Remove** cmdlets display the log file when the operation completes. You can suppress the display by setting the **-DisplayLog** parameter to **\$false**. Notepad.exe is used by default to display the log file. You can specify a different editor in the CustomerInfo.yml file.

Editor: C:\Program Files\Notepad++\notepad++.exe

### Derived state results by command

Derived state can result in one of the three combinations of actions on the cloud site:

- Add, Update, and Delete
- Add and Update only
- Add only

The following table shows the derived state for each cmdlet and optional parameters that can change the derived state behavior of the cmdlet.

Command	Add, Update, Delete	Add, Update	Add
Import	<code>-Merge \$false</code>	Default	<code>-AddOnly \$true</code>
Merge	N/A	Default	<code>-AddOnly \$true</code>
New	N/A	N/A	Default
Sync	<code>-Merge \$false</code>	Default	<code>-AddOnly \$true</code>
Restore	<code>-Merge \$false</code>	Default	<code>-AddOnly \$true</code>

### Cmdlet return values

#### ActionResult

All cmdlets return the following value.

```

1 public class ActionResult
2 {
3
4     public bool Overall_Success;
5     public Dictionary<string, string> Individual_Success;
6     public object CustomResult;
7 }
8
9 <!--NeedCopy-->

```

`Overall_Success` returns a single boolean showing the overall success of the cmdlet across all selected components: true meaning successful and false meaning unsuccessful.

`Individual_Success` returns one or three values for each main component. A component's result can be Success, Failure, or Skipped. Skipped indicates the component was not selected for execution by the cmdlet.

`CustomResult` is cmdlet specific.

### CustomResult

`Import`, `Merge`, `Restore`, `Sync`, `Compare`, `Compare File`, and `Remove` return the following custom result information to a single instance of `EvaluationResultData`.

#### Note:

`Export` and `Template` cmdlets do not return a custom result.

```

1 public class EvaluationResultData
2 {
3
4     public Dictionary<string, Dictionary<string, ActionResultValues
5         >> EvaluationResults;
6     public int Added;
7     public int Updated;
8     public int Deleted;
9     public int NoChange;
10    public int TotalChanged;
11    public EvaluationResults OverallResult;
12    public string CloudBackupFolder;
13    public string SourceBackupFolder;
14 }
15 Where:

```



```
16     public enum ActionResultValues
17     {
18
19         Add,
20         Update,
21         Delete,
22         Identical,
23         DoNothing
24     }
25
26     public enum EvaluationResults
27     {
28
29         Success,
30         Failure,
31         Skipped
32     }
33
34 <!--NeedCopy-->
```

`EvaluationResults` displays a list with one entry per selected component. The key is the component name and the value is a list of each component member and the action taken on that component member. Actions can be any one of the `ActionResultValues` values.

`Added`, `Updated`, `Deleted`, and `NoChange` indicate the total number of component members added, updated, deleted, or no action taken, in that order.

`TotalChanged` is the sum of `Added`, `Updated`, and `Deleted`.

`OverallResult` is a single boolean indicating the result of the cmdlet. True indicates total success across all components and false indicates failure in processing one or more components.

`CloudBackupFolder` is the fully qualified file specification of the cloud site configuration backup before the cmdlet performing any cloud-modifying actions.

`SourceBackupFolder` is the fully qualified file specification of the source file backup made after completion of the cmdlet. By default, these files are at `%HOMEPATH%\Documents\Citrix\AutoConfig`.

## Help

- `Get-Help` - PowerShell help is available for each cmdlet. All parameters are documented with each cmdlet along with a brief explanation of the cmdlet. To access help for any cmdlet, type `Get-Help` in front of the cmdlet. For example, `Get-Help Import-CvAdAcToSite`.

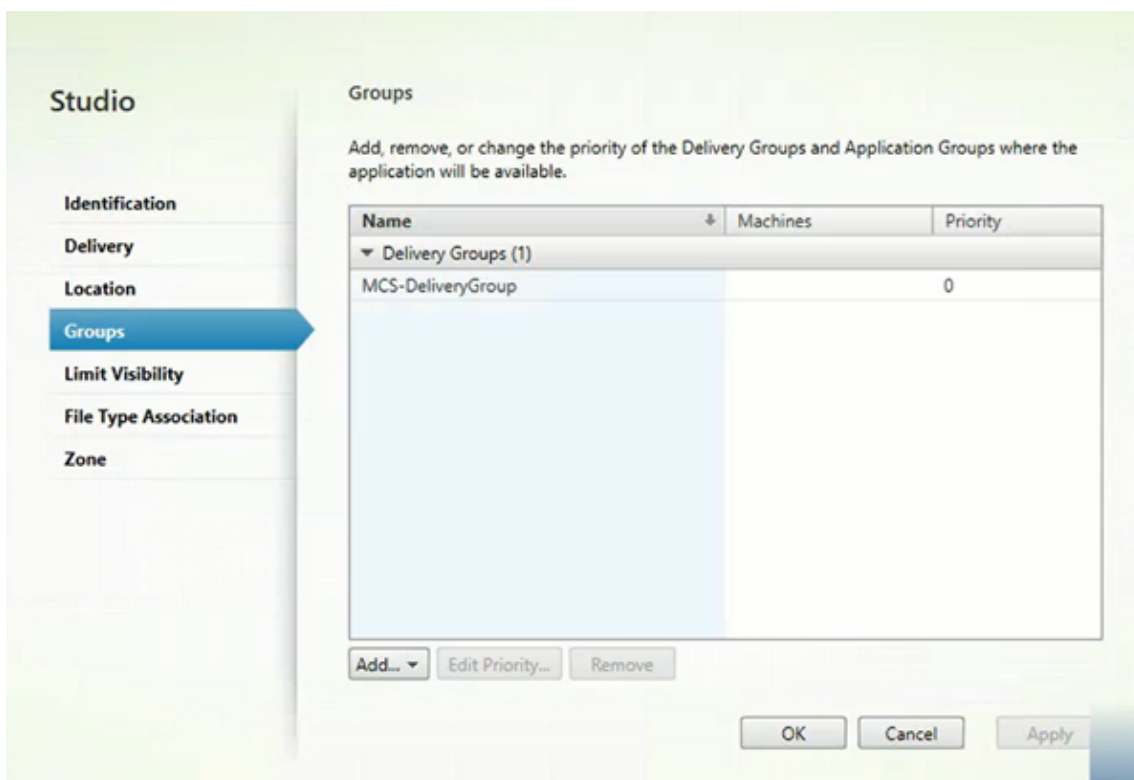
## Additional information

- Importing objects with MCS machine catalogs
- Citrix Cloud Government
- Host connections
- Automation
- Exporting from PCs other than the Delivery Controller
- Citrix Cloud data collection
- Folders
- Backups
- Logging
- Diagnostic files
- Dependency problems
- Common problems
- Recommendations
- Resources

## Importing other objects when you have Machine Creation Services machine catalogs


Currently, Automated Configuration does not support importing MCS machine catalogs or their corresponding delivery groups in an automated way. However, you can still import other configuration options such as applications and policies using Automated Configuration. You must create the machine catalog and delivery group using the same name as the on-premises setup. Follow the steps to prepare your environment, before proceeding to import the application settings:

1. In your Citrix Cloud portal, click the hamburger menu and go to **My Services > Virtual Apps and Desktops Service > Manage**. Create your MCS machine catalog as you normally would. Be sure to name your catalog **exactly** the same way your existing on-premises catalog is named.
2. Create the corresponding delivery group for the new catalog. Be sure you name it **exactly** after the corresponding on-premises delivery group.
3. In your on-premises environment's Citrix Studio, in the **Applications** node, confirm the applications belong to the matching delivery groups by selecting the application, right-clicking the app, and going to **Properties**.
4. Click **Groups** to confirm the groups the app belongs to:

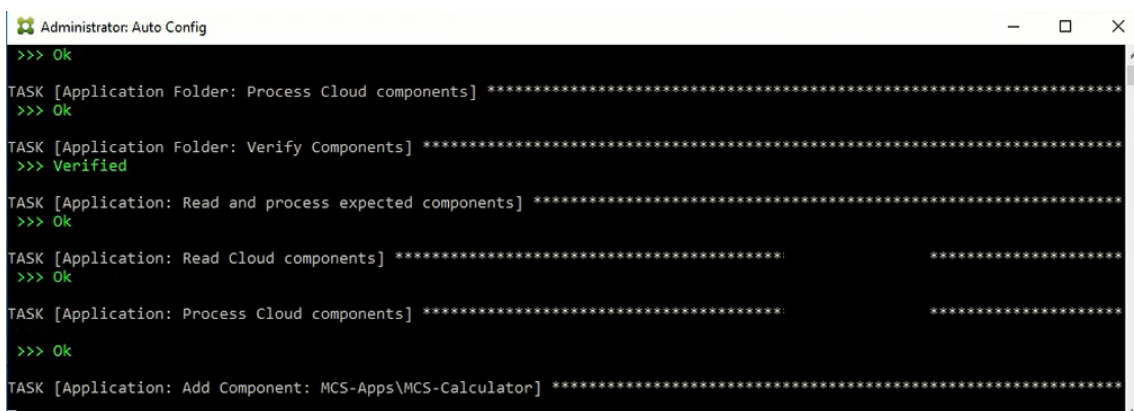


5. In PowerShell, run the **Merge** command and use the `byDeliveryGroupName` flag, which filters the applications by delivery group name.

`Merge-CvadAcToSite -Applications $true -ByDeliveryGroupName <DG_name>`



6. Type **Yes** to continue.



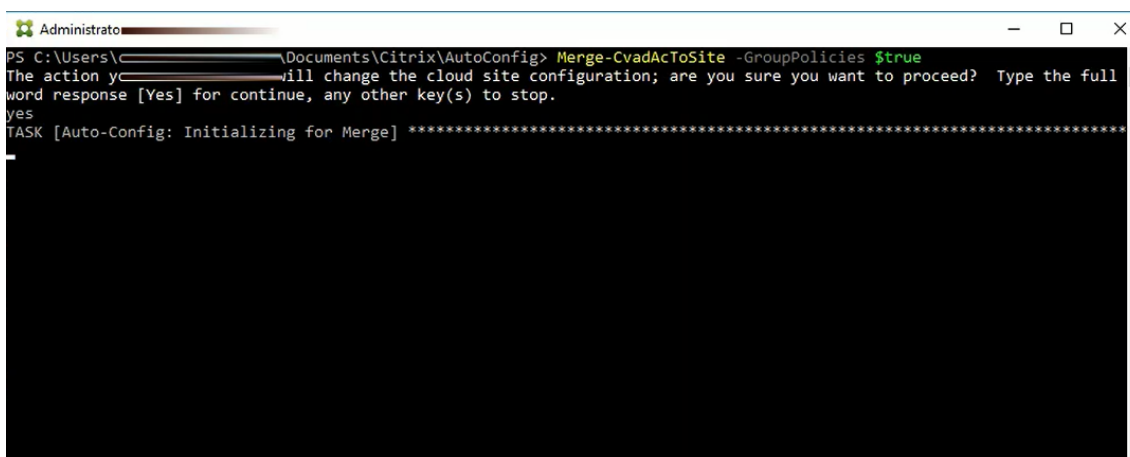
7. From **Manage > Full Configuration**, select **Applications** in the left pane.

Refresh the display to make sure the apps are listed as expected. Select the applications and select **Application Properties > Groups** to check.

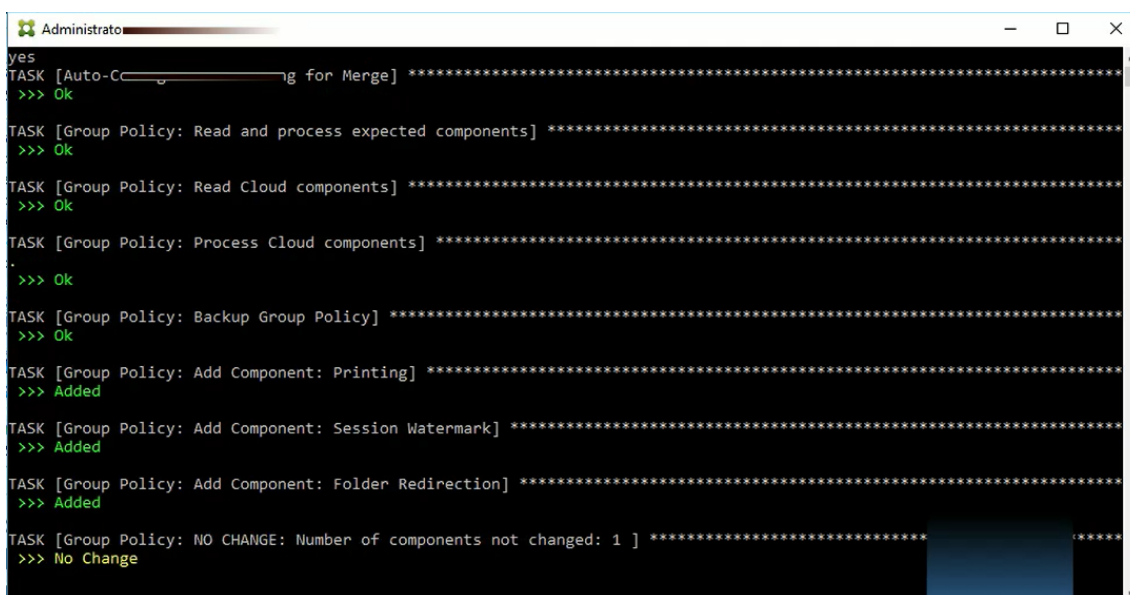
### Importing MCS-related policies

If you need to import policies associated with your MCS catalogs or groups, follow these instructions:

1. Run the `Merge-CvadAcToSite -GroupPolicies $true` command in PowerShell and type **yes** to continue.



Successful execution displays a similar output to the preceding screenshot (Added values). The following screenshot shows the result of a line for which there were no changes (No change).



2. Refresh the **Manage > Full Configuration** display and select **Policies** in the left pane.
3. Check the Policies **Assigned to** tab and compare it with your on-premises policy assignment.

## Moving to Citrix Cloud Government

The Citrix Cloud Government environment uses different access points to authenticate and allocate access tokens. This unique requirement applies to any Automated Configuration tool accessing the cloud. Perform the following steps to use Automated Configuration in Citrix Cloud Government environments.

1. In the `%HOMEPATH%\Documents\Citrix\AutoConfig` folder, edit `CustomerInfo.yml`.
2. Add the following line to `CustomerInfo.yml` (or change it, if already present.)

```
Environment: "ProductionGov"
```

Automated Configuration is now able to be used on Citrix Cloud Government environments.

## Host connections

Host connections and their associated hypervisors can be exported and imported using Automated Configuration.

Adding a hypervisor to a host connection requires security information specific to the type of hypervisor. This information cannot be exported from the on-premises site for security considerations. You must manually provide the information so that Automated Configuration can successfully import host connections and hypervisors to the cloud site.

The export process creates the `HostSecurity.yml` file in `%HOMEPATH%\Documents\Citrix\AutoConfig` containing placeholders for each security item needed for the specific hypervisor type. You must update the `HostSecurity.yml` file before importing into the cloud site. Administrator updates are retained over multiple exports with new security placeholders added as needed. Security items are never removed.

### Note:

Be sure to double quote all values entered in the `HostSecurity.yml` file.

```
1 ---
2
3 HostConn1:
4 ConnectionType: XenServer
5 UserName: root
6 PasswordKey: rootPassword
7 HostCon2:
8 ConnectionType: AWS
9 ApiKey: 78AB6083-EF60-4D26-B2L5-BZ35X00DA5CH
10 SecretKey: TwBLaaaaaaaaaaaaaaaaaw==
11 Region: East
12 <!--NeedCopy-->
```

## Per-hypervisor security information

The following lists the security information required for each hypervisor type.

- XenServer, Hyper-V, VMware
  - User Name
  - Clear-text Password
- Microsoft Azure
  - Subscription ID
  - Application ID
  - Application Secret
- Amazon Web Services
  - Service Account ID
  - Application Secret
  - Region

## Special security considerations

All security information is entered as clear text. If clear text is not recommended, the host connections and associated hypervisors can be manually created using the service's **Manage > Full Configuration** interface. The host connections and hypervisor names must match their on-premises counterparts exactly so that machine catalogs that use the host connections can be successfully imported.

## Automation

Automated Configuration tool cmdlets can be run in automation scripts without administrator intervention by suppressing prompts and the display of the log results at cmdlet completion. You can also set parameters to do the same by using the CustomerInfo.yml file.

Add the following parameter to cloud modifying cmdlets to suppress the display of prompts.

```
-Confirm $false
```

Add the following parameter to cmdlets to suppress the display of log at the completion of the cmdlet.

```
-DisplayLog $false
```

Add the following parameter to cmdlets to suppress logging to the PowerShell command window.

```
-Quiet
```

As another method, the following parameters can be placed in the CustomerInfo.yml file.

```
Confirm: False
```

```
DisplayLog: False
```

## Exporting from PCs other than the Delivery Controller

The Automated Configuration tool uses multiple Citrix PowerShell SDKs to export the on-premises site configuration to files. These SDKs are automatically installed on the Delivery Controller, enabling the tool to run on the Delivery Controller without extra actions. When running on non-Delivery Controller machines, it is necessary to install the set of Citrix PowerShell SDKs needed by the tool. This SDK set is part of Citrix Studio which can be installed from the Citrix Virtual Apps and Desktops installation media.

### Note:

Automated Configuration cannot be run on the Cloud Connector.

## Citrix Cloud data collection

For information on what information Citrix Cloud collects, see [Citrix Cloud Services Customer Content and Log Handling](#).

## Folders

### Default folder root location

All Automated Configuration tool operations occur in the root folder or in subfolders inside it. The root folder is located in `%HOMEPATH%\Documents\Citrix\AutoConfig`.

## Export

All exported files are placed in two folder locations, providing ease-of-use and a history of exports. Exports are always placed in the root folder. Copies are placed in a subfolder named **Export** with the date and time of the export.

The root folder always contains the most recent exported on-premises site configuration. Each **Export** subfolder contains the export done on the indicated date and time, which maintains a history of exports. You can use any **Export** subfolder to configure the cloud site. Automated Configuration does not delete or modify existing export subfolders.

## Import/Merge/Sync/Compare

[Import](#), [Merge](#), and [Compare](#) operations always sourced from files located in the root folder. Each operation results in the creation of a subfolder to which files in the root folder are copied, providing a history of cloud site changing source files.

## Restore

The `Restore` operation uses an existing subfolder to configure the cloud site. The source folder is specified on the required `-RestoreFolder` parameter. Unlike with other commands, no new subfolder is created because the `Restore` operation uses an existing subfolder. The restore folder can be the root folder but still must be specified on the `-RestoreFolder` parameter.

## Backups

Automated Configuration initializes, updates, and backs up a cloud site configuration. When used over time, many different configurations can change on the cloud site. To facilitate long-term use and preserve history changes, Automated Configuration uses a preservation scheme to save this history of changes and provide a method to restore earlier states.

Cloud site configuration backups are always made to a subfolder named **Backup** with the data and time of the backup. Automated Configuration does not delete or modify existing export subfolders.

You can use the backups to restore specific components or your entire configuration. To restore the entire delivery group and machine catalog components, use the cmdlet:

```
Restore-CvadAcToSite -RestoreFrom %HOMEPATH%\Documents\Citrix\AutoConfig/  
Backup_yyyy_mm_dd_hh_mm_ss -DeliveryGroups -MachineCatalogs
```

### Note:

The backup file information in the preceding cmdlet is based on your own backups.

To restore the entire cloud site configuration, use the cmdlet:

```
Restore-CvadAcToSite -RestoreFrom %HOMEPATH%\Documents\Citrix\AutoConfig/  
Backup_yyyy_mm_dd_hh_mm_ss
```

### Note:

The backup file information in the preceding cmdlet is based on your own backups.

## Changing the default root folder

The `Export`, `Import`, `Merge`, `Sync`, and `Compare` operations can change the default root folder by using the `-AlternateFolder` parameter. The creation and management of per-operation subfolders remains the same as previously described.

## Files copied to subfolders

All files having an “.yml” extension are copied to operation subfolders except for the following:

- CustomerInfo.yml



- ZoneMapping.yml
- HostSecurity.yml

### **Automated fail-safe cloud site backups**

A backup of the current cloud site configuration is made before running operations that change the configuration. This includes **Import**, **Merge**, **Sync**, and **Restore** parameters. The backup is always in a subfolder beneath the operational subfolder.

In the case of **Restore**, the backup folder is a subfolder of the folder specified on the **-RestoreFrom** parameter.

### **Logging**

Running any cmdlet results in a created log file and an entry in the main history log file. All operation log files are placed in a backup folder. All log file names begin with **CitrixLog**, then show the auto-config operation and the date and timestamp of the cmdlet execution. Logs do not auto-delete.

The main history log is located in `%HOMEPATH%\Documents\Citrix\AutoConfig`, in the file named **History.Log**. Each cmdlet execution results in a main log entry containing the date, operation, result, backup, and log file locations of the execution.

### **Logging details**

Each log file includes the following:

- The name of the operation and whether check mode is enabled
- The start and end date and time
- Multiple entries for each component's actions and success/failure notifications
- Summary of actions taken including various counts of created objects
- Suggested fixes where applicable
- Backup folder location where applicable
- Main log location
- Duration

### **Diagnostic files**

Diagnostic files assist you in determining and resolving problems. The following files are created when their operation is run. They are located in the action-specific subfolder under `%HOMEPATH%\Documents\Citrix\AutoConfig`. Include these files when providing information for problem resolution support.

## **Export**

`PoshSdk_YYYY_MM_DD_HH_MM_SS.ps1`

This file enumerates all Broker PowerShell SDK calls made to export the site configuration to files.

## **Import, Merge, Restore, Sync, Backup, Compare**

`Transaction_YYYY_MM_DD_HH_MM_SS.txt`

This file documents each Rest API call and related information.

`RestApiContent_YYYY_MM_DD_HH_MM_SS.txt`

This file contains the all `Add`, `Update`, and `Delete` Rest API content.

## **Problems resulting from dependencies**

Imports and merges might fail due to missing dependencies. Some common problems are:

1. Group Policies are missing delivery group filters. The usual causes are delivery groups that have not been imported.
2. Applications fail to import or merge. The usual cause is missing delivery groups or application groups that have not been imported.
3. Application groups are missing a `RestrictToTag`. The usual causes are tags that have not been imported.
4. Host connections fail. The usual cause is missing security information in the `CvadAcSecurity.yml` file.
5. Machine catalogs fail. The usual cause is host connections that were not imported.
6. Machines missing from machine catalogs and delivery groups. The usual cause is machines that were not found in Active Directory.
7. Users missing from delivery groups. The usual cause is users that were not found in Active Directory.

## **Common problems**

### **The cloud site is empty**

An empty site indicates that a successful import/merge/restore operation has not completed. If an operation was attempted, review the logs from the operation.

### **The cloud site has more items than expected**

This can happen if enhancements were made to the cloud site configuration and a [Merge](#) operation was run. It can also happen if the cloud site configuration had previous values and an import/merge/restore operation had issues. Try rerunning the operation.

### **The cloud site has fewer items than expected**

This can happen if enhancements were made to the cloud site configuration and an [Import](#) operation with the [Merge](#) parameter set to `$false`, resulting in an identical derived state. The original state is backed up in the **Automated Fail-Safe Cloud Site Backup** folder associated with the import operation. To correct:

1. Restore the **Automated Fail-Safe Cloud Site Backup** using the Restore operation setting the `RestoreFolder` parameter to the **Automated Fail-Safe Cloud Site Backup** folder.
2. Use the [Merge](#) operation instead of the [Import](#) operation.

### **Cloud site edits after import have been lost**

This can occur if an [Import](#) operation was done with the `-Merge` parameter set to `$false`, resulting in an identical derived state import. The original state is backed up in the **Automated Fail-Safe Cloud Site Backup** folder associated with the import operation. To correct:

1. Restore the **Automated Fail-Safe Cloud Site Backup** using the Restore operation setting the `RestoreFolder` parameter to the **Automated Fail-Safe Cloud Site Backup** folder.
2. Use the [Merge](#) operation instead of the [Import](#) operation.

### **The cloud site does not match the current on-premises site**

This can occur when a [Merge](#) operation was done resulting in a merged derived state instead of an [Import](#) with the [Merge](#) parameter set to `$false` resulting in an identical derived state. Repeat using [Import](#) with the [Merge](#) parameter set to `$false`.

### **Recommendations**

- Do not run more than one instance of Automated Configuration at a time. Running multiple concurrent instances produces unpredictable results in the cloud site. If this occurs, rerun one instance of Automated Configuration to bring the site to the expected state.
- Do not work in Full Configuration while running Automated Configuration. Making changes in Full Configuration while running Automated Configuration can produce unpredictable results in the cloud site. If this occurs, rerun one instance of Automated Configuration to bring the site to the expected state.

- Always visually verify the merge/import/restore results in Full Configuration to ensure the cloud site meets expectations.

## Resources

- For more detailed and up-to-date troubleshooting and support information, see Knowledge Center article [CTX277730](#).
- For more information on Automated Configuration, see [Proof of Concept: Automated Configuration Tool](#).
- Visit the [Citrix Discussion forum for Automated Configuration](#).
- Watch [Under the Hood of the Automated Configuration Tool for Citrix Virtual Apps and Desktops](#) on YouTube.
- The Cloud Learning Center contains step-by-step video guides to building a service deployment, including the tasks described in this article. See [Migrating Citrix Virtual Apps and Desktops to Citrix Cloud Learning Path](#).

## Print

October 20, 2020

Managing printers in your environment is a multistage process:

1. Become familiar with printing concepts, if you are not already.
2. Plan your printing architecture. This includes analyzing your business needs, your existing printing infrastructure, how your users and applications interact with printing today, and which printing management model best applies to your environment.
3. Configure your printing environment by selecting a printer provisioning method and then creating policies to deploy your printing design. Update policies when new employees or servers are added.
4. Test a pilot printing configuration before deploying it to users.
5. Maintain your Citrix printing environment by managing printer drivers and optimizing printing performance.
6. Troubleshoot issues that may arise.

For complete information about printing in a Citrix Virtual Apps and Desktops environment, begin with [Print](#). From that article, you can move on to:

- [Printing configuration examples](#)
- [Best practices](#)

- [Printing policies and preferences](#)
- [Provision printers](#)
- [Maintain the printing environment](#)

### **Install the Universal Print Server on your print servers**

1. Ensure that each print server has Microsoft Visual C++ Runtime 2017, 32-bit and 64-bit installed.
2. Navigate to the Citrix Universal Print Server [download page](#) and click **Download File**.
3. Run one of the following commands on each print server:
  - For a 32-bit operating system: **UpsServer\_x86.msi**.
  - For a 64-bit operating system: **UpsServer\_x64.msi**.

After you install the Universal Print Server, configure it using the guidance in [Provision printers](#).

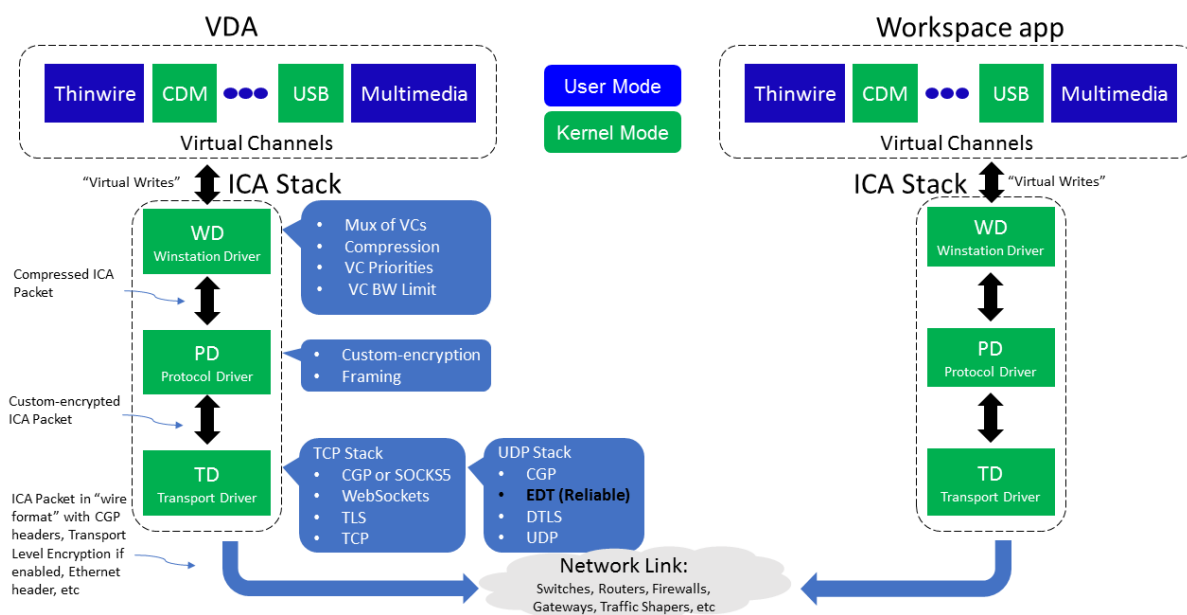
## **HDX**

May 7, 2021

### **Warning:**

Editing the registry incorrectly can cause serious problems that might require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

Citrix HDX represents a broad set of technologies that deliver a high-definition experience to users of centralized applications and desktops, on any device and over any network.




HDX is designed around three technical principles:

- Intelligent redirection
- Adaptive compression
- Data de-duplication

Applied in different combinations, they optimize the IT and user experience, decrease bandwidth consumption, and increase user density per hosting server.

- **Intelligent redirection** - Intelligent redirection examines screen activity, application commands, endpoint device, and network and server capabilities to instantly determine how and where to render an application or desktop activity. Rendering can occur on either the endpoint device or hosting server.
- **Adaptive compression** - Adaptive compression allows rich multimedia displays to be delivered on thin network connections. HDX first evaluates several variables, such as the type of input, device, and display (text, video, voice, and multimedia). It chooses the optimal compression codec and the best proportion of CPU and GPU usage. It then intelligently adapts based on each unique user and basis. This intelligent adaptation is per user, or even per session.



- **Data de-duplication** - De-duplication of network traffic reduces the aggregate data sent between client and server. It does so by taking advantage of repeated patterns in commonly accessed data such as bitmap graphics, documents, print jobs, and streamed media. Caching these patterns allows only the changes to be transmitted across the network, eliminating duplicate traffic. HDX also supports multicasting of multimedia streams, where a single transmission from the source is viewed by multiple subscribers at one location, rather than a one-to-one connection for each user.

For more information, see [Boost productivity with a high-definition user workspace](#).

### At the device

HDX uses the computing capacity of user devices to enhance and optimize the user experience. HDX technology ensures that users receive a smooth, seamless experience with multimedia content in their virtual desktops or applications. Workspace control enables users to pause virtual desktops and applications and resume working from a different device at the point where they left off.

### On the network

HDX incorporates advanced optimization and acceleration capabilities to deliver the best performance over any network, including low-bandwidth and high-latency WAN connections.

HDX features adapt to changes in the environment. The features balance performance and bandwidth. They apply the best technologies for each user scenario, whether the desktop or application is accessed locally on the corporate network or remotely from outside the corporate firewall.

### **In the data center**

HDX uses the processing power and scalability of servers to deliver advanced graphical performance, regardless of the client device capabilities.

HDX channel monitoring provided by Citrix Director displays the status of connected HDX channels on user devices.

### **HDX Insight**

HDX Insight is the integration of NetScaler Network Inspector and Performance Manager with Director. It captures data about ICA traffic and provides a dashboard view of real time and historical details. This data includes client-side and server-side ICA session latency, bandwidth use of ICA channels, and the ICA round-trip time value of each session.

You can enable NetScaler to use the HDX Insight virtual channel to move all the required data points in an uncompressed format. If you disable this feature, the NetScaler device decrypts and decompresses the ICA traffic spread across various virtual channels. Using the single virtual channel lessens complexity, enhances scalability, and is more cost effective.


#### **Minimum requirements:**

- Citrix Virtual Apps and Desktops 7 v1808
- XenApp and XenDesktop 7.17
- NetScaler version 12.0 Build 57.x
- Citrix Workspace app for Windows 1808
- Citrix Receiver for Windows 4.10
- Citrix Workspace app for Mac 1808
- Citrix Receiver for Mac 12.8

#### **Enable or disable HDX Insight virtual channel**

To disable this feature, set the Citrix NetScaler Application Flow service properties to Disabled. To enable, set the service to Automatic. In either case, we recommend that you restart the server machine after changing these properties. By default, this service is enabled (Automatic).





### Experience HDX capabilities from your virtual desktop

- To see how browser content redirection, one of four HDX multimedia redirection technologies, accelerates delivery of HTML5 and WebRTC multimedia content:
  1. Download the [Chrome browser extension](#) and install it on the virtual desktop.
  2. To experience how browser content redirection accelerates the delivery of multimedia content to virtual desktops, view a video on your desktop from a website containing HTML5 videos, such as YouTube. Users don't know when browser content redirection is running. To see whether browser content redirection is being used, drag the browser window quickly. You'll see a delay or out of frame between the viewport and the user interface. You can also right-click on the webpage and look for **About HDX Browser Redirection** in the menu.
- To see how HDX delivers high definition audio:
  1. Configure your Citrix client for maximum audio quality; see the Citrix Workspace app documentation for details.
  2. Play music files by using a digital audio player (such as iTunes) on your desktop.

HDX provides a superior graphics and video experience for most users by default, and configuration isn't required. Citrix policy settings that provide the best experience for most use cases are enabled by default.

- HDX automatically selects the best delivery method based on the client, platform, application, and network bandwidth, and then self-tunes based on changing conditions.
- HDX optimizes the performance of 2D and 3D graphics and video.
- HDX enables user devices to stream multimedia files directly from the source provider on the internet or intranet, rather than through the host server. If the requirements for this client-side content fetching are not met, media delivery falls back to server-side content fetching and multimedia redirection. Usually, adjustments to the multimedia redirection feature policies aren't needed.

- HDX delivers rich server-rendered video content to virtual desktops when multimedia redirection is not available: View a video on a website containing high definition videos, such as <http://www.microsoft.com/silverlight/iis-smooth-streaming/demo/>.

Good to know:

- For support and requirements information for HDX features, see the [System requirements](#) article. Except where otherwise noted, HDX features are available for supported Windows Multi-session OS and Windows Single-session OS machines, plus Remote PC Access desktops.
- This content describes how to optimize the user experience, improve server scalability, or reduce bandwidth requirements. For information about using Citrix policies and policy settings, see the [Citrix policies](#) documentation for this release.
- For instructions that include editing the registry, use caution: editing the registry incorrectly can cause serious problems that might require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

## Auto client reconnect and session reliability

When accessing hosted applications or desktops, network interruption might occur. To experience a smoother reconnection, we offer auto client reconnect and session reliability. In a default configuration, session reliability starts and then auto client reconnect follows.

### Auto client reconnect:

Auto client reconnect relaunches the client engine to reconnect to a disconnected session. Auto client reconnect closes (or disconnects) the user session after the time specified in the setting. If auto client reconnect is in progress, the system sends application and desktops network interruption notification to the user as follows:

- **Desktops.** The session window is grayed out and a countdown timer shows the time until the reconnections occur.
- **Applications.** The session window closes and a dialog appears to the user containing a countdown timer showing the time until the reconnections are attempted.

During auto client reconnect, sessions relaunch expecting network connectivity. User cannot interact with sessions while auto client reconnect is in progress.

On reconnection, the disconnected sessions reconnect using saved connection information. The user can interact with the applications and desktops normally.

Default auto client reconnect settings:

- Auto client reconnect timeout: 120 seconds
- Auto client reconnect: Enabled

- Auto client reconnect authentication: Disabled
- Auto client reconnect Logging: Disabled

For more information, see [Auto client reconnect policy settings](#).

### **Session reliability:**

Session reliability reconnects ICA sessions seamlessly across network interruptions. Session reliability closes (or disconnects) the user session after the time specified in the setting. After the session reliability timeout, the auto client reconnect settings take effect, attempting to reconnect the user to the disconnected session. When session reliability is in progress, application and desktops network interruption notification are sent to the user as follows:

- **Desktops.** The session window becomes translucent and a countdown timer shows the time until the reconnections occur.
- **Applications.** The window becomes translucent along with connection interrupted pop ups from the notification area.

While session reliability is active, the user cannot interact with the ICA sessions. However, user actions like keystrokes are buffered for few seconds immediately after the network interruption and retransmitted when the network is available.

On reconnection, the client and the server resume at the same point where they were in their exchange of protocol. The session windows lose translucency and appropriate notification area pop ups are shown for applications.

#### Default session reliability settings

- Session reliability timeout: 180 seconds
- Reconnection UI opacity level: 80%
- Session reliability connection: Enabled
- Session reliability port number: 2598

For more information, see [Session reliability policy settings](#).

### **NetScaler with auto client reconnect and session reliability:**

If Multistream and Multiport policies are enabled on the server and any or all these conditions are true, auto client reconnect does not work:

- Session reliability is disabled on NetScaler Gateway.
- A failover occurs on the NetScaler appliance.
- NetScaler SD-WAN is used with NetScaler Gateway.

### **HDX adaptive throughput**

HDX adaptive throughput intelligently fine-tunes the peak throughput of the ICA session by adjusting output buffers. The number of output buffers is initially set at a high value. This high value allows

data to be transmitted to the client more quickly and efficiently, especially in high latency networks. Providing better interactivity, faster file transfers, smoother video playback, higher framerate and resolution results in an enhanced user experience.

Session interactivity is constantly measured to determine whether any data streams within the ICA session are adversely affecting interactivity. If that occurs, the throughput is decreased to reduce the impact of the large data stream on the session and allow interactivity to recover.

**Important:**

HDX adaptive throughput changes the way that output buffers are set by moving this mechanism from the client to the VDA, and no manual configuration is necessary.

This feature has the following requirements:

- VDA version 1811 or later
- Workspace app for Windows 1811 or later

## **Improve the image quality sent to user devices**

The following visual display policy settings control the quality of images sent from virtual desktops to user devices.

- Visual quality. Controls the visual quality of images displayed on the user device: medium, high, always lossless, build to lossless (default = medium). The actual video quality using the default setting of medium depends on available bandwidth.
- Target frame rate. Specifies the maximum number of frames per second that are sent from the virtual desktop to the user device (default = 30). For devices that have slower CPUs, specifying a lower value can improve the user experience. The maximum supported frame rate per second is 60.
- Display memory limit. Specifies the maximum video buffer size for the session in kilobytes (default = 65536 KB). For connections requiring more color depth and higher resolution, increase the limit. You can calculate the maximum memory required.

## **Improve video conference performance**

Several popular video conferencing applications are optimized for delivery from Citrix Virtual Apps and Desktops through multimedia redirection (see, for example, [HDX RealTime Optimization Pack](#)). For applications that are not optimized, HDX webcam video compression improves bandwidth efficiency and latency tolerance for webcams during video conferencing in a session. This technology streams webcam traffic over a dedicated multimedia virtual channel. This technology uses less bandwidth compared to the isochronous HDX Plug-n-Play USB redirection support, and works well over WAN connections.

Citrix Workspace app users can override the default behavior by choosing the Desktop Viewer Mic & Webcam setting **Don't use my microphone or webcam**. To prevent users from switching from HDX webcam video compression, disable USB device redirection by using the policy settings under ICA policy settings > USB Devices policy settings.

HDX webcam video compression requires that the following policy settings be enabled (all are enabled by default).

- Client audio redirection
- Client microphone redirection
- Multimedia conferencing
- Windows Media Redirection

If a webcam supports hardware encoding, HDX video compression uses the hardware encoding by default. Hardware encoding might consume more bandwidth than software encoding. To force software compression, add the following DWORD key value to the registry key: `HKCU\Software\Citrix\HdxRealTime: DeepCompress_ForceSWEncode=1`.

## Network traffic priorities

Priorities are assigned to network traffic across multiple connections for a session using Quality of Service supported routers. Four TCP streams and two User Datagram Protocol (UDP) streams are available to carry ICA traffic between the user device and the server:

- TCP streams - real time, interactive, background, and bulk
- UDP streams - voice and Framehawk display remoting

Each virtual channel is associated with a specific priority and transported in the corresponding connection. You can set the channels independently, based on the TCP port number used for the connection.

Multiple channel streaming connections are supported for Virtual Delivery Agents (VDAs) installed on Windows 10, Windows 8, and Windows 7 machines. Work with your network administrator to ensure the Common Gateway Protocol (CGP) ports configured in the Multi-Port Policy setting are assigned correctly on the network routers.

Quality of Service is supported only when multiple session reliability ports, or the CGP ports, are configured.

### **Warning:**

Use transport security when using this feature. Citrix recommends using Internet Protocol Security (IPsec) or Transport Layer Security (TLS). TLS connections are supported only when the connections traverse a NetScaler Gateway that supports multi-stream ICA. On an internal corporate network, multi-stream connections with TLS are not supported.

To set Quality of Service for multiple streaming connections, add the following Citrix policy settings to a policy (see [Multi-stream connections policy settings](#) for details):

- Multi-Port policy - This setting specifies ports for ICA traffic across multiple connections, and establishes network priorities.
  - Select a priority from the CGP default port priority list. By default, the primary port (2598) has a High priority.
  - Type more CGP ports in CGP port1, CGP port2, and CGP port3 as needed, and identify priorities for each. Each port must have a unique priority.

Explicitly configure the firewalls on VDAs to allow the additional TCP traffic.

- Multi-Stream computer setting - This setting is disabled by default. If you use Citrix NetScaler SD-WAN with Multi-Stream support in your environment, you do not need to configure this setting. Configure this policy setting when using third-party routers or legacy Branch Repeaters to achieve the desired Quality of Service.
- Multi-Stream user setting - This setting is disabled by default.

For policies containing these settings to take effect, users must log off and then log on to the network.

## Show or hide the remote language bar

The language bar displays the preferred input language in an application session. If this feature is enabled (default), you can show or hide the language bar from the **Advanced Preferences > Language bar** UI in Citrix Workspace app for Windows. By using a registry setting on the VDA side, you can disable client control of the language bar feature. If this feature is disabled, the client UI setting doesn't take effect, and the per user current setting determines the language bar state. For more information, see [Improve the user experience](#).

To disable client control of the language bar feature from the VDA:

1. In the registry editor, navigate to `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Citrix\wfshell\TWI`.
2. Create a DWORD value key, `SeamlessFlags`, and set it to `0x40000`.

## Unicode keyboard mapping

Non-Windows Citrix Receivers use the local keyboard layout (Unicode). If a user changes the local keyboard layout and the server keyboard layout (scan code), they might not be in sync and the output is incorrect. For example, User1 changes the local keyboard layout from English to German. User1 then changes the server-side keyboard to German. Even though both keyboard layouts are German, they might not be in sync causing incorrect character output.

### **Enable or disable Unicode keyboard layout mapping**

By default, the feature is disabled on the VDA side. To enable the feature, toggle on the feature by using registry editor regedit on the VDA. Add the following registry key:

KEY\_LOCAL\_MACHINE/SOFTWARE/Citrix/CtxKlMap

Name: EnableKlMap

Type: DWORD

Value: 1

To disable this feature, set **EnableKlMap** to 0 or delete the **CtxKlMap** key.

### **Enable Unicode keyboard layout mapping compatible mode**

By default, Unicode keyboard layout mapping automatically hooks some windows API to reload the new Unicode keyboard layout map when you change the keyboard layout on the server side. A few applications cannot be hooked. To keep compatibility, you can change the feature to compatible mode to support these non-hooked applications. Add the following registry key:

HKEY\_LOCAL\_MACHINE/SOFTWARE/Citrix/CtxKlMap

Name: DisableWindowHook

Type: DWORD

Value: 1


To use normal Unicode keyboard layout mapping, set **DisableWindowHook** to 0.

## **Adaptive transport**

August 16, 2021

Adaptive Transport is a mechanism in Citrix Virtual Apps and Desktops that provides the ability to use Enlightened Data Transport (EDT) as the transport protocol for ICA connections. Adaptive Transport switches to TCP when EDT is not available.


EDT is a Citrix-proprietary transport protocol built on top of User Datagram Protocol (UDP). It delivers a superior user experience on challenging long-haul connections while maintaining server scalability. EDT improves data throughput for all ICA virtual channels on unreliable networks, providing a better and more consistent user experience.



When Adaptive Transport is set to **Preferred**, EDT is used as the primary transport protocol and TCP is used for fallback. By default, Adaptive Transport is set to **Preferred**. You can set Adaptive Transport to **Diagnostic mode** for testing purposes, which only allows EDT and disables the fallback to TCP.

With Citrix Workspace app for Windows, Mac, and iOS, EDT and TCP connections are attempted in parallel during the initial connection, session reliability reconnection, and automatic client reconnection. Doing so reduces the connection time if the underlying UDP transport is unavailable and TCP must be used instead. If Adaptive Transport is set to **Preferred** and the connection is established using TCP, Adaptive Transport continues to attempt to switch to EDT every five minutes.

With Citrix Workspace app for Linux and Android, EDT connections are attempted first. If the connection is unsuccessful, Citrix Workspace app tries to connect using TCP after the EDT request times out.





## System requirements

The following are the requirements for using Adaptive Transport and EDT:

- Control plane
  - Citrix Virtual Apps and Desktops Service
  - Citrix Virtual Apps and Desktops 1912 or later
- Virtual Delivery Agent
  - Version 1912 or later (2103 or later recommended)
  - Version 2012 is the minimum required for using EDT with Citrix Gateway Service
- StoreFront
  - Version 3.12.x
  - Version 1912.0.x
- Citrix Workspace app
  - Windows: version 1912 or later (2105 or later recommended)
  - Linux: version 1912 or later (2104 or later recommended)
  - Mac: version 1912 or later
  - iOS: latest version available in Apple App Store
  - Android: latest version available in Google Play
- Citrix Gateway (ADC)
  - 13.0.52.24 or later
  - 12.1.56.22 or later
- Firewall (from VDA perspective)
  - UDP 1494 inbound – if session reliability is disabled
  - UDP 2598 inbound – if session reliability is enabled
  - UDP 443 inbound – if VDA SSL is enabled for ICA encryption (DTLS)
  - UDP 443 outbound – if using Citrix Gateway Service. For more information, see the [Citrix Gateway service](#) documentation.

## Considerations

- Enable session reliability to use EDT MTU Discovery and to use EDT with Citrix Gateway and Citrix Gateway service.
- Ensure that the EDT MTU is adequately set to avoid fragmentation. Otherwise, performance can be impacted or sessions might fail to launch in some situations. For more information, see the [EDT MTU Discover](#) section.
- For details on requirements and considerations for using EDT with Citrix Gateway service, see [HDX Adaptive Transport with EDT support for Citrix Gateway service](#).
- For details on Citrix Gateway configuration to support EDT, see [Configure Citrix Gateway to support Enlightened Data Transport and HDX Insight](#).
- IPv6 is not supported currently.

## Configuration

Adaptive Transport is enabled by default. You can configure the following options using the **HDX Adaptive Transport** setting in Citrix policy.

- **Preferred.** This is the default setting. Adaptive Transport is enabled, and it uses EDT as the preferred transport protocol, with fallback to TCP.
- **Diagnostic mode.** Adaptive Transport is enabled, and it forces the use of EDT. Fallback to TCP is disabled. This setting is recommended for testing and troubleshooting only.
- **Off.** Adaptive Transport is disabled, and only TCP is used for transport.

To confirm that EDT is being used as the transport protocol for the session, you can use Director or the CtxSession.exe command-line utility on the VDA.

In Director, look up the session and select **Details**. If the **Connection type** is **HDX** and the **Protocol** is **UDP**, EDT is being used as the transport protocol for the session. If the **Connection type** is **RDP**, ICA is not in use, and the **Protocol** displays N/A. For more information, see [Monitor sessions](#).

## Session Details

Session Control ▾   Shadow   Send Message

<b>ID</b>	2
<b>Session State</b>	Active
<b>Application State</b>	Desktop
<b>Anonymous</b>	No
<b>Time in state</b>	0 minutes
<b>Endpoint name</b>	
<b>Endpoint IP</b>	
<b>Connection type</b>	HDX
<b>Protocol</b>	UDP
<b>Citrix Workspace App Version</b>	21.5.0.48
<b>ICA RTT</b>	67 ms
<b>ICA Latency</b>	65 ms
<b>Launched via</b>	n/a
<b>Connected via</b>	

To use the CtxSession.exe utility, launch a Command Prompt or PowerShell within the session and run `ctxsession.exe`. To see verbose statistics, run `ctxsession.exe -v`. If EDT is in use, the transport protocol shows one of the following:

- **UDP > ICA** (Session Reliability disabled)
- **UDP > CGP > ICA** (Session Reliability enabled)
- **UDP > DTLS > CGP > ICA** (ICA is DTLS-encrypted end-to-end)

```
Administrator: Windows PowerShell
PS C:\windows\system32> ctxsession -v

Session Id 2:
Transport Protocols: UDP -> DTLS -> CGP -> ICA
  Local Address:
  Remote Address:
  Client Address:
Security Protocol: UNKNOWN VALUE - 131072
Security Cipher: 128 bit AES
Cipher Strength: 128 bits
ICA Encryption: Basic

EDT Reliable Statistics:
Bandwidth 121.777 Mbps, Send Rate 0 bps, Recv Rate 0 bps, RTT 65.531 ms
Sent 0, Sent Lost 0 (0.00%), Rcvd 0, Rcvd Lost 0 (0.00%)
Sent ACKs 0, Sent NAKs 0, Rcvd ACKs 0, Rcvd NAKs 0
Flow Window 16383, Congest Window 4050, Delivery Rate 7591
EDT MTU: 1400

ICA Statistics:
SentBandwidth (bps) = 6376 RecvBandwidth (bps) = 568
SentPreCompression = 1800688 RecvPreExpansion = 32864
SentPostCompression = 1429125 RecvPostExpansion = 137041
Compression Ratio % = 79 Expansion Ratio % = 23
LastLatency = 67 AverageLatency = 53
IcaBufferLength = 980
```

## EDT MTU Discovery

MTU Discovery allows EDT to automatically determine the Maximum Transmission Unit (MTU) when establishing a session. Doing so prevents EDT packet fragmentation that might result in performance degradation or failure to establish a session.

## Requirements

- VDA minimum version 1912 (2103 or later recommended)
- Citrix Workspace app
  - Windows: version 1912 or later (2105 or later recommended)
  - Mac: version 2108 or later
  - Android: version 21.5 or later
- Citrix ADC:
  - 13.0.52.24
  - 12.1.56.22
- Session Reliability must be enabled

If you use client platforms or versions that don't support this feature, see [CTX231821](#) for details about configuring a custom EDT MTU that is appropriate for your environment.

### Important:

MTU Discovery is not supported with Multi-Stream ICA.

### To control EDT MTU Discovery on the VDA

MTU Discovery is enabled by default. To disable this feature, delete the **EDT MTU Discovery** registry value and restart the VDA. For more information, see the [EDT MTU Discovery](#) setting in the list of HDX features managed through the registry.

#### Warning:

Editing the registry incorrectly can cause serious problems that might require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

### To control EDT MTU Discovery on the client

You can control EDT MTU Discovery selectively on the client by adding the **MtuDiscovery** parameter in the ICA file. To disable the feature, set the following under the **Application** section:

```
MtuDiscovery=Off
```

To re-enable the feature, remove the **MtuDiscovery** parameter from the ICA file.

#### Important:

For this ICA file parameter to work, enable the feature on the VDA. If the feature is not enabled on the VDA, the ICA file parameter has no effect.

### Loss tolerant mode

#### Important:

- The feature requires a minimum of Citrix Workspace app 2002 for Windows.
- Loss tolerant mode is not supported on Citrix Gateway or Citrix Gateway Service. This mode is available only with direct connections.

Loss tolerant mode uses the EDT Lossy transport protocol to enhance the user experience for users connecting through networks with high latency and packet loss.

Initially, sessions are established using EDT. If the latency and packet loss thresholds are reached or surpassed, the applicable virtual channels switch from EDT to EDT Lossy, while leaving the other virtual channels on EDT. If the latency and packet loss decrease below the thresholds, the applicable virtual channels switch back to EDT.

The default thresholds are:

- Packet loss: 5%
- Latency: 300 ms (RTT)

Loss tolerant mode is enabled by default. You can disable the mode or adjust the packet loss and latency thresholds using the loss tolerant mode thresholds setting.

## Requirements

- Citrix Virtual Delivery Agent (VDA) 2003
- Citrix Workspace app 2002 for Windows
- Session reliability enabled. For more information about session reliability, see [Session reliability policy settings](#).

## Known issues

Adaptive Transport and EDT contain the following issues:

- Packet fragmentation can cause performance degradation or even failure to launch sessions. You can adjust the EDT MTU to avoid this. Use MTU Discovery or the workaround described in [CTX231821](#).
- A grey or black screen may appear when launching a session from a Windows client if MTU Discovery is enabled. To address this issue, upgrade to Workspace app for Windows 2105 or later or Workspace app for Windows 1912 CU4 or later.
- Fallback to TCP may fail on Linux and Android clients when connecting through Citrix Gateway or Citrix Gateway Service. This happens when there is a successful EDT negotiation between the client and the Gateway, and the EDT negotiation fails between the Gateway and the VDA. To address this issue, upgrade to Workspace app for Linux 2104 or later and Workspace app for Android 21.5 or later.
- Asymmetrical network paths can cause MTU Discovery to fail for connections that do not go through Citrix Gateway or Citrix Gateway Service. To address this issue, upgrade to VDA version 2103 or later. [CVADHELP-16654]
- When using Citrix Gateway or Citrix Gateway Service, asymmetrical network paths can cause MTU Discovery to fail. This is due to an issue on Gateway that causes the Don't Fragment (DF) bit in the EDT packets' header not to be propagated. A fix for this issue is not yet available. [CGOP-18438]
- MTU Discovery may fail for users that connect through a DS-Lite network. Some modems fail to honor the DF bit when packet processing is enabled, preventing MTU Discovery from detecting fragmentation. In this situation, these are the available options:
  - Disable packet processing on the user's modem.
  - Disable MTU Discovery and use a hardcoded MTU as described in [CTX231821](#).
  - Disable Adaptive Transport to force sessions to use TCP. If only a subset of users is affected, consider disabling it on the client-side so that other users can continue to use EDT.

## Troubleshoot

To troubleshoot Adaptive Transport and EDT, we suggest the following:

1. Thoroughly review and validate the [requirements](#), [considerations](#), and [known issues](#).
2. Check if there are Citrix policies in Studio or GPO overwriting the desired **HDX Adaptive Transport** setting.
3. Check if there are settings on the client overwriting the desired HDX Adaptive Transport setting. This can be a GPO preference, a setting configured using the optional Workspace app administrative template, or a manual configuration of the **HDXoverUDP** setting in the registry or client's configuration file.
4. On multi-session VDA machines, ensure that the UDP listeners are active. Open a command prompt in the VDA machine and run `netstat -a -p udp`. For more information, see [How to Confirm HDX Enlightened Data Transport Protocol](#).
5. Launch a direct session internally, bypassing the Citrix Gateway, and check the protocol in use. If the session uses EDT, the VDA is ready to use EDT for external connections through Citrix Gateway.
6. If EDT works for direct internal connections and not for sessions going through Citrix Gateway:
  - Ensure that Session Reliability is enabled
  - Ensure that the Gateway has DTLS enabled
7. Check if the appropriate firewall rules have been configured in both network firewalls and firewalls running on the VDA machines.
8. Check if your users' connections require a non-standard MTU. Connections with an effective MTU lower than 1500 bytes cause EDT packet fragmentation, which in turn can affect performance or even cause session launch failures. This issue is common when using VPN, some Wi-Fi access points, and mobile networks, such as 4G and 5G. For information on how to address this issue, see the [MTU Discovery](#) section.

## Interoperability with Citrix SD-WAN

Citrix SD-WAN WAN optimization (WANOP) offers cross-session tokenized compression (data deduplication), including URL-based video caching, providing significant bandwidth reduction. The reduction occurs if two or more people at the office location watch the same client-fetched video or transfer or print significant portions of the same file or document. Furthermore, by running the processes for ICA data reduction and print job compression on the branch office appliance, WANOP offers VDA server CPU offload and enables higher Citrix Virtual Apps and Desktops server scalability.

Currently, SD-WAN WANOP does not support EDT. However, there is no need to disable Adaptive Transport if SD-WAN WANOP is in use. When a user launches a session that goes through an SD-WAN with

WANOP enabled, it automatically sets the session to use TCP as the transport protocol. Non-WANOP sessions continue to use EDT whenever possible.

## Rendezvous protocol

September 9, 2021

In environments that use the Citrix Gateway service, the Rendezvous protocol allows HDX sessions to bypass the Citrix Cloud Connector and connect directly and securely to the Citrix Gateway service.

### Requirements

- Access to environment using Citrix Workspace and Citrix Gateway service.
- Control Plane: Citrix Virtual Apps and Desktops Service (Citrix Cloud).
- VDA: Version 1912 or later.
  - Version 2012 is the minimum required for EDT Rendezvous.
  - Version 2012 is the minimum required for non-transparent proxy support (no PAC file support).
  - Version 2103 is the minimum required for proxy configuration with a PAC file.
- Enable the Rendezvous protocol in the Citrix policy. For more information, see [Rendezvous protocol policy setting](#).
- The VDAs must have access to [https://\\*.nssvc.net](https://*.nssvc.net), including all subdomains. If you can't add all subdomains to the allow list in that manner, use [https://\\*.c.nssvc.net](https://*.c.nssvc.net) and [https://\\*.g.nssvc.net](https://*.g.nssvc.net) instead. For more information, see the [Internet Connectivity Requirements](#) section of the Citrix Cloud documentation (under Virtual Apps and Desktops service) and the Knowledge Center article [CTX270584](#).
- The VDAs must be able to connect to the addresses mentioned above on TCP 443 and UDP 443 for TCP Rendezvous and EDT Rendezvous, respectively.
- Cloud Connectors must obtain the VDAs' FQDNs when brokering a session. Accomplish this in one of these two ways:
  - **Enable DNS resolution for the site.** Navigate to **Full Configuration > Settings** and turn on the **Enable DNS resolution** setting. Alternatively, use the Citrix Virtual Apps and Desktops Remote PowerShell SDK and run the command `Set-BrokerSite -DnsResolutionEnabled $true`. For more information about the Citrix Virtual Apps and Desktops Remote PowerShell SDK, see [SDKs and APIs](#).



- **DNS Reverse Lookup Zone with PTR records for the VDAs.** If you choose this option, we recommend that you configure VDAs to always attempt to register PTR records. To do so, use the Group Policy Editor or Group Policy Object, navigate to **Computer Configuration > Administrative Templates > Network > DNS Client**, and set **Register PTR Records** to **Enabled and Register**. If the connection's DNS suffix does not match the domain's DNS suffix, you must also configure the **Connection-specific DNS suffix** setting for the machines to register PTR records successfully.

**Note:**

If using the DNS resolution option, the Cloud Connectors must be able to resolve the fully qualified domain names (FQDNs) of the VDA machines. In the case that internal users connect directly to the VDA machines, the client devices also must be able to resolve the VDA machines' FQDNs.

If using a DNS reverse lookup zone, the FQDNs in the PTR records must match the FQDNs of the VDA machines. If the PTR record contains a different FQDN, the Rendezvous connection fails. For example, if the machine's FQDN is `vda01.domain.net`, the PTR record must contain `vda01.domain.net`. A different FQDN such as `vda01.sub.domain.net` does not work.

## Proxy configuration

The VDA supports establishing Rendezvous connections through a proxy.

### Proxy considerations

Consider the following when using proxies with Rendezvous:

- Transparent proxies, non-transparent HTTP proxies, and SOCKS5 proxies are supported.
- Packet decryption and inspection are not supported. Configure an exception so that the ICA traffic between the VDA and the Gateway Service is not intercepted, decrypted, or inspected. Otherwise, the connection breaks.
- HTTP proxies support machine-based authentication by using Negotiate and Kerberos or NT LAN Manager (NTLM) authentication protocols.

When you connect to the proxy server, the Negotiate authentication scheme automatically selects the Kerberos protocol. If Kerberos isn't supported, Negotiate falls back to NTLM for authentication.

**Note:**

To use Kerberos, you must create the service principal name (SPN) for the proxy server and

associate it with the proxy's Active Directory account. The VDA generates the SPN in the format `HTTP/<proxyURL>` when establishing a session, where the proxy URL is retrieved from the **Rendezvous proxy** policy setting. If you don't create an SPN, authentication falls back to NTLM. In both cases, the VDA machine's identity is used for authentication.

- Authentication with a SOCKS5 proxy is not currently supported. If using a SOCKS5 proxy, you must configure an exception so that traffic destined to Gateway Service addresses (specified in the requirements) can bypass authentication.
- Only SOCKS5 proxies support data transport through EDT. For an HTTP proxy, use TCP as transport protocol for ICA.

### Transparent proxy

If using a transparent proxy in your network, no additional configuration is required on the VDA.

### Non-transparent proxy

If using a non-transparent proxy in your network, configure the [Rendezvous proxy configuration](#) setting. When the setting is enabled, specify the HTTP or SOCKS5 proxy address, or enter the path to the PAC file for the VDA to know which proxy to use. For example:

- Proxy address: `http://<URL or IP>:<port>` or `socks5://<URL or IP>:<port>`
- PAC file: `http://<URL or IP>/<path>/<filename>.pac`

If you use the PAC file to configure the proxy, define the proxy using the syntax required by the Windows HTTP service: `PROXY [<scheme>=]<URL or IP>:<port>`. For example, `PROXY socks5=<URL or IP>:<port>`.

### Rendezvous validation

If you meet all requirements, follow these steps to validate if Rendezvous is in use:

1. Launch PowerShell or a command prompt within the HDX session.
2. Run `ctxsession.exe -v`.
3. The transport protocols in use indicate the type of connection:
  - TCP Rendezvous: **TCP > SSL > CGP > ICA**
  - EDT Rendezvous: **UDP > DTLS > CGP > ICA**
  - Proxy through Cloud Connector: **TCP > CGP > ICA**

## Additional considerations

### Windows cipher suite order

For a custom cipher suite order, make sure that you include the VDA-supported cipher suites from the following list:

- TLS\_ECDHE\_RSA\_WITH\_AES\_256\_GCM\_SHA384
- TLS\_ECDHE\_RSA\_WITH\_AES\_256\_CBC\_SHA384
- TLS\_ECDHE\_RSA\_WITH\_AES\_128\_CBC\_SHA

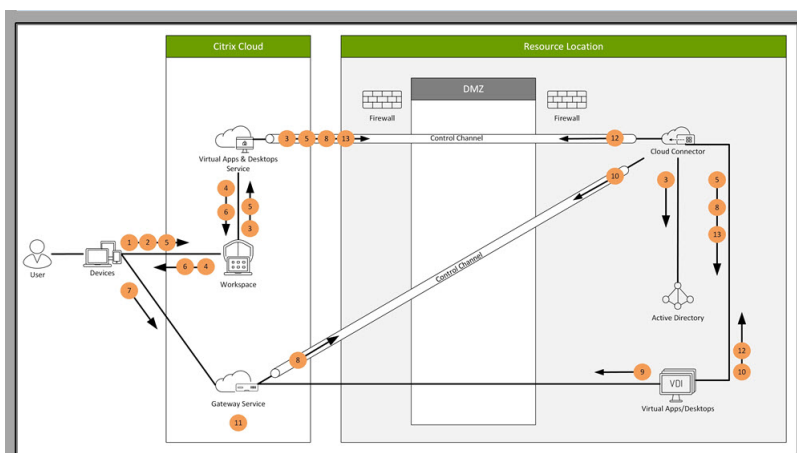
If the custom cipher suite order does not contain these cipher suites, the Rendezvous connection fails.

### Zscaler Private Access

If using Zscaler Private Access (ZPA), it is highly recommended that you configure bypass settings for the Gateway Service to avoid increased latency and the associated performance impact. To do so, you must define application segments for the Gateway Service addresses – specified in the requirements – and set them to always bypass. For information on configuring application segments to bypass ZPA, see the [Zscaler documentation](#).

## How Rendezvous works

This diagram is an overview of the Rendezvous connection flow.



Follow the steps to understand the Rendezvous connection flow:

1. Navigate to Citrix Workspace.
2. Enter credentials in Citrix Workspace.
3. If using on-premises Active Directory, the Citrix Virtual Apps and Desktops service authenticates credentials with Active Directory using the Cloud Connector channel.
4. Citrix Workspace displays counted resources from the Citrix Virtual Apps and Desktops service.

5. Select resources from Citrix Workspace. The Citrix Virtual Apps and Desktops service sends a message to the VDA to prepare for an incoming session.
6. Citrix Workspace sends an ICA file to the endpoint that contains an STA ticket generated by Citrix Cloud.
7. The endpoint connects to the Citrix Gateway service, provides the ticket to connect to the VDA, and Citrix Cloud validates the ticket.
8. The Citrix Gateway service sends connection information to the Cloud Connector. The Cloud Connector determines if the connection is supposed to be a Rendezvous connection and sends the information to the VDA.
9. The VDA establishes a direct connection to the Citrix Gateway service.
10. If a direct connection between the VDA and the Citrix Gateway service isn't possible, the VDA proxies its connection through the Cloud Connector.
11. The Citrix Gateway service establishes a connection between the endpoint device and the VDA.
12. The VDA verifies its license with the Citrix Virtual Apps and Desktops service through the Cloud Connector.
13. The Citrix Virtual Apps and Desktops service sends and applies the session policies to the VDA through the Cloud Connector.

## Citrix ICA virtual channels

April 9, 2021

### **Warning:**

Editing the registry incorrectly can cause serious problems that might require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

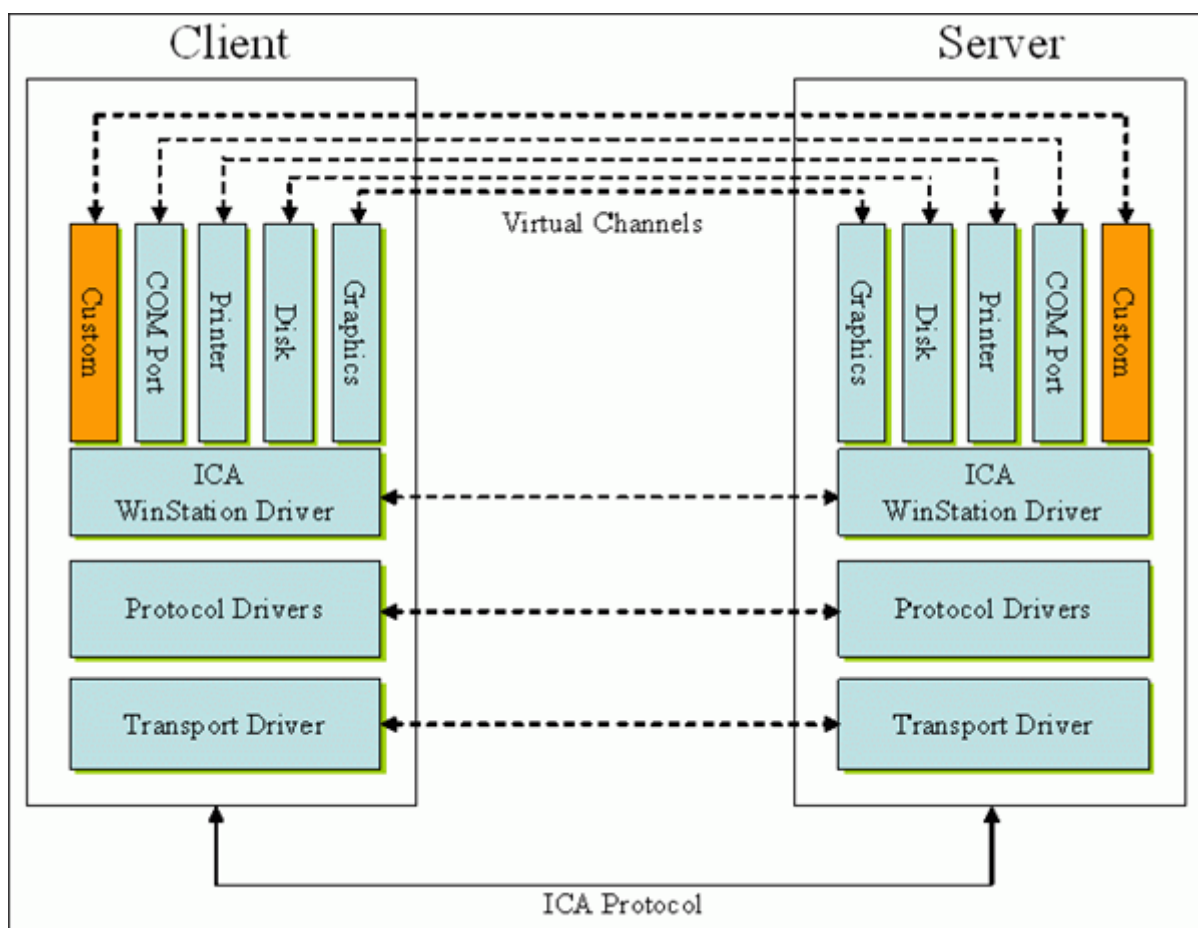
### **What are ICA virtual channels?**

A large portion of the functionality and communication between the Citrix Workspace app and the Citrix Virtual Apps and Desktops servers occurs over virtual channels. Virtual channels are a necessary part of the remote computing experience with the Citrix Virtual Apps and Desktops servers. Virtual channels are used for:

- Audio
- COM ports
- Disks
- Graphics

- LPT ports
- Printers
- Smart cards
- Third-party custom virtual channels
- Video

New virtual channels are sometimes released with new versions of the Citrix Virtual Apps and Desktops servers and Citrix Workspace app products to provide more functionality.



A virtual channel consists of a client-side virtual driver that communicates with a server-side application. Citrix Virtual Apps and Desktops ship with various virtual channels included. They're designed to allow customers and third-party vendors to create their own virtual channels by using one of the provided Software Development Kits (SDKs).

Virtual channels provide a secure way to accomplish various tasks. For example, an application that is running on a Citrix Virtual Apps server that is communicating with a client-side device or an application that is communicating with the client-side environment.

On the client side, virtual channels correspond to virtual drivers. Each virtual driver provides a specific function. Some are required for normal operation, and others are optional. Virtual drivers operate at

the presentation layer protocol level. There can be several protocols active at any time by multiplexing channels that are provided by the Windows Station (WinStation) protocol layer.

The following functions are contained in the VirtualDriver registry value under this registry path:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\ICA Client\Engine\Configuration\Advanced  
\Modules\ICA 3.0
```

or

```
HKEY_LOCAL_MACHINE\SOFTWARE\WOW6432Node\Citrix\ICA Client\Engine\Configuration  
\Advanced\Modules\ICA 3.0 (for 64-bit)
```

- Thinwire3.0 (Required)
- ClientDrive
- ClientPrinterQueue
- ClientPrinterPort
- Clipboard
- ClientComm
- ClientAudio
- LicenseHandler (Required)
- TWI (Required)
- SmartCard
- ICACTL (Required)
- SSPI
- TwainRdr
- UserEXperience
- Vd3d

**Note:**

You can disable specific client functionality by removing one or more of these values from the registry key. For example, if you wanted to remove the Client Clipboard, remove the word **Clipboard**.

This list contains the client virtual driver files and their respective functions. Citrix Virtual Apps and Citrix Workspace app for Windows use these files. They are in the form of Dynamic Link Libraries (user mode), and not Windows drivers (kernel mode) except for Generic USB as described in Generic USB virtual channel.

- vd3dn.dll – Direct3D virtual channel used for desktop composition redirection
- vdcamN.dll – Bidirectional audio
- vdcdm30n.dll – Client drive mapping
- vdcom30N.dll - Client COM port mapping
- vdcpm30N.dll – Client printer mapping
- vdctlN.dll – ICA controls channel

- vddvc0n.dll – Dynamic virtual channel
- vdeuemn.dll - End user experience monitoring
- vdgusbn.dll – Generic USB virtual channel
- vdkbhook.dll – Transparent key pass-through
- vdlfpn.dll – Framehawk display channel over UDP like transport
- vdmmn.dll – Multimedia support
- vdmrvc.dll – Mobile Receiver virtual channel
- vdmtn.dll - Multi-touch support
- vdscardn.dll – Smartcard support
- vdsens.dll – Sensors virtual channel
- vdspl30n.dll – Client UPD
- vdsspin.dll – Kerberos
- vdtuin.dll – Transparent UI
- vdtw30n.dll – Client Thinwire
- vdtwin.dll – Seamless
- vdtwn.dll – Twain

Some virtual channels are compiled into other files. For example Clipboard Mapping is available in wfica32.exe

### **64-bit compatibility**

Citrix Workspace app for Windows is 64-bit compatible. As with most of the binaries compiled for 32 bit, these client files have 64-bit compiled equivalents:

- brapi64.dll
- confmgr.dll
- ctxlogging.dll
- ctxmui.dll
- icaconf.exe
- icaconfs.dll
- icafile.dll
- pnipcn64.dll
- pnsson.dll
- ssoncom.exe
- ssonstub.dll
- vdkbhook64.dll

### **Generic USB virtual channel**

Generic USB virtual channel implementation uses two kernel mode drivers along with virtual channel driver vdgusbn.dll:

- ctxusbm.sys
- ctxusbr.sys

### **How ICA virtual channels work**

Virtual channels are loaded in multiple ways. The Shell (WfShell for the server and PicaShell for the workstation) load some virtual channels. Some virtual channels are hosted as windows services.

Virtual channel modules loaded by the Shell, for example:

- EUEM
- Twain
- Clipboard
- Multimedia
- Seamless session sharing
- Time Zone

Some are loaded as kernel mode, for example:

- CtxDvcs.sys – Dynamic virtual channel
- Icausbb.sys – Generic USB redirection
- Picadm.sys – Client drive mapping
- Picaser.sys – COM port redirection
- Picapar.sys – LPT port redirection

### **Graphics virtual channel on the server side**

Starting with XenApp 7.0 and XenDesktop7.0, `ctxgfx.exe` hosts the graphics virtual channel for both workstation and terminal server based sessions. `Ctxgfx` hosts platform specific modules that interact with the corresponding driver (`Icardd.dll` for RDSH and `vdod.dll` and `vidd.dll` for workstation).

For XenDesktop 3D Pro deployments an OEM graphics driver is installed for the corresponding GPU on the VDA. `Ctxgfx` loads specialized adaptor modules to interact with the OEM graphics driver.

### **Hosting specialized channels in windows services**

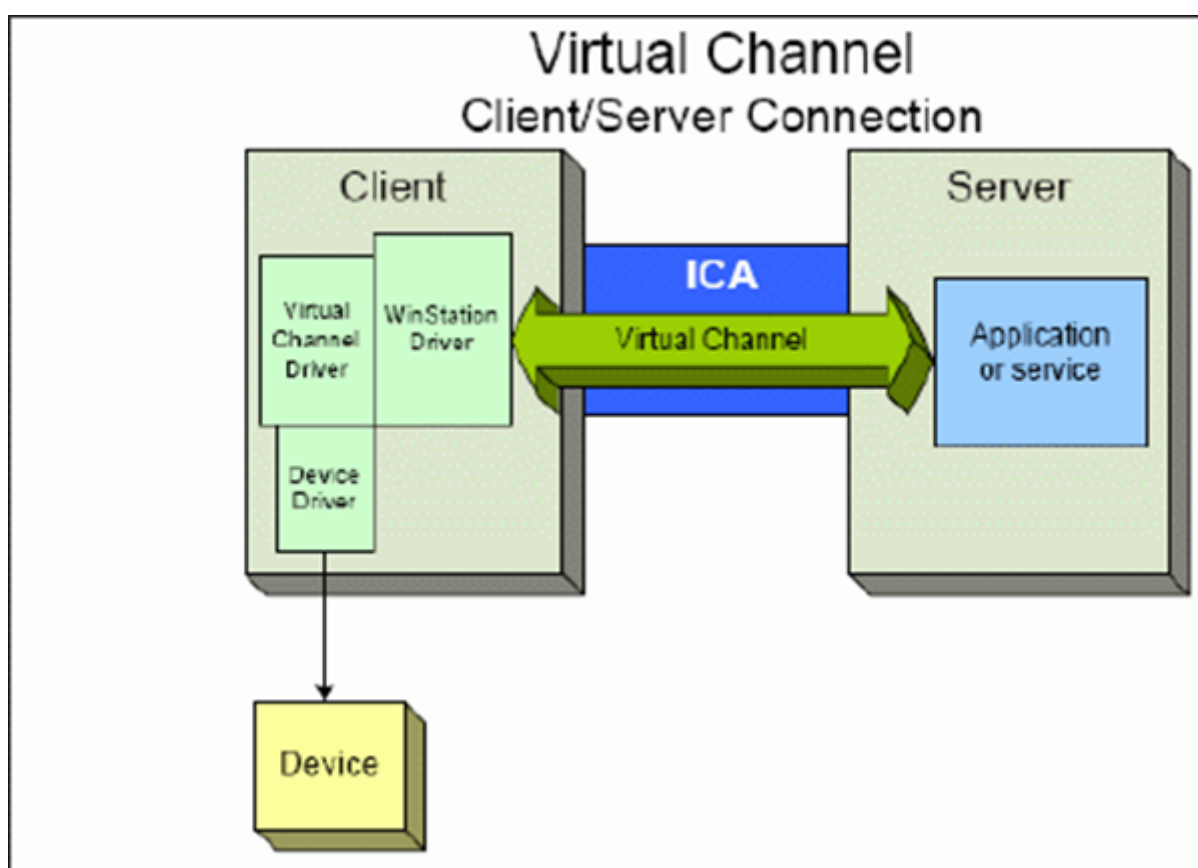
On Citrix Virtual Apps and Desktops servers, various channels are hosted as windows services. Such hosting provides one-to-many semantics for multiple applications in a session and multiple sessions on the server. Examples of such services include:



- Citrix Device Redirector Service
- Citrix Dynamic Virtual Channel Service
- Citrix End User Experience Monitoring Service
- Citrix Location and Sensor Virtual Channel Service
- Citrix MultiTouch Redirection Service
- Citrix Print Manager Service
- Citrix Smartcard Service
- Citrix Audio Redirection Service (Citrix Virtual Desktops only)

The audio virtual channel on Citrix Virtual Apps is hosted using Windows Audio service.

On the server side, all client virtual channels are routed through the WinStation driver, Wdica.sys. On the client side, the corresponding WinStation driver, built into wfica32.exe, polls the client virtual channels. This image illustrates the virtual channel client-server connection.



This overview contains a client-server data exchange using a virtual channel.

1. The client connects to the Citrix Virtual Apps and Desktops server. The client passes information about the virtual channels it supports to the server.
2. The server-side application starts, obtains a handle to the virtual channel, and optionally queries for additional information about the channel.

3. The client virtual driver and server-side application pass data using the following two methods:
  - If the server application has data to send to the client, the data is sent to the client immediately. When the client receives the data, the WinStation driver de-multiplexes the virtual channel data from the ICA stream and immediately passes it to the client virtual driver.
  - If the client virtual driver has data to send to the server, the data is sent the next time the WinStation driver polls it. When the server receives the data, it is queued until the virtual channel application reads it. There is no way to alert the server virtual channel application that data was received.
4. When the server virtual channel application is completed, it closes the virtual channel and frees any allocated resources.

### Creating your own virtual channel using the Virtual Channel SDK

Creating a virtual channel using the Virtual Channel SDK requires intermediate programming knowledge. Use this method to provide a major communication path between the client and the server. For example, if you are implementing usage of a device on the client side, such as a scanner, to be used with a process in the session.

#### Note:

- The Virtual Channel SDK requires the WFAPI SDK to write the server side of the virtual channel.
- Because of enhanced security for Citrix Virtual Apps and Desktops, you must specify which virtual channels are allowed to be opened in an ICA session. For more information, see [Virtual channel allow list policy settings](#).

### Creating your own virtual channel using the ICA Client Object SDK

Creating a virtual channel using the ICA Client Object (ICO) is easier than using the Virtual Channel SDK. Use the ICO by creating a named object in your program using the **CreateChannels** method.

#### Important:

Because of enhanced security starting with the 10.00 version of the Citrix Receiver for Windows and later (and Citrix Workspace apps for Windows), you must take an extra step when creating an ICO virtual channel.

For more information, see [Client Object API Specification Programmer's Guide](#).

## Pass-through functionality of virtual channels

Most virtual channels that Citrix provides operate unmodified when you use the Citrix Workspace app for Windows within an ICA session (also known as a pass-through session). There are considerations when using the client in extra hops.

The following functions operate the same way in single or multiple hops:

- Client COM port mapping
- Client drive mapping
- Client printer mapping
- Client UPD
- End user experience monitoring
- Generic USB
- Kerberos
- Multimedia support
- Smartcard support
- Transparent key pass-through
- Twain

As the inherent nature of latency and factors such as compression and decompression and rendering being performed at each hop, performance might be affected with each additional hop that the client undergoes. The affected areas are:

- Bidirectional audio
- File transfers
- Generic USB redirection
- Seamless
- Thinwire

### Important:

By default, the client drives mapped by an instance of the client running in a pass-through session are restricted to the client drives of the connecting client.

## Pass-through functionality of virtual channels between a Citrix Virtual Desktop session and a Citrix Virtual App session

Most virtual channels provided by Citrix operate unmodified when you use Citrix Workspace app for Windows within an ICA session on a Citrix Virtual Desktops server (also known as a pass-through session).

Specifically, on the Citrix Virtual Desktops server, there is a VDA hook that runs **picaPassthruHook**. This hook makes the client think it's running on a CPS server, and placing the client into its traditional pass-through mode.

We support the following traditional virtual channels and their functionality:

- Client
- Client COM port mapping
- Client drive mapping
- Client printer mapping
- Generic USB (limited due to performance)
- Multimedia support
- Smartcard support
- SSON
- Transparent key pass-through

### **Security and ICA virtual channels**

Securing usage is an important part of planning, developing, and implementing virtual channels. There are several references to specific areas of security located throughout this document.

### **Best practices**

Open virtual channels when you **Connect** and **Reconnect**. Close virtual channels when you log off and **Disconnect**.

Keep the following guidelines in mind when you create scripts that use virtual channel functions.

#### **Naming the Virtual Channels:**

You can create a maximum of 32 virtual channels. Seventeen of the 32 channels are reserved for special purposes.

- Virtual channel names must not be more than seven characters in length.
- The first three characters are reserved for the vendor name, and the next four for the channel type. For example, **CTXAUD** represents the Citrix audio virtual channel.

Virtual channels are referred to by a seven-character (or shorter) ASCII name. In some previous versions of the ICA protocol, virtual channels were numbered. The numbers are now assigned dynamically based on the ASCII name, making implementation easier. Users who are developing virtual channel code for internal use only can use any seven-character name that does not conflict with existing virtual channels. Use only numbers and upper and lowercase ASCII. Follow the existing naming convention when adding your own virtual channels. There are several predefined channels. The predefined channels begin with the OEM identifier CTX and are for use only by Citrix.

#### **Double-Hop Support:**

Virtual Channel	Is double hop supported?
Audio	No
Browser Content Redirection	No
CDM	Yes
CEIP	No
Clipboard	Yes
Continuum (MRVC)	No
Control VC	Yes
HTML5 Video Redirection (v1)	Yes
Keyboard, Mouse	Yes
MultiTouch	No
NSAPVC	No
Printing	Yes
SensVC	No
Smartcard	Yes
Twain	Yes
USB VC	Yes
WAYCOM devices -K2M using USB VC	Yes
Webcam Video Compression	Yes
Windows Media Redirection	Yes


### See also

- [ICA Virtual Channel SDK](#)
- The [Citrix Developer Network](#) is the home for all technical resources and discussions involving the use of Citrix SDKs. In this network, you can find access to SDKs, sample code and scripts, extensions and plug-ins, and SDK documentation. Also included are the Citrix Developer Network forums, where technical discussions take place around each of the Citrix SDKs.

## Double hop in Citrix Virtual Apps and Desktops

April 8, 2021

In the context of a Citrix client session, the term “double hop” refers to a Citrix Virtual App session that is running within a Citrix Virtual Desktop session. The following diagram illustrates a double hop.



In a double hop scenario, when the user connects to a Citrix Virtual Desktop running on a single-session OS VDA (known as VDI) or a multi-session OS VDA (known as a published desktop), that is considered the first hop. After the user connects to the virtual desktop, the user can launch a Citrix Virtual Apps session. That is considered the second hop.

You can use a double hop deployment model to support various use cases. The case where the Citrix Virtual Desktop and the Citrix Virtual Apps environments are managed by different entities is one common example. This method can also be effective in resolving application compatibility issues.

### System requirements

All Citrix Virtual Apps and Desktops editions including the Citrix Cloud service support double hop.

The first hop must use a supported version of the single-session or multi-session OS VDA and the Citrix Workspace App. The second hop must use a supported version of the multi-session OS VDA. See the [Product Matrix](#) page for supported versions.

For best performance and compatibility, Citrix recommends using a Citrix client of the same version or newer than the VDA versions in use.

In environments where the first hop involves a third-party (non-Citrix) virtual desktop solution in combination with a Citrix Virtual Apps session, support is limited to the Citrix Virtual Apps environment. In the event of any issues related to the third-party virtual desktop, including - but not limited to - Citrix Workspace app compatibility, redirection of hardware devices, and session performance, Citrix can provide technical support in a limited capacity. A Citrix Virtual Desktop at the first hop might be required as part of troubleshooting.

## Deployment considerations for HDX in double hop

In general, each session in a double hop is unique and client-server functions are isolated to a given hop. This section includes areas that require special consideration by Citrix administrators. Citrix recommends that customers conduct thorough testing of required HDX capabilities to ensure user experience and performance is adequate for a given environment configuration.

### Graphics

Use default graphics settings (selective encoding) on the first and second hops. In the case of [HDX 3D Pro](#), Citrix highly recommends that all applications that require graphics acceleration run locally in the first hop with the appropriate GPU resources available to the VDA.

### Latency

End-to-end latency can impact the overall user experience. Consider the added latency between the first and second hops. This is especially important with redirection of hardware devices.

### Multimedia

Server-side (in session) rendering of audio and video content performs best in the first hop. Video playback in the second hop requires decoding and re-encoding at the first hop, increasing bandwidth and hardware resource utilization as a result. Audio and video content must be limited to the first hop whenever possible.

### USB device redirection

HDX includes generic and optimized redirection modes to support a wide array of USB device types. Pay special attention to the mode in use at each hop and use the following table as reference for best results. For more information about generic and optimized redirection modes, see [Generic USB devices](#).

First hop (VDI or published desktop)	Second hop (Virtual apps)	Support notes
Optimized	Optimized	Recommended (based on device support). For example, USB mass storage, TWAIN scanners, Webcam, Audio.
Generic	Generic	For devices where the optimized option is not available.

First hop (VDI or published desktop)	Second hop (Virtual apps)	Support notes
Generic	Optimized	While technically possible, it is recommended to use the optimized mode across both hops when device support is available.
Optimized	Generic	Not supported

**Note:**

Due to the inherent chattiness of USB protocols, performance may decrease across hops. Functionality and results vary depending on specific device and application requirements. Validation testing is highly recommended in all cases of device redirection and especially important in double hop scenarios.

**Support exceptions**

Double hop sessions support most HDX features and capabilities except for the following:

- [Browser content redirection](#)
- [Local App Access](#)
- [RealTime Optimization Pack for Skype for Business](#)
- [Optimization for Microsoft Teams](#)

**Devices**

June 18, 2021

HDX provides a high-definition user experience on any device, at any location. The articles in the Devices section describe these devices:

- [Generic USB device](#)
- [Mobile and touch screen devices](#)
- [Serial devices](#)
- [Specialty keyboards](#)
- [TWAIN devices](#)
- [Webcams](#)
- [WIA devices](#)



## Optimized vs. generic USB device

An optimized USB device is one for which Citrix Workspace app has specific support. For example, the ability to redirect webcams using the HDX Multimedia virtual channel. A generic device is a USB device for which there is no specific support in Citrix Workspace app.

By default, generic USB redirection can't redirect USB devices with optimized virtual channel support unless put into Generic mode.

In general, you get better performance for USB devices in Optimized mode than in Generic mode. However, there are cases where a USB device doesn't have full functionality in Optimized mode. It might be necessary to switch to Generic mode to gain full access to its features.

With USB mass storage devices, you can use either client drive mapping or generic USB redirection, or both, controlled by Citrix policies. The main differences are:

If both generic USB redirection and the client drive mapping policies are enabled and a mass storage device is inserted either before or after a session starts, it's redirected using client drive mapping.

When these conditions are true, the mass storage device is redirected using generic USB redirection:

- Both generic USB redirection and the client drive mapping policies are enabled.
- A device is configured for automatic redirection.
- A mass storage device is inserted either before or after a session starts.

For more information, see <http://support.citrix.com/article/CTX123015>.

---

Feature	Client drive mapping	Generic USB redirection
Enabled by default	Yes	No
Read-only access configurable	Yes	No
Encrypted device access	Yes, if encryption is unlocked before the device is accessed on the virtual session.	Citrix Virtual Desktops only

---

## Mixed DPIs with multi-monitors

Citrix Virtual Apps and Desktops environments do not support the use of different DPIs between monitors. You can verify the DPI (% scaling) using Windows Control Panel > Display options. If using a Windows 8.1 or Windows 10 client device, enabling the **Let me choose one scaling level for all my displays option** in the Windows Control Panel > Display options configures the monitors appropriately. For more information, see Knowledge Center article [CTX201696](#).

## Generic USB devices

April 7, 2020

HDX technology provides **optimized support** for most popular USB devices. These devices include:

- Monitors
- Mice
- Keyboards
- Voice over Internet Protocol phones
- Headsets
- Webcams
- Scanners
- Cameras
- Printers
- Drives
- Smart card readers
- Drawing tablets
- Signature pads

Optimized support offers an improved user experience with better performance and bandwidth efficiency over a WAN. Optimized support is usually the best option, especially in high latency or security-sensitive environments.

HDX technology provides **generic USB redirection** for specialty devices that don't have optimized support or where it is unsuitable. For more information about generic USB redirection, see [Generic USB redirection](#).

For more information about USB devices and Citrix Workspace app for Windows, see [Configuring composite USB device redirection](#) and [Configuring USB support](#).

## Mobile and touch screen devices

June 18, 2021

### Tablet mode for touch screen devices using Windows Continuum

Continuum is a Windows 10 feature that adapts to the way the client device is used. This version of Continuum support, including dynamic change of modes, is available starting at VDA version 7.16 and Citrix Receiver for Windows version 4.10.

Windows 10 VDA detects the presence of a keyboard or mouse on a touch enabled client and puts the client in to desktop mode. If a keyboard or mouse is not present, Windows 10 VDA puts the client in to tablet/mobile mode. This detection occurs on connection and reconnection. It also occurs at dynamic attachment or detachment of the keyboard or mouse.

The feature is enabled by default. To disable this version of the feature, edit the [Tablet mode toggle policy settings](#) in the ICA policy settings article.

For the feature version included in XenApp 7.14 and 7.15 LTSR and XenDesktop 7.14 and 7.15 LTSR, use the registry settings to disable the feature. For more information, see [Tablet mode for touch screen devices](#).

The **tablet mode** offers a user interface that is better suited to touch screens:

- Slightly larger buttons.
- The Start screen and any apps you start open in a full screen.
- Taskbar contains a back button.
- Icons deleted from the task bar.

You have access to the File Explorer.

The **desktop mode** offers the traditional user interface where you interact in the same manner as using PC and a keyboard and mouse.

Tablet mode requires a minimum of version XenServer 7.2. XenServer 7.2 integrates with the Citrix Virtual Desktops VDA, changing the hypervisor to enable the virtual firmware settings for 2-in-1 devices. Windows 10 loads the GPIO driver on the target virtual machine based on this updated BIOS. It is used for toggling between tablet and desktop modes within the virtual machine. For more information, see the [release notes](#).

Citrix Workspace app for HTML5 (the light version) does not support Windows Continuum features.



Run the XenServer CLI command to allow laptop/tablet switching:

**xe vm-param-set uuid=<VM\_UUID> platform:acpi\_laptop\_slate=1**

**Important:**

Updating the base image for an existing machine catalog after changing the metadata setting doesn't affect any previously provisioned VMs. After changing the XenServer VM base image, create a catalog, choose the base image, and provision a new Machine Creation Services (MCS) machine.

**Before starting a session:**

We recommend that you navigate to **Settings > System > Tablet Mode** on the VDA before starting a session and set the following options from the drop-down menus:

- Use the appropriate mode for my hardware
- Don't ask me and always switch

If you don't set these options before starting the session, set the options after you start the session and restart the VDA.

## Tablet mode

When I sign in

Use the appropriate mode for my hardware ▾

When this device automatically switches tablet mode on or off

Don't ask me and always switch ▾

### Microsoft Surface Pro and Surface Book pens

We support standard pen functionality with Windows Ink-based applications. This functionality requires a Virtual Delivery Agent running on a minimum of Microsoft Windows 10 version 1809 and client devices using a minimum of Citrix Workspace app for Windows version 1902. Support includes pointing, erasing, pen pressure, Bluetooth signals, and other features depending on the operating system firmware and pen model. For example, pen pressure can be up to 4096 levels. This feature is enabled by default.

For a demonstration of Windows Ink and the pen functionality, click this graphic:



### System requirements

- Citrix Virtual Apps and Desktops minimum version 1903
- Citrix Workspace app for Windows minimum version 1902
- Microsoft Windows 10 minimum version 1809

To disable or enable this feature, see [Microsoft Surface Pro and Surface Book pens](#) in the list of features managed through the registry.

### Serial ports

July 8, 2021

Most new PCs don't have built-in serial (COM) ports. The ports are easy to add by using USB converters. Applications suited for serial ports often involve sensors, controllers, old check readers, pads, and so forth. Some USB virtual COM-port devices use vendor-specific drivers in place of the Windows-provided drivers (usbser.sys). These drivers allow you to force the virtual COM port of the USB device so that it doesn't change even if connected to different USB sockets. This might be done from the **Device Manager > Ports (COM & LPT) > Properties** or from the application that controls the device.

Client COM port mapping allows devices attached to the COM ports on the user's endpoint to be used during virtual sessions. You can use these mappings like any other network mappings.

For each COM port, a driver in the operating system assigns a symbolic link name such as COM1 and COM2. The applications then use the link to access the port.

**Important:**

Because a device can attach to the endpoint by using USB directly, doesn't mean it can be redirected using generic USB redirection. Some USB devices function as virtual COM ports, which applications can access in the same way as physical serial port. The operating system can abstract COM ports and treat them like fileshares. Two common protocols for virtual COM are CDC ACM or MCT. When connected through an RS-485 port, applications might not work at all. Get an RS-485-to-RS232 converter to use RS-485 as a COM port.

**Important:**

Some applications recognize the device (for example, a signature pad) consistently only if it is connected to COM1 or COM2 on the client workstation.

## Map a client COM port to a server COM port

You can map client COM ports to a Citrix session in three ways:

- Manage console policies. For more information about policies, see [Port redirection policy settings](#).
  - VDA command prompt.
  - Remote Desktop (Terminal Services) configuration tool.
1. Enable the **Client COM port redirection** and the **Auto connect client COM ports Studio** policies. After applied, some information is available in HDX Monitor.

Name	Value
HardwareId	1591092831
InternetClient	False
LastError	
Name	FTLLFERNANDOK02
Policy_AutoConnectClientComPorts	False
Policy_AutoConnectClientLptPorts	False
...	*

2. If **Auto connect client COM ports** failed to map the port, you can map the port manually or use logon scripts. Log on to the VDA, and at a command prompt window, type:

```
NET USE COMX: \\CLIENT\COMZ:
```

Or

NET USE COMX: \\CLIENT\CLIENTPORT:COMZ:

**X** is the number of the COM port on the VDA (ports 1 through 9 are available for mapping). **Z** is the number of the client COM port you want to map.

To confirm that the operation was successful, type **NET USE** at a VDA command prompt. The list that appears contains mapped drives, LPT ports, and mapped COM ports.

```
C:\Windows\system32>net use
New connections will be remembered.
```

Status	Local	Remote	Network
	COM3	\\Client\COM3:	Citrix Client Network

- To use this COM port in a virtual desktop or application, install your user device application and point it to the mapped COM port name. For example, if you map COM1 on the client to COM3 on the server, install your COM port device application in the VDA and point it to COM3 during the session. Use this mapped COM port as you would a COM port on the user device.

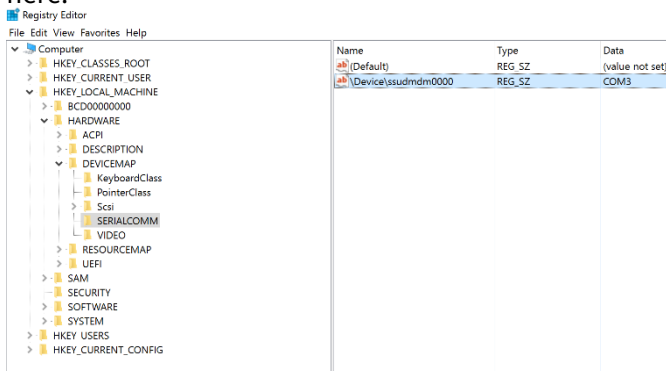
**Important:**

COM port mapping is not TAPI-compatible. You can't map Windows Telephony Application Programming Interface (TAPI) devices to client COM ports. TAPI defines a standard way for applications to control telephone functions for data, fax, and voice calls. TAPI manages signaling, including dialing, answering, and ending calls. Also, supplemental services such as holding, transferring, and conference calls.

**Troubleshoot**

- Ensure you can access the device directly from the endpoint, bypassing Citrix. While the port is not mapped to the VDA, you are not connected to a Citrix session. Follow any troubleshooting instructions that came with the device and verify that it works locally first.

When a device is connected to a serial COM port, a registry key is created on the hive shown here:



You can also find this information from the command prompt by running **chgport /query**.

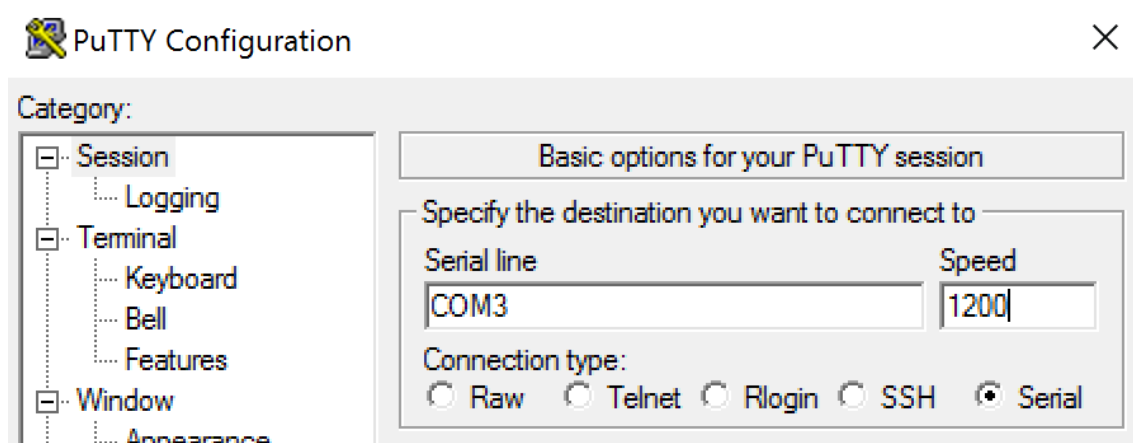
```
C:\Windows\system32\cmd.exe
C:\Users\fernandok>chgport /query
COM3 = \Device\ssudmdm0000

C:\Users\fernandok>mode

Status for device COM3:
-----
      Baud:                1200
      Parity:               Even
      Data Bits:           7
      Stop Bits:           1
      Timeout:              OFF
      XON/XOFF:             OFF
      CTS handshaking:     OFF
      DSR handshaking:     OFF
      DSR sensitivity:     OFF
      DTR circuit:         ON
      RTS circuit:         ON
```

If troubleshooting instructions for the device aren't available, try opening a PuTTY session. Choose **Session** and in **Serial line** specify your COM Port.





You can run **MODE** in a local command window. The output might display the COM port in use and the Baud/Parity/Data Bits/Stop Bits, which you need in your PuTTY session. If the PuTTY connection is successful, press **Enter** to see feedback from the device. Whatever characters you type might be repeated on the screen, or responded to. If this step is unsuccessful, you can't access the device from a virtual session.

2. Map the local COM port to the VDA (using policies or **NET USE COMX: \\CLIENT\COMZ:**) and repeat the same PuTTY procedures in the previous step, but this time from the VDA PuTTY. If PuTTY fails to show the error **Unable to open connection to COM1. Unable to open serial port**, another device might be using COM1.
3. Run **chgbport /query**. If the built-in Windows serial driver on the VDA is auto-assigning \Device\Serial0 to a COM1 port of your VDA, do the following:
  - A. Open CMD on the VDA and type **NET USE**.
  - B. Delete any existing mapping (for example, COM1) on the VDA.

#### **NET USE COM1 /DELETE**

- C. Map the device to the VDA.

#### **NET USE COM1: \\CLIENT\COM3:**

- D. Point your application on the VDA to COM3.

Lastly, try to map your local COM port (for example, COM3) to a different COM port on the VDA (other than COM1, for example COM3). Ensure that your application is pointing to it:

#### **NET USE COM3: \\CLIENT\COM3**

4. If now you do see the port mapped, PuTTY is working but no data passing, it might be a race condition. The application might connect and open the port before it is mapped, locking it from being mapped. Try one of the following:
  - Open a second application published on the same server. Wait a few seconds for the port to be mapped, and then open the real application that tries to use the port.

- Enable the COM port redirection policies from the Group Policy Editor in Active Directory instead of the service's Manage > Full Configuration interface. Those policies are **Client COM port redirection** and **Auto connect client COM ports**. Policies applied this way might be processed before the Manage console policies, guaranteeing that the COM port is mapped. Citrix policies are pushed to the VDA and stored in:  
`HKLN\SOFTWARE\Policies\Citrix \<user session ID\>`
- Use this logon script for the user or instead of publishing the application, publish a .bat script that first deletes any mapping on the VDA, remaps the virtual COM port, and then starts the application:

```
@echo off
NET USE COM1 /delete
NET USE COM2 /delete
NET USE COM1: \\CLIENT\COM1:
NET USE COM2: \\CLIENT\COM2:
MODE COM1: BAUD=1200 (or whatever value needed)
MODE COM2: BAUD=9600 PARITY=N Data=8 Stop=1 (or whatever value needed)
START C:\Program Files\<Your Software Path>\<your_software.exe>
```

5. Process Monitor from Sysinternals is the tool of last resort. When running the tool on the VDA, find and filter objects like COM3, picaser.sys, CdmRedirector, but especially <your\_app>.exe. Any errors might appear as Access Denied or similar.

## Specialty keyboards

June 18, 2021

### Bloomberg keyboards

**Warning:**

Editing the registry incorrectly can cause serious problems that might require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

Citrix Virtual Apps and Desktops support the Bloomberg model 4 Starboard keyboard (and earlier model 3). This keyboard enables customers in the financial sector to use the special features of the keyboard to access financial market data and perform trading quickly.

This keyboard is compatible with the KVM switch boxes and can work in two modes:

- PC (One USB cable with no KVM)
- KVM mode (Two USB Cables with one routed through KVM)

**Important:**

We recommend that you use the Bloomberg keyboard with only one session. We don't recommend using the keyboard with multiple concurrent sessions (one client to multiple sessions).

The Bloomberg keyboard 4 is a USB composite device comprising four USB devices in one physical shell:

- Keyboard.
- Fingerprint reader.
- Audio device with keys to increase and decrease volume and mute the speaker and the microphone. This device includes onboard speaker, microphone, and jack for the microphone and headset.
- USB hub to connect all of these devices to the system.

**Requirements:**

- The session to which Citrix Workspace app for Windows is connecting must support USB devices.
- Minimum of Citrix Workspace app 1808 for Windows or Citrix Receiver for Windows 4.8 to support Bloomberg keyboard model 3 and 4.
- Minimum of Citrix Workspace app 1808 for Windows or Citrix Receiver for Windows 4.12 to use KVM mode (two USB cables with one routed through KVM) for Model 4.

For information about configuring Bloomberg keyboards on Citrix Workspace app for Windows, see [Configuring Bloomberg keyboards](#).

To enable Bloomberg keyboard support, see [Bloomberg keyboards](#) in the list of features managed through the registry.

**Verify support:**

To determine if Bloomberg keyboard support is enabled in Citrix Workspace app, check if the Desktop Viewer correctly reports the Bloomberg keyboard's devices.

Desktop scenario:

Open the Desktop Viewer. If support for Bloomberg keyboard is enabled, the Desktop Viewer shows see three devices under the USB icon:

- Bloomberg Fingerprint Scanner
- Bloomberg Keyboard Features
- Bloomberg LP Keyboard 2013

Seamless Application only scenario:

Open the **Connection Center** menu from the Citrix Workspace app notification area icon. If support for the Bloomberg keyboard is enabled, the three devices appear in the **Devices** menu.

The check mark against each of these devices indicates that they are remoted to the session.

## TWAIN devices

April 6, 2020

### Requirements

- The scanner must be TWAIN compliant.
- Install the TWAIN drivers on the local device. They are not required on the server.
- Attach the scanner locally (for example, through USB).
- Ensure that the scanner is using the local TWAIN driver and not the Windows Image Acquisition service.
- Ensure that there is no policy applied to the user account that is used for the test, and which is limiting the bandwidth within the ICA session. For example, client USB redirection bandwidth limit.

For information about policy settings, see [TWAIN devices policy settings](#).

## Webcams

June 18, 2021

### High definition webcam streaming

Webcams can be used by video conferencing applications running within the virtual session. The application on the server selects the webcam format and resolution based on the supported format types. When a session starts, the client sends the webcam information to the server. Choose a webcam from the video conferencing application. When the webcam and the application both support high-definition rendering, the application uses high-definition resolution. We support webcam resolutions up to 1920x1080.

This feature requires the Citrix Receiver for Windows, minimum version 4.10. For a list of Citrix Workspace app platforms that support HDX webcam redirection, see [Citrix Workspace app feature matrix](#).

For more information about high-definition webcam streaming, see [HDX video conferencing and webcam video compression](#).

You can use a registry key to disable and enable the feature and then configure a specific resolution. For information, see [High-definition webcam streaming and High-definition webcam resolution](#) in the list of features managed through the registry.

## WIA devices

May 27, 2021

### Requirements

- The scanner must be WIA compliant.
- Install the WIA drivers on the local device. They are not required on the server.
- Attach the scanner locally (for example, through USB).
- Ensure that the scanner is using the local Windows Image Acquisition service and not the TWAIN driver.
- Ensure that there is no policy applied to the user account that is used for the test, and which is limiting the bandwidth within the ICA session. For example, client USB redirection bandwidth limit.

### Windows Image Acquisition application allow list

An allow list lets you control which applications on the VDA can access the Windows Image Acquisition scanner redirection. The Registry Editor uses input from the allow list setting on each VDA that contains Windows Image Acquisition. By default, no applications have access to Windows Image Acquisition.

To adjust Windows Image Acquisition for applications on the VDA, see the [Windows Image Acquisition application allow list](#) setting in the list of features managed through the registry.

For information about policy settings, see [WIA devices policy settings](#).

## Graphics

September 3, 2020

Citrix HDX graphics include an extensive set of graphics acceleration and encoding technologies that optimizes the delivery of rich graphics applications from Citrix Virtual Apps and Desktops. The graphic technologies provide the same experience as using a physical desktop when working remotely with virtual applications that are graphics intensive.

You can use software or hardware for graphics rendering. Software rendering requires a third-party library called software rasterizer. For example, Windows includes the WARP rasterizer for DirectX based graphics. Sometimes, you might want to use an alternative software renderer. Hardware rendering (hardware acceleration) requires a graphics processor (GPU).

HDX Graphics offers a default encoding configuration that is optimized for the most common use cases. By using Citrix policies, IT administrators can also configure various graphics-related settings to meet different requirements and provide the desired user experience.

### **Thinwire**

Thinwire is the Citrix default display remoting technology used in Citrix Virtual Apps and Desktops.

Display remoting technology allows graphics generated on one machine to be transmitted, typically across a network, to another machine for display. Graphics are generated as a result of user input, for example, keystrokes or mouse actions.

### **HDX 3D Pro**

The HDX 3D Pro capabilities in Citrix Virtual Apps and Desktops enable you to deliver desktops and applications that perform best using a graphics processing unit (GPU) for hardware acceleration. These applications include 3D professional graphics applications based on OpenGL and DirectX. The standard VDA supports GPU acceleration of DirectX only.

### **GPU acceleration for Windows Single-session OS**

By using HDX 3D Pro, you can deliver graphically intensive applications as part of hosted desktops or applications on Single-session OS machines. HDX 3D Pro supports physical host computers (including desktop, blade, and rack workstations) and GPU Passthrough and GPU virtualization technologies offered by XenServer, vSphere, and Hyper-V (passthrough only) hypervisors.

Using GPU Passthrough, you can create VMs that have exclusive access to dedicated graphics processing hardware. You can install multiple GPUs on the hypervisor and assign VMs to each of these GPUs on a one-to-one basis.

Using GPU virtualization, multiple virtual machines can directly access the graphics processing power of a single physical GPU.

### **GPU acceleration for Windows Multi-session OS**

HDX 3D Pro allows graphics-heavy applications running in Windows Multi-session OS sessions to render on the server graphics processing unit (GPU). By moving OpenGL, DirectX, Direct3D, and Windows Presentation Foundation (WPF) rendering to the server GPU, graphics rendering doesn't slow down

the server CPU. Also, the server is able to process more graphics because the workload is split among the CPU and GPU.

### Framehawk

**Important:**

As of Citrix Virtual Apps and Desktops 7 1903, Framehawk is no longer supported. Instead, use [Thinwire](#) with [adaptive transport](#) enabled.

Framehawk is a display remoting technology for mobile workers on broadband wireless connections (Wi-Fi and 4G/LTE cellular networks). Framehawk overcomes the challenges of spectral interference and multipath propagation and delivers a fluid and interactive user experience to users of virtual apps and desktops.

### Text-based session watermark

Text-based session watermarks help to deter and enable tracking data theft. This traceable information appears on the session desktop as a deterrent to those using photographs and screen captures to steal data. You can specify a watermark that is a layer of text. The watermark can display over the entire session screen without changing the content of the original document. Text-based session watermarks require VDA support.

### Related information

- [HDX 3D Pro](#)
- [GPU acceleration for Windows Single-session OS](#)
- [GPU acceleration for Windows Multi-session OS](#)
- [Thinwire](#)
- [Text-based session watermark](#)

## HDX 3D Pro

July 8, 2020

The HDX 3D Pro capabilities of Citrix Virtual Apps and Desktops enable you to deliver desktops and applications that perform best using a graphics processing unit (GPU) for hardware acceleration. These applications include 3D professional graphics applications based on OpenGL and DirectX. The standard VDA supports GPU acceleration of DirectX only.

For the HDX 3D Pro policy settings, see [Optimize for 3D graphics workload](#).

All supported Citrix Workspace apps can be used with 3D graphics. For best performance with complex 3D workloads, high-resolution monitors, multi-monitor configurations, and high frame rate applica-

tions, we recommend the latest versions of Citrix Workspace app for Windows and Citrix Workspace app for Linux. For more information on supported versions of Citrix Workspace app, see [Lifecycle Milestones for Citrix Workspace app](#).

Examples of 3D professional applications include:

- Computer-aided design, manufacturing, and engineering (CAD/CAM/CAE) applications
- Geographical Information System (GIS) software
- Picture Archiving Communication System (PACS) for medical imaging
- Applications using the latest OpenGL, DirectX, NVIDIA CUDA, and OpenCL and WebGL versions
- Computationally intensive non-graphical applications that use NVIDIA Compute Unified Device Architecture (CUDA) GPUs for parallel computing

HDX 3D Pro provides the best user experience over any bandwidth:

- On WAN connections: Deliver an interactive user experience over WAN connections with bandwidths as low as 1.5 Mbps.
- On LAN connections: Deliver a user experience equivalent to that of a local desktop on LAN connections.

You can replace complex and expensive workstations with simpler user devices by moving the graphics processing into the data center for centralized management.

HDX 3D Pro provides GPU acceleration for Windows single-session OS machines and Windows multi-session OS machines. For more information, see [GPU acceleration for Windows single-session OS](#) and [GPU acceleration for Windows multi-session OS](#).

HDX 3D Pro is compatible with GPU passthrough and GPU virtualization technologies offered by the following hypervisors, in addition to bare metal:

- Citrix Hypervisor
  - GPU passthrough with NVIDIA GRID, AMD, and Intel GVT-d
  - GPU virtualization with NVIDIA GRID, AMD, and Intel GVT-g
  - See hardware compatibility at [Hypervisor Hardware Compatibility List](#).

Use the HDX Monitor tool to validate the operation and configuration of HDX visualization technologies and to diagnose and troubleshoot HDX issues. To download the tool and learn more about it, see <https://taas.citrix.com/hdx/download/>.

## **GPU acceleration for Windows multi-session OS**

June 18, 2021



HDX 3D Pro allows graphics-heavy applications running in Windows Multi-session OS sessions to render on the server's graphics processing unit (GPU). By moving OpenGL, DirectX, Direct3D, and Windows Presentation Foundation (WPF) rendering to the server's GPU, graphics rendering does not slow the server's CPU. Also, the server is able to process more graphics because the workload is split between the CPU and GPU.

Since Windows Server is a multi-user operating system, multiple users can share a GPU accessed by Citrix Virtual Apps without the need for GPU virtualization (vGPU).

For procedures that involve editing the registry, use caution: Editing the registry incorrectly can cause serious problems that may require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

## GPU sharing

GPU Sharing enables GPU hardware rendering of OpenGL and DirectX applications in remote desktop sessions. It has the following characteristics:

- Can be used on bare metal or virtual machines to increase application scalability and performance.
- Enables multiple concurrent sessions to share GPU resources (most users do not require the rendering performance of a dedicated GPU).
- Requires no special settings.

A GPU can be assigned to the Windows Server virtual machine in either full pass-through or virtual GPU (vGPU) modes following Hypervisor and GPU vendor requirements. Bare-metal deployments on physical Windows Server machines are also supported.

GPU Sharing does not depend on any specific graphics card.

- For virtual machines, select a graphics card that is compatible with the Hypervisor in use. For a Citrix Hypervisor hardware compatibility list, see [Hypervisor Hardware Compatibility List](#).
- When running on bare metal, it is recommended to have a single display adapter enabled by the operating system. If multiple GPUs are installed on the hardware, disable all but one of them using Device Manager.

Scalability using GPU Sharing depends on several factors:

- The applications being run
- The amount of video RAM they consume
- The graphics card's processing power

Some applications handle video RAM shortages better than others. If the hardware becomes overloaded, instability or a crash of the graphics card driver might occur. Limit the number of concurrent users to avoid such issues.

To confirm that GPU acceleration is occurring, use a third-party tool such as GPU-Z. GPU-Z is available at <http://www.techpowerup.com/gpuz/>.

- Access to a high-performance video encoder for NVIDIA GPUs and Intel Iris Pro graphics processors. A policy setting (enabled by default) controls this feature and allows the use of hardware encoding for H.264 encoding (where available). If such hardware is not available, the VDA falls back to CPU-based encoding using the software video codec. For more information, see [Graphics policy settings](#).

## DirectX, Direct3D, and WPF rendering

DirectX, Direct3D, and WPF rendering are only available on servers with a GPU that supports a display driver interface (DDI) version of 9ex, 10, or 11.

- On Windows Server 2008 R2, DirectX and Direct3D require no special settings to use a single GPU.
- On Windows Server 2012 and later, Remote Desktop Services (RDS) sessions on the RD Session Host server use the Microsoft Basic Render Driver as the default adapter. To use the GPU in RDS sessions on Windows Server 2012 and later, enable the **Use the hardware default graphics adapter for all Remote Desktop Services sessions** setting in the group policy **Local Computer Policy > Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Remote Session Environment**.
- To enable WPF applications to render using the server's GPU, create the settings in the registry of the server running Windows Multi-session OS sessions. For information on the registry setting, see [Windows Presentation Foundation \(WPF\) rendering](#) in the list of features managed through the registry.

## GPU acceleration for CUDA or OpenCL applications

GPU acceleration of CUDA and OpenCL applications running in a user session is disabled by default.

To use the CUDA acceleration POC features, enable the registry settings. For information, see [GPU acceleration for CUDA or OpenCL applications](#) in the list of features managed through the registry.

## GPU acceleration for Windows single-session OS

March 17, 2021

With HDX 3D Pro, you can deliver graphically intensive applications as part of hosted desktops or applications on Single-session OS machines. HDX 3D Pro supports physical host computers (including

desktop, blade, and rack workstations) and GPU Passthrough and GPU virtualization technologies offered by Citrix Hypervisor, vSphere, Nutanix, and Hyper-V (passthrough only) hypervisors.

HDX 3D Pro offers the following features:

- Adaptive H.264-based or H.265-based deep compression for optimal WAN and wireless performance. HDX 3D Pro uses CPU-based full-screen H.264 compression as the default compression technique for encoding. Hardware encoding with H.264 is used with NVIDIA, Intel, and AMD cards that support NVENC. Hardware encoding with H.265 is used with NVIDIA cards that support NVENC.
- Lossless compression option for specialized use cases. HDX 3D Pro also offers a CPU-based lossless codec to support applications where pixel-perfect graphics are required, such as medical imaging. True lossless compression is recommended only for specialized use cases because it consumes more network and processing resources.

When using lossless compression:

- The lossless indicator, a notification area icon, notifies the user if the screen displayed is a lossy frame or a lossless frame. This icon helps when the **Visual Quality** policy setting specifies **Build to lossless**. The lossless indicator turns green when the frames sent are lossless.
- The lossless switch enables the user to change to Always Lossless mode anytime within the session. To select or deselect **Lossless anytime within a session**, right-click the icon and click **Switch to pixel perfect** or use the shortcut ALT+SHIFT+1.

For lossless compression: HDX 3D Pro uses the lossless codec for compression regardless of the codec selected through policy.

For lossy compression: HDX 3D Pro uses the original codec, either the default or the one selected through policy.

Lossless switch settings are not retained for subsequent sessions. To use a lossless codec for every connection, select **Always lossless** in the **Visual quality** policy setting.

- You can override the default shortcut, ALT+SHIFT+1, to select or deselect Lossless within a session. Configure a new registry setting at HKEY\_LOCAL\_MACHINE\SOFTWARE\Citrix\HDX3D\LLIndicator.
  - Name: HKEY\_LOCAL\_MACHINE\_HotKey, Type: String
  - The format to configure a shortcut combination is C=0|1, A=0|1, S=0|1, W=0|1, K=val. Keys must be comma “,” separated. The order of the keys does not matter.
  - A, C, S, W and K are keys, where C=Control, A=ALT, S=SHIFT, W=Win, and K=a valid key. Allowed values for K are 0–9, a–z, and any virtual key code.
  - For example:
    - \* For F10, set K=0x79
    - \* For Ctrl + F10, set C=1, K=0x79

- \* For Alt + A, set A=1, K=a or A=1, K=A or K=A, A=1
- \* For Ctrl + Alt + 5, set C=1, A=1, K=5 or A=1, K=5, C=1
- \* For Ctrl + Shift + F5, set A=1, S=1, K=0x74

**Caution:**

Editing the registry incorrectly can cause serious problems that might require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

- Multiple and high resolution monitor support. For Single-session OS machines, HDX 3D Pro supports user devices with up to four monitors. Users can arrange their monitors in any configuration and can mix monitors with different resolutions and orientations. The number of monitors is limited by the capabilities of the host computer GPU, the user device, and the available bandwidth. HDX 3D Pro supports all monitor resolutions and is limited only by the capabilities of the GPU on the host computer.
- Dynamic resolution. You can resize the virtual desktop or application window to any resolution. **Note:** The only supported method to change the resolution is by resizing the VDA session window. Changing resolution from within the VDA session (using **Control Panel > Appearance and Personalization > Display > Screen Resolution**) is not supported.
- Support for NVIDIA vGPU architecture. HDX 3D Pro supports NVIDIA vGPU cards. For information, see [NVIDIA vGPU](#) for GPU passthrough and GPU sharing. NVIDIA vGPU enables multiple VMs to have simultaneous, direct access to a single physical GPU, using the same NVIDIA graphics drivers that are deployed on non-virtualized operating systems.
- Support for VMware vSphere and VMware ESX using Virtual Direct Graphics Acceleration (vDGA)
  - You can use HDX 3D Pro with vDGA for both RDS and VDI workloads.
- Support for VMware vSphere/ESX using NVIDIA vGPU and AMD MxGPU.
- Support for Microsoft HyperV using Discrete Device Assignment in Windows Server 2016.
- Support for Data Center Graphics with Intel Xeon Processor E3 Family. HDX 3D Pro supports multi-monitors (up to 3), console blanking, custom resolution, and high frame-rate with the supported family of Intel processors. For more information, see <http://www.citrix.com/intel> and <http://www.intel.com/content/www/us/en/servers/data-center-graphics.html>.
- Support for AMD RapidFire on the AMD FirePro S-series server cards. HDX 3D Pro supports multi-monitors (up to 6), console blanking, custom resolution, and high frame-rate. Note: HDX 3D Pro support for AMD MxGPU (GPU virtualization) works with VMware vSphere vGPUs only. Citrix Hypervisor and Hyper-V are supported with GPU passthrough. For more information, see [AMD Virtualization Solution](#).

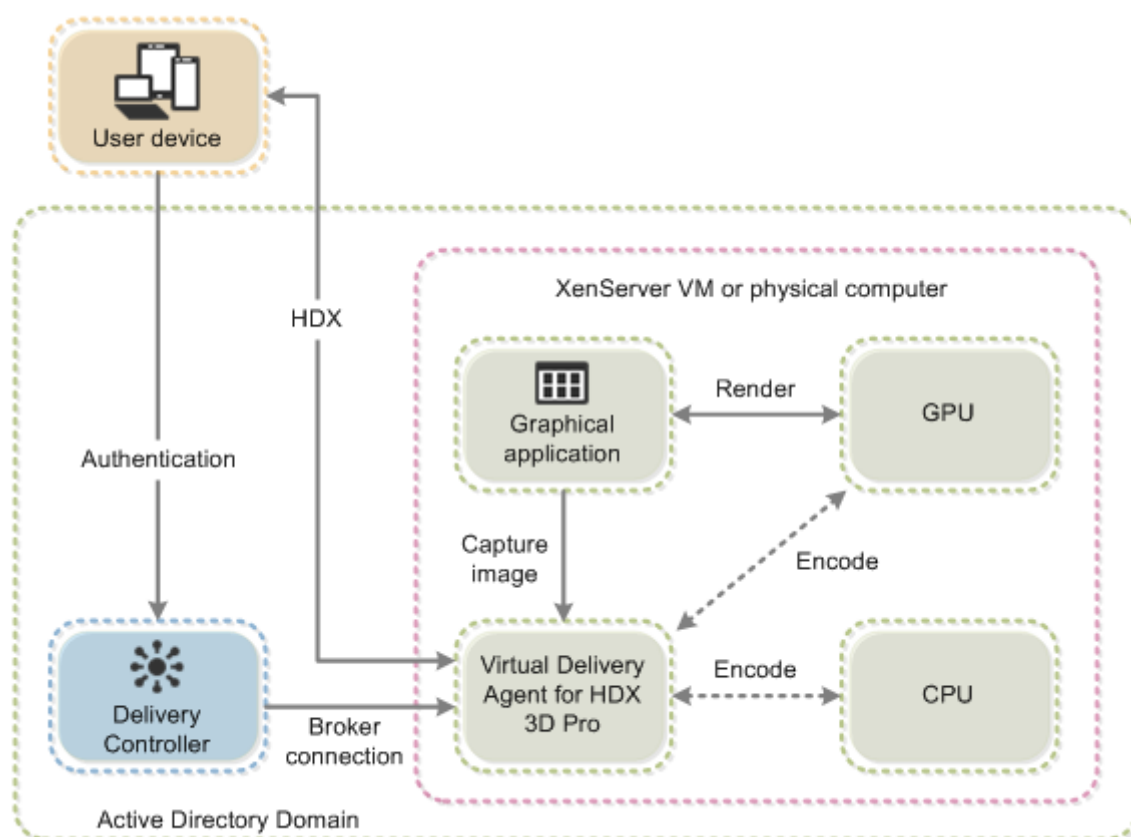
- Access to a high-performance video encoder for NVIDIA GPUs, AMD GPUs, and Intel Iris Pro graphics processors. A policy setting (enabled by default) controls this feature. The feature allows the use of hardware encoding for H.264 encoding (where available). If such hardware is not available, the VDA falls back to CPU-based encoding using the software video codec. For more information, see [Graphics policy settings](#).

As shown in the following figure:

- When a user logs on to Citrix Workspace app and accesses the virtual application or desktop, the Controller authenticates the user. The Controller then contacts the VDA for HDX 3D Pro to broker a connection to the computer hosting the graphical application.

The VDA for HDX 3D Pro uses the appropriate hardware on the host to compress views of the complete desktop or of just the graphical application.

- The desktop or application views and the user interactions with them are transmitted between the host computer and the user device. This transmission is done through a direct HDX connection between Citrix Workspace app and the VDA for HDX 3D Pro.



### Optimize the HDX 3D Pro user experience

To use HDX 3D Pro with multiple monitors, ensure that the host computer is configured with at least as many monitors as are attached to user devices. The monitors attached to the host computer can

be either physical or virtual.

Do not attach a monitor (either physical or virtual) to a host computer while a user is connected to the virtual desktop or application providing the graphical application. Doing so can cause instability during a user's session.

Let your users know that changes to the desktop resolution (by them or an application) are not supported while a graphical application session is running. After closing the application session, a user can change the resolution of the Desktop Viewer window in the Citrix Workspace app - Desktop Viewer Preferences.

When multiple users share a connection with limited bandwidth (for example, at a branch office), we recommend that you use the **Overall session bandwidth limit** policy setting to limit the bandwidth available to each user. Using this setting ensures that the available bandwidth does not fluctuate widely as users log on and off. Because HDX 3D Pro automatically adjusts to use all the available bandwidth, large variations in the available bandwidth over the course of user sessions can negatively impact performance.

For example, if 20 users share a 60 Mbps connection, the bandwidth available to each user can vary between 3 Mbps and 60 Mbps, depending on the number of concurrent users. To optimize the user experience in this scenario, determine the bandwidth required per user at peak periods and limit users to this amount always.

For users of a 3D mouse, we recommend that you increase the priority of the Generic USB Redirection virtual channel to 0. For information about changing the virtual channel priority, see the Knowledge Center article [CTX128190](#).

## Thinwire

June 18, 2021

### Introduction

Thinwire, a part of Citrix HDX technology, is the Citrix default display remoting technology used in Citrix Virtual Apps and Desktops.

Display remoting technology allows graphics generated on one machine to be transmitted, typically across a network, to another machine for display.

A successful display remoting solution provides a highly interactive user experience that is similar to that of a local PC. Thinwire achieves this experience by using a range of complex and efficient image analysis and compression techniques. Thinwire maximizes server scalability and consumes less bandwidth than other display remoting technologies.

Because of this balance, Thinwire meets most general business use cases and is used as the default display remoting technology in Citrix Virtual Apps and Desktops.

## HDX 3D Pro

In its default configuration, Thinwire can deliver 3D or highly interactive graphics and use a graphics processing unit (GPU), if present. However, we recommend enabling HDX 3D Pro mode using the **Optimize for 3D graphics workload** or **Visual quality > Build to lossless** policies for scenarios when GPUs are present. These policies configure Thinwire to use a video codec (H.264 or H.265) to encode the entire screen using hardware acceleration if a GPU is present. Doing so provides a more fluid experience for 3D professional graphics. For more information, see [H.264 Build to lossless](#), [HDX 3D Pro](#), and [GPU acceleration for Windows Single-session OS](#).

## Requirements

Thinwire is optimized for modern operating systems, including Windows Server 2012 R2, Windows Server 2016, Windows Server 2019, Windows 7, and Windows 10. For Windows Server 2008 R2, legacy graphics mode is recommended. Use the built-in [Citrix policy templates](#), High Server Scalability-Legacy OS and Optimized for WAN-Legacy OS to deliver the Citrix recommended combinations of policy settings for these use cases.

### Note:

We do not support legacy graphics mode in this release. It is included for backward compatibility when using XenApp 7.15 LTSR, XenDesktop 7.15 LTSR, and previous VDA releases with Windows 7 and Windows 2008 R2.

- The policy setting which drives the behavior of Thinwire, **Use video codec for compression**, is available on VDA versions in Citrix Virtual Apps and Desktops 7 1808 or later and XenApp and XenDesktop 7.6 FP3 and later. The **Use video codec when preferred** option is the default setting on VDA versions Citrix Virtual Apps and Desktops 7 1808 or later and XenApp and XenDesktop 7.9 and later.
- All Citrix Workspace apps support Thinwire. Some Citrix Workspace apps might support features of Thinwire that others do not, for example, 8-bit or 16-bit graphics for reduced bandwidth usage. Support for such features is automatically negotiated by Citrix Workspace app.
- Thinwire uses more server resources (CPU, memory) in multi-monitor and high-resolution scenarios. It is possible to tune the amount of resources Thinwire uses, however, bandwidth usage might increase as a result.
- In low bandwidth or high latency scenarios, consider enabling 8-bit or 16-bit graphics to improve interactivity. Visual quality might be affected, especially at 8-bit color depth.

## Encoding methods

Thinwire can operate in two different encoding modes depending on policy and client capabilities:

- Thinwire full screen H.264 or H.265
- Thinwire with selective H.264 or H.265

Legacy GDI remoting uses the XPDM remoting driver and not a Thinwire bitmap encoder.


## Configuration

Thinwire is the default display remoting technology.

The following Graphics policy setting sets the default and provides alternatives for different use cases:

- [Use video codec for compression](#)
  - **Use video codec when preferred.** This is the default setting. No additional configuration is required. Keeping this setting as the default ensures that Thinwire is selected for all Citrix connections, and is optimized for scalability, bandwidth, and superior image quality for typical desktop workloads. This is functionally equivalent to **For actively changing regions**.
  - Other options in this policy setting continue to use Thinwire with other technologies for different use cases. For example:
    - **For actively changing regions.** The adaptive display technology in Thinwire identifies moving images (video, 3D in motion) and uses H.264 or H.265 only in the part of the screen where the image is moving.
    - **For the entire screen.** Delivers Thinwire with full-screen H.264 or H.265 to optimize for improved user experience and bandwidth in cases with heavy use of 3D graphics. In the case of H.264 4:2:0 (the **Visually lossless** policy is disabled), the final image is not pixel perfect (lossless) and might not be suitable for certain scenarios. In such cases, consider using [H.264 Build to lossless](#) instead.





Various other policy settings, including the following Visual display policy settings can be used to fine tune the performance of display remoting technology. Thinwire supports them all.

- [Preferred color depth for simple graphics](#)
- [Target frame rate](#)
- [Visual quality](#)

To get the Citrix recommended combinations of policy settings for different business use cases, use the built-in [Citrix Policy templates](#). The **High Server Scalability** and **Very High Definition User Experience** templates both use Thinwire with the optimum combinations of policy settings for your organization's priorities and your users' expectations.

## Monitoring Thinwire

You can monitor the use and performance of Thinwire from Citrix Director. The HDX virtual channel details view contains useful information for troubleshooting and monitoring Thinwire in any session. To view Thinwire-related metrics:

1. In Director, search for a user, machine or endpoint, open an active session and click **Details**. Or, you can select **Filters > Session > All Sessions**, open an active session and click **Details**.

2. Scroll down to the **HDX** panel.

HDX

Download System Report

	Adobe® Flash®	Virtual channel: Idle Flash redirection: Inactive
	Graphics - Framehawk	Virtual channel: Idle Current FPS: 0
	Scanner	Virtual channel: Idle Compression level: Medium
	Smart Cards	Virtual channel: Idle Number of devices: 0
	Legacy Graphics	Virtual channel: Active Still image compression: Medium
	Audio	Virtual channel: Idle Number of devices: 1
	Graphics - Thinwire	Virtual channel: Active Current FPS: 1
	Mapped Client Drives	Virtual channel: Idle Client drives available: 0
	Network	Bandwidth used: 0% Average latency: 47 ms
	Printing	Mapped printers: 4 Virtual channel: Idle
	VDA	Version: Session ID: 3
	Windows Media	Virtual channel: Idle Active streams: 2

3. Select **Graphics - Thinwire**.

Graphics - Thinwire

There are no alerts at this time.

▼ Status

Virtual channel state	Idle
Virtual channel priority	High
Monitors	1
Frames Per Second	1
Provider	Standard (RDS)
Video codec use	None

**Monitor 0**

Monitor ID	0
Primary	True
Left	0
Top	0
Right	1280
Bottom	800

## Lossless compression codec (MDRLE)

In a typical desktop session, most of the imagery is simple graphics or text regions. Thinwire determines where these regions are and selects these areas for lossless encoding using the 2DRLE codec. At the Citrix Workspace app client side, these elements are decoded using the Citrix Workspace app-side 2DRLE decoder for session display.

In XenApp and XenDesktop 7.17, we added a higher compression ratio MDRLE codec that consumes less bandwidth in typical desktop sessions than the 2DRLE codec. This new codec does not impact server scalability.

Lower bandwidth usually means improved session interactivity (especially on shared or constrained links) and reduced costs. For example, the expected bandwidth consumption when using the MDRLE codec is approximately 10–15% less compared with XenApp and XenDesktop 7.15 LTSR for typical Office-like workloads.

Configuration isn't required for the MDRLE codec. If Citrix Workspace app supports MDRLE decoding, the VDA uses the VDA MDRLE encoding and the Citrix Workspace app MDRLE decoding. If Citrix Workspace app doesn't support MDRLE decoding, the VDA automatically falls back to 2DRLE encoding.

### MDRLE Requirements:

- Citrix Virtual Apps and Desktops minimum version 7 1808 VDAs
- XenApp and XenDesktop minimum version 7.17 VDAs
- Citrix Workspace app for Windows minimum version 1808
- Citrix Receiver for Windows minimum version 4.11

## Progressive Mode

Citrix Virtual Apps and Desktops 1808 introduced progressive mode and enabled it by default. In constrained network conditions (default: bandwidth < 2 Mbps, or latency > 200 ms), Thinwire increased the compression of text and static imagery to improve interactivity during screen activity. The heavily compressed text and images are then progressively sharpened, in a random block fashion, when screen activity stopped. While compressing and sharpening this way improves overall interactivity, it reduces cache efficiency and increases bandwidth usage.

As of Citrix Virtual Apps and Desktops 1906, progressive mode is disabled by default. We now use a different approach. The quality of still images is now based on network conditions and floats between a pre-defined minimum and maximum value for each **Visual quality** setting. Because there is no explicit sharpening step, Thinwire optimizes image delivery and maintains cache efficiency, while providing nearly all of the benefits of progressive mode.

## Changing progressive mode behavior

You can change the progressive mode state with the registry key. For information, see [Progressive mode](#) in the list of features managed through the registry.

### H.264 Build to lossless

**Build to lossless** is a special Thinwire configuration that optimizes graphics delivery for interactivity and final image quality. You can enable this setting by setting the **Visual quality** policy to **Build to lossless**.

Build to lossless compresses the screen using H.264 (or H.265) during screen activity and sharpens to pixel perfect (lossless) when activity stops. The H.264 (or H.265) image quality adapts to available resources to maintain the best possible frame rate. The sharpening step is performed gradually, giving an immediate response if the user begins screen activity shortly after sharpening starts. For example, selecting a model and rotating it.

H.264 **Build to lossless** offers all the advantages of full screen H.264 or H.265, including hardware acceleration, but with the added benefit of a final, guaranteed lossless screen. This is critical for 3D-type workloads that require a final pixel-perfect image. For example, manipulating medical imagery. Also, H.264 **Build to lossless** uses fewer resources than full screen H.264 4:4:4. As a result, using **Build to lossless** usually results in a higher frame rate than Visually lossless H.264 4:4:4.

#### Note:

In addition to the **Visual quality** policy, set the **Use video codec** policy to **Use when preferred** (default) or **For actively changing regions**. You can revert to non-H.264 Build to lossless by setting the **Use video codec** policy to **Do not use video codec**. This results in moving images being encoded with JPEG instead of H.264 (or H.265).

## Text-based session watermark

December 7, 2020

Text-based session watermarks help to deter and enable tracking data theft. This traceable information appears on the session desktop as a deterrent to those using photographs and screen captures to steal data. You can specify a watermark that is a layer of text, which displays over the entire session screen without changing the content of the original document. Text-based session watermarks require VDA support.

**Important:**

Text-based session watermarking is not a security feature. The solution does not prevent data theft completely, but it provides some level of deterrent and traceability. Although we do not guarantee complete information traceability when using this feature, we recommend that you combine this feature with other security solutions as applicable.

The session watermark is text and is applied to the session that is delivered to the user. The session watermark carries information for tracking data theft. The most important data is the identity of the logon user of the current session in which the screen image was taken. To trace the data leakage more effectively, include other information such as server or client internet protocol address and a connect time.

To adjust the user experience, use the [Session Watermark policy settings](#) to configure the placement and watermark appearance on the screen.

**Requirements:**

Virtual Delivery Agents:

Multi-session OS 7.17

Single-session OS 7.17

**Limitations:**

- Session watermarks are not supported in sessions where Local App Access, Windows media redirection, MediaStream, browser content redirection, and HTML5 video redirection are used. To use session watermark, ensure that these features are disabled.
- Session watermark is not supported and doesn't appear if the session is running in full-screen hardware accelerated modes (full-screen H.264 or H.265 encoding).
- If you set these HDX policies, watermark settings don't take effect and a watermark isn't displayed in the session display.

**Use hardware encoding for video codec to Enabled**

**Use video codec for compression to For the entire screen**

- If you set these HDX policies, the behavior is undetermined and the watermark might not display.

**Use hardware encoding for video codec to Enabled**

**Use video codec for compression to Use video codec when preferred**

To ensure the watermark displays, set **Use hardware encoding for video codec** to **Disabled**, or set **Use video codec for compression** to **For actively changing regions** or **Do not use video codec**.

- Session watermark supports only the Thinwire graphics mode.

- If you use Session Recording, the recorded session doesn't include the watermark.
- If you use Windows remote assistance, the watermark is not shown.
- If a user presses the **Print Screen** key to capture the screen, the screen captured at the VDA side doesn't include the watermarks. We recommend that you take measures to avoid the captured image being copied.

## Multimedia

April 7, 2020

The HDX technology stack supports the delivery of multimedia applications through two complementary approaches:

- Server-side rendering multimedia delivery
- Client-side rendering multimedia redirection

This strategy ensures that you can deliver a full range of multimedia formats, with a great user experience, while maximizing server scalability to reduce the cost-per-user.

With server-rendered multimedia delivery, audio and video content is decoded and rendered on the Citrix Virtual Apps and Desktops server by the application. The content is then compressed and delivered using ICA protocol to Citrix Workspace app on the user device. This method provides the highest rate of compatibility with various applications and media formats. Because video processing is compute-intensive, server-rendered multimedia delivery benefits greatly from the onboard hardware acceleration. For example, support for DirectX Video Acceleration (DXVA) offloads the CPU by performing H.264 decoding in separate hardware. Intel Quick Sync, AMD RapidFire, and NVIDIA NVENC technologies provide hardware-accelerated H.264 encoding.

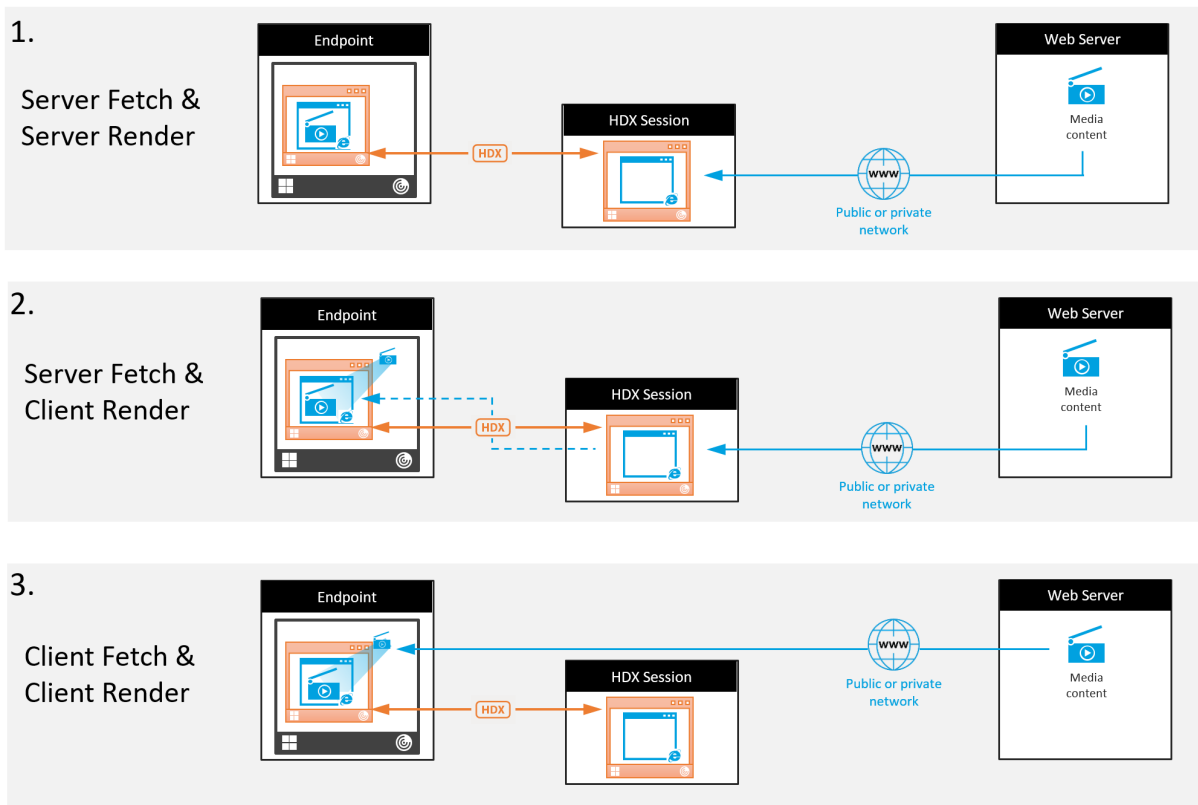
Because most servers do not offer any hardware acceleration for video compression, server scalability is negatively impacted if all video processing is done on the server CPU. You can maintain high server scalability, by redirecting many multimedia formats to the user device for local rendering.

- Windows Media redirection offloads the server for a wide variety of media formats typically associated with the Windows Media Player.
- HTML5 video has become popular, and Citrix introduced a redirection technology for this type of content. We recommend the browser content redirection for websites using HTML5, HLS, DASH, or WebRTC.
- You can apply the general content redirection technologies Host-to-client redirection and Local App Access to the multimedia content.

Putting these technologies together, if you don't configure redirection, HDX does Server-Side Rendering.

If you configure redirection, HDX uses either Server Fetch and Client Render or Client Fetch and Client Render. If those methods fail, HDX falls back to Server-Side Rendering as needed and is subject to the Fallback Prevention Policy.

### Example scenarios



#### Scenario 1. (Server Fetch and Server Rendering):

1. The server fetches the media file from its source, decodes, and then presents the content to an audio device or display device.
2. The server extracts the presented image or sound from the display device or audio device respectively.
3. The server optionally compresses it, and then transmits it to the client.

This approach incurs a high CPU cost, high bandwidth cost (if the extracted image/sound isn't compressed efficiently), and has low server scalability.

Thinwire and Audio virtual channels handle this approach. The advantage of this approach is that it reduces the hardware and software requirements for the clients. Using this approach the decoding happens on the server and it works for a wider variety of devices and formats.

#### Scenario 2. (Server Fetch and Client Render):

This approach relies on being able to intercept the media content before it is decoded and presented to the audio or display device. The compressed audio/video content is instead sent to the client where it is then decoded and presented locally. The advantage of this approach is that the are offloaded to the client devices, saving CPU cycles on the server.

However, it also introduces some additional hardware and software requirements for the client. The client must be able to decode each format that it might receive.

### **Scenario 3. (Client Fetching and Client Rendering):**

This approach relies on being able to intercept the media content URL before it's fetched from the source. The URL is sent to the client where the media content is fetched, decoded, and presented locally. This approach is conceptually simple. Its advantage is that it saves both CPU cycles on the server and bandwidth because the server sends only control commands. However, the media content is not always accessible to the clients.

### **Framework and platform:**

Single-session operating systems (Windows, Mac OS X, and Linux) provide multimedia frameworks that enable the faster development of multimedia applications. This table lists some of the more popular multimedia frameworks. Each framework divides media processing into several stages and uses a pipelined-based architecture.

---

Framework	Platform
DirectShow	Windows (98 and later)
Media Foundation	Windows (Vista and later)
Gstreamer	Linux
Quicktime	Mac OS X

---

### **Double hop support with media redirection technologies**

---

Audio redirection	No
Browser content redirection	No
HDX webcam redirection	Yes
HTML5 Video redirection	Yes
Windows Media redirection	Yes

---



## Audio features

July 8, 2021

You can configure and add the following Citrix policy settings to a policy that optimizes HDX audio features. For usage details plus relationships and dependencies with other policy settings, see [Audio policy settings](#) and [Bandwidth policy settings](#) and [Multi-stream connections policy settings](#).

### Important:

We recommend delivering audio using User Datagram Protocol (UDP) rather than TCP. Only Windows Virtual Delivery Agent (VDA) supports audio over UDP.

UDP audio encryption using DTLS is available only between Citrix Gateway and Citrix Workspace app. Therefore, sometimes it might be preferable to use TCP transport. TCP supports end-to-end TLS encryption from the VDA to Citrix Workspace app.

## Audio quality

In general, higher sound quality consumes more bandwidth and server CPU utilization by sending more audio data to user devices. Sound compression allows you to balance sound quality against overall session performance; use Citrix policy settings to configure the compression levels to apply to sound files.

By default, the **Audio quality policy** setting is set to High - high definition audio when TCP transport is used. The policy is set to Medium - optimized-for-speech when UDP transport (recommended) is used. The **High Definition audio** setting provides high fidelity stereo audio, but consumes more bandwidth than other quality settings. Do not use this audio quality for non-optimized voice chat or video chat applications (such as softphones). The reason being that it might introduce latency into the audio path that is not suitable for real-time communications. We recommend the optimized for speech policy setting for real-time audio, regardless of the selected transport protocol.

When the bandwidth is limited, for example satellite or dial-up connections, reducing audio quality to **Low** consumes the least possible bandwidth. In this situation, create separate policies for users on low-bandwidth connections so that users on high-bandwidth connections are not adversely impacted.

For setting details, see [Audio policy settings](#). Remember to enable Client audio settings on the user device.

Bandwidth guidelines for audio playback and recording:

- High quality (default)
  - Bitrate: ~100 kbps (min 75, max 175 kbps) for playback / ~70 kbps for microphone capture
  - Number of Channels: 2 (Stereo) for playback, 1 (mono) for microphone capture

- Frequency: 44100 Hz
- Bit-depth: 16-bit
- Medium quality (recommended for VoIP)
  - Bitrate: ~16 kbps (min 20, max 40 kbps) for playback, ~16 kbps for microphone capture
  - Number of Channels: 1 (Mono) for both playback and capture
  - Frequency: 16000 Hz (wideband)
  - Bit-depth: 16-bit
- Low quality
  - Bitrate: ~ 11 kbps (min 10; max 25 kbps) for playback, ~11 kbps for microphone capture
  - Number of Channels: 1 (Mono) for both playback and capture
  - Frequency: 8000 Hz (narrowband)
  - Bit-depth: 16-bit

### Client audio redirection

To allow users to receive audio from an application on a server through speakers or other sound devices on the user device, leave the **Client audio redirection** setting at **Allowed**. This is the default.

Client audio mapping puts extra load on the servers and the network. However, prohibiting client audio redirection disables all HDX audio functionality.

For setting details, see [Audio policy settings](#). Remember to enable client audio settings on the user device.

### Client microphone redirection

To allow users to record audio using input devices such as microphones on the user device, leave the **Client microphone redirection** setting at its default (Allowed).

For security, user devices alert their users when servers they don't trust try to access microphones. Users can choose to accept or reject access before using the microphone. Users can disable this alert on Citrix Workspace app.

For setting details, see [Audio policy settings](#). Remember to enable Client audio settings on the user device.

### Audio Plug N Play

The Audio Plug N Play policy setting allows or prevents the use of multiple audio devices to record and play sound. This setting is **Enabled** by default. Audio Plug N Play enables audio devices to be recognized. The devices are recognized even if they are not plugged in until after the user session has started.

This setting applies only to Windows Multi-session OS machines.

For setting details, see [Audio policy settings](#).

### **Audio redirection bandwidth limit and audio redirection bandwidth limit percent**

The Audio redirection bandwidth limit policy setting specifies the maximum bandwidth (in kilobits per second) for a playing and recording audio in a session.

The Audio redirection bandwidth limit percent setting specifies the maximum bandwidth for audio redirection as a percentage of the total available bandwidth.

By default, zero (no maximum) is specified for both settings. If both settings are configured, the one with the lowest bandwidth limit is used.

For setting details, see [Bandwidth policy settings](#). Remember to enable Client audio settings on the user device.

### **Audio over UDP Real-time Transport and Audio UDP port range**

By default, Audio over User Datagram Protocol (UDP) Real-time Transport is allowed (when selected at the time of installation). It opens up a UDP port on the server for connections that use Audio over UDP Real-time Transport. If there is network congestion or packet loss, we recommend configuring UDP/RTP for audio to ensure the best possible user experience. For any real time audio such as soft-phone applications, UDP audio is preferred to EDT. UDP allows for packet loss without retransmission, ensuring that no latency is added on connections with high packet loss.

#### **Important:**

When Citrix Gateway is not in the path, audio data transmitted with UDP is not encrypted. If Citrix Gateway is configured to access Citrix Virtual Apps and Desktops resources, then audio traffic between the endpoint device and Citrix Gateway is secured using DTLS protocol.

The Audio UDP port range specifies the range of port numbers that the Windows VDA uses to exchange audio packet data with the user device.

By default, the range is 16500 through 16509.

For setting details about Audio over UDP Real-time Transport, see [Audio policy settings](#). For details about Audio UDP port range, see [Multi-stream connections policy settings](#). Remember to enable Client audio settings on the user device.

Audio over UDP requires the Windows VDA. For supported policies on the Linux VDA, see [Policy support list](#).

## Audio setting policies for user devices

1. Load the group policy templates by following [Configuring the Group Policy Object administrative template](#).
2. In the Group Policy Editor, expand **Administrative Templates > Citrix Components > Citrix Workspace > User Experience**.
3. For **Client audio settings**, select **Not Configured**, **Enabled**, or **Disabled**.
  - **Not Configured**. By default, Audio Redirection is enabled using high quality audio or the previously configured custom audio settings.
  - **Enabled**. Enables audio redirection using the selected options.
  - **Disabled**. Disables audio redirection.
4. If you select **Enabled**, choose a sound quality. For UDP audio, use **Medium** (default).
5. For UDP audio only, select **Enable Real-Time Transport** and then set the range of incoming ports to open in the local Windows firewall.
6. To use UDP Audio with Citrix Gateway, select **Allow Real-Time Transport Through gateway**. Configure Citrix Gateway with DTLS. For more information, see [this article](#).

As an Administrator, if you do not have control on endpoint devices to make these changes, use the default.ica attributes from StoreFront to enable UDP Audio. For example, for bring your own devices or home computers.

1. On the StoreFront machine, open C:\inetpub\wwwroot\Citrix\\App\_Data\default.ica with an editor such as notepad.
2. Make the following entries under the [Application] section.

; This text enables Real-Time Transport

```
EnableRtpAudio=true
```

; This text allows Real-Time Transport Through gateway

```
EnableUDPThroughGateway=true
```

; This text sets audio quality to Medium

```
AudioBandwidthLimit=1
```

; UDP Port range

```
RtpAudioLowestPort=16500
```

```
RtpAudioHighestPort=16509
```

If you enable User Datagram Protocol (UDP) audio by editing default.ica, then UDP audio is enabled for all users who are using that store.

## Avoid echo during multimedia conferences

Users in audio or video conferences might hear an echo. Echoes usually occur when speakers and microphones are too close to each other. For that reason, we recommend the use of headsets for audio and video conferences.

HDX provides an echo cancellation option (enabled by default) that minimizes any echo. The effectiveness of echo cancellation is sensitive to the distance between the speakers and the microphone. Ensure that the devices aren't too close or too far away from each other.

You can change a registry setting to disable echo cancellation. For information, see [Avoid echo during multimedia conferences](#) in the list of features managed through the registry.

## Softphones

A softphone is software acting as a phone interface. You use a softphone to make calls over the internet from a computer or other smart device. By using a softphone, you can dial phone numbers and carry out other phone-related functions using a screen.

Citrix Virtual Apps and Desktops support several alternatives for delivering softphones.

- **Control mode.** The hosted softphone controls a physical telephone set. In this mode, no audio traffic goes through the Citrix Virtual Apps and Desktops server.
- **HDX RealTime optimized softphone support (recommended).** The media engine runs on user device, and Voice over Internet Protocol traffic flows peer-to-peer. For examples, see:
  - [HDX Optimization for Microsoft Teams](#)
  - [HDX RealTime Optimization Pack](#), which optimizes the delivery of Microsoft Skype for Business
  - [Cisco Jabber Softphone for VDI](#) (formerly known as VXME)
  - [Cisco Webex Meetings for VDI](#)
  - [Avaya VDI Equinox](#) (formerly known as [VDI Communicator](#))
  - [Zoom VDI Plugin](#)
  - [Genesys PureEngage Cloud](#)
  - [Nuance Dragon PowerMic dictation device](#)
- **Local App Access.** A Citrix Virtual Apps and Desktops feature that allows an application such as a softphone to run locally on the Windows user device yet appear seamlessly integrated with their virtual/published desktop. This feature offloads all audio processing to the user device. For more information, see [Local App Access and URL redirection](#).
- **HDX RealTime generic softphone support.** Voice over Internet Protocol-over-ICA.

### **Generic softphone support**

Generic softphone support, enables you to host an unmodified softphone on XenApp or XenDesktop in the data center. The audio traffic goes over the Citrix ICA protocol (preferably using UDP/RTP) to

the user device running the Citrix Workspace app.

Generic softphone support is a feature of HDX RealTime. This approach to softphone delivery is especially useful when:

- An optimized solution for delivering the softphone is not available and the user is not on a Windows device where Local App Access can be used.
- The media engine that is needed for optimized delivery of the softphone isn't installed on the user device or isn't available for the operating system version running on the user device. In this scenario, Generic HDX RealTime provides a valuable fallback solution.

There are two softphone delivery considerations using Citrix Virtual Apps and Desktops:

- How the softphone application is delivered to the virtual/published desktop.
- How the audio is delivered to and from the user headset, microphone, and speakers, or USB telephone set.

Citrix Virtual Apps and Desktops include numerous technologies to support generic softphone delivery:

- Optimized-for-Speech codec for fast encode of the real-time audio and bandwidth efficiency.
- Low latency audio stack.
- Server-side jitter buffer to smooth out the audio when the network latency fluctuates.
- Packet tagging (DSCP and WMM) for Quality of Service.
  - DSCP tagging for RTP packets (Layer 3)
  - WMM tagging for Wi-Fi

The Citrix Workspace app versions for Windows, Linux, Chrome, and Mac also are Voice over Internet Protocol capable. Citrix Workspace app for Windows offers these features:

- Client-side jitter buffer - Ensures smooth audio even when the network latency fluctuates.
- Echo cancellation - Allows for greater variation in the distance between microphone and speakers for workers who do not use a headset.
- Audio plug-n-play - Audio devices do not need to be plugged in before starting a session. They can be plugged in at any time.
- Audio device routing - Users can direct ringtone to speakers but the voice path to their headset.
- Multi-stream ICA - Enables flexible Quality of Service-based routing over the network.
- ICA supports four TCP and two UDP streams. One of the UDP streams supports the real-time audio over RTP.

For a summary of Citrix Workspace app capabilities, see [Citrix Receiver Feature Matrix](#).

### **System configuration recommendations**

#### *Client Hardware and Software:*

For optimal audio quality, we recommend the latest version of Citrix Workspace app and a good quality headset that has acoustic echo cancellation (AEC). Citrix Workspace app versions for Windows,

Linux, and Mac support Voice over Internet Protocol. Also, Dell Wyse offers Voice over Internet Protocol support for ThinOS (WTOS).

*CPU Considerations:*

Monitor CPU usage on the VDA to determine if it is necessary to assign two virtual CPUs to each virtual machine. Real-time voice and video are data intensive. Configuring two virtual CPUs reduces the thread switching latency. Therefore, we recommend that you configure two vCPUs in a Citrix Virtual Desktops VDI environment.

Having two virtual CPUs does not necessarily mean doubling the number of physical CPUs, because physical CPUs can be shared across sessions.

Citrix Gateway Protocol (CGP), which is used for the Session Reliability feature, also increases CPU consumption. On high-quality network connections, you can disable this feature to reduce CPU consumption on the VDA. Neither of the preceding steps might be necessary on a powerful server.

*UDP Audio:*

Audio over UDP provides excellent tolerance of network congestion and packet loss. We recommend it instead of TCP when available.

*LAN/WAN configuration:*

Proper configuration of the network is critical for good real-time audio quality. Typically, you must configure virtual LANs (VLANs) because excessive broadcast packets can introduce jitter. IPv6-enabled devices might generate many broadcast packets. If IPv6 support is not needed, you can disable IPv6 on those devices. Configure to support Quality of Service.

*Settings for use WAN connections:*

You can use voice chat over LAN and WAN connections. On a WAN connection, audio quality depends on the latency, packet loss, and jitter on the connection. If delivering softphones to users on a WAN connection, we recommend using the NetScaler SD-WAN between the data center and the remote office. Doing so maintains a high Quality of Service. NetScaler SD-WAN supports Multi-Stream ICA, including UDP. Also, for a single TCP stream, it's possible to distinguish the priorities of various ICA virtual channels to ensure that high priority real-time audio data receives preferential treatment.

Use Director or the [HDX Monitor](#) to validate your HDX configuration.

*Remote user connections:*

Citrix Gateway supports DTLS to deliver UDP/RTP traffic natively (without encapsulation in TCP). Open firewalls bidirectionally for UDP traffic over Port 443.

*Codec selection and bandwidth consumption:*

Between the user device and the VDA in the data center, we recommend using the **Optimized-for-Speech** codec setting, also known as Medium Quality audio. Between the VDA platform and the IP-PBX, the softphone uses whatever codec is configured or negotiated. For example:

- G711 provides good voice quality but has a bandwidth requirement of from 80 kilobits per second through 100 kilobits per second per call (depending on Network Layer2 overheads).
- G729 provides good voice quality and has a low bandwidth requirement of from 30 kilobits per second through 40 kilobits per second per call (depending on Network Layer 2 overheads).

### ***Delivering softphone applications to the virtual desktop***

There are two methods by which you can deliver a softphone to the XenDesktop virtual desktop:

- The application can be installed in the virtual desktop image.
- The application can be streamed to the virtual desktop using Microsoft App-V. This approach has manageability advantages because the virtual desktop image is kept uncluttered. After being streamed to the virtual desktop, the application runs in that environment as if it was installed in the usual manner. Not all applications are compatible with App-V.

### ***Delivering audio to and from the user device***

Generic HDX RealTime supports two methods of delivering audio to and from the user device:

- **Citrix Audio Virtual Channel.** We generally recommend the Citrix Audio Virtual Channel because it's designed specifically for audio transport.
- **Generic USB Redirection.** Supports audio devices having buttons or a display (or both), human interface device (HID), if the user device is on a LAN or LAN-like connection back to the Citrix Virtual Apps and Desktops server.

### ***Citrix audio virtual channel***

The bidirectional Citrix Audio Virtual Channel (CTXCAM) enables audio to be delivered efficiently over the network. Generic HDX RealTime takes the audio from the user headset or microphone and compresses it. Then, it sends it over ICA to the softphone application on the virtual desktop. Likewise, the audio output of the softphone is compressed and sent in the other direction to the user headset or speakers. This compression is independent of the compression used by the softphone itself (such as G.729 or G.711). It is done using the Optimized-for-Speech codec (Medium Quality). Its characteristics are ideal for Voice over Internet Protocol. It features quick encode time, and it consumes only approximately 56 Kilobits per second of network bandwidth (28 Kbps in each direction), peak. This codec must be explicitly selected in the service's Manage console because it is not the default audio codec. The default is the HD Audio codec (High Quality). This codec is excellent for high fidelity stereo soundtracks but is slower to encode compared to the Optimized-for-Speech codec.

### ***Generic USB Redirection***

Citrix Generic USB Redirection technology (CTXGUSB virtual channel) provides a generic means of remoting USB devices, including composite devices (audio plus HID) and isochronous USB devices. This approach is limited to LAN-connected users. This reason being that the USB protocol tends to be sensitive to network latency and requires considerable network bandwidth. Isochronous USB redirection works well when using some softphones. This redirection provides excellent voice quality and low la-



tency. However, Citrix Audio Virtual Channel is preferred because it is optimized for audio traffic. The primary exception is when you're using an audio device with buttons. For example, a USB telephone attached to the user device that is LAN-connected to the data center. In this case, Generic USB Redirection supports buttons on the phone set or headset that control features by sending a signal back to the softphone. There isn't an issue with buttons that work locally on the device.

## Limitation

After you install an audio device on your client, enable the audio redirection, and start an RDS session, the audio files might not play audio. As a workaround, add the registry key on the RDS machine, and then restart the machine. For information, see [Audio limitation](#) in the list of features managed through the registry.

## Browser content redirection

July 8, 2021


Browser content redirection prevents the rendering of webpages in the allow list on the VDA side. This feature uses Citrix Workspace app to instantiate a corresponding rendering engine on the client side, which fetches the HTTP and HTTPS content from the URL.

### Note:

You can specify that webpages be redirected to the VDA side (and not redirected on the client side) by using a block list.

This overlay web layout engine runs on the endpoint device instead of on the VDA and uses the endpoint CPU, GPU, RAM, and Network.

Only the browser viewport is redirected. The viewport is the rectangular area in your browser where content displays. The viewport doesn't include things like the Address Bar, Favorites Toolbar, Status Bar. Those items are in the user interface, which are still running on the browser in the VDA.




1. Configure a policy in the Manage > Full Configuration interface that specifies the Access Control List containing the URLs for redirection from the allow or block lists. For the browser on the VDA to detect that the URL that the user is navigating to matches the allow list or does not match a block list, a browser extension performs the comparison. The browser extension (BHO) for Internet Explorer 11 is included in the installation media and is installed automatically. For Chrome, the browser extension is available in the Chrome Web Store, and you can deploy it using the Group Policies and ADMX files. Chrome extensions are installed on a per-user basis. Updating a golden image to add or remove an extension is not required.
2. If a match is found in the allow list (for example <https://www.mycompany.com/>), and there is no match to a URL in the block list (for example <https://www.mycompany.com/engineering>), a virtual channel (CTXCSB) instructs Citrix Workspace app that a redirection is required and relays the URL. Citrix Workspace app then instantiates a local rendering engine and displays the website.
3. Citrix Workspace app then blends back the website into the virtual desktop browser content area seamlessly.

The color of the logo specifies the status of the Chrome extension. It is one of these three colors:

- Green: Active and connected.
- Gray: Not active/idle on the current tab.
- Red: Broken/Not working.

You can debug logging by using **Options** in the extensions menu.



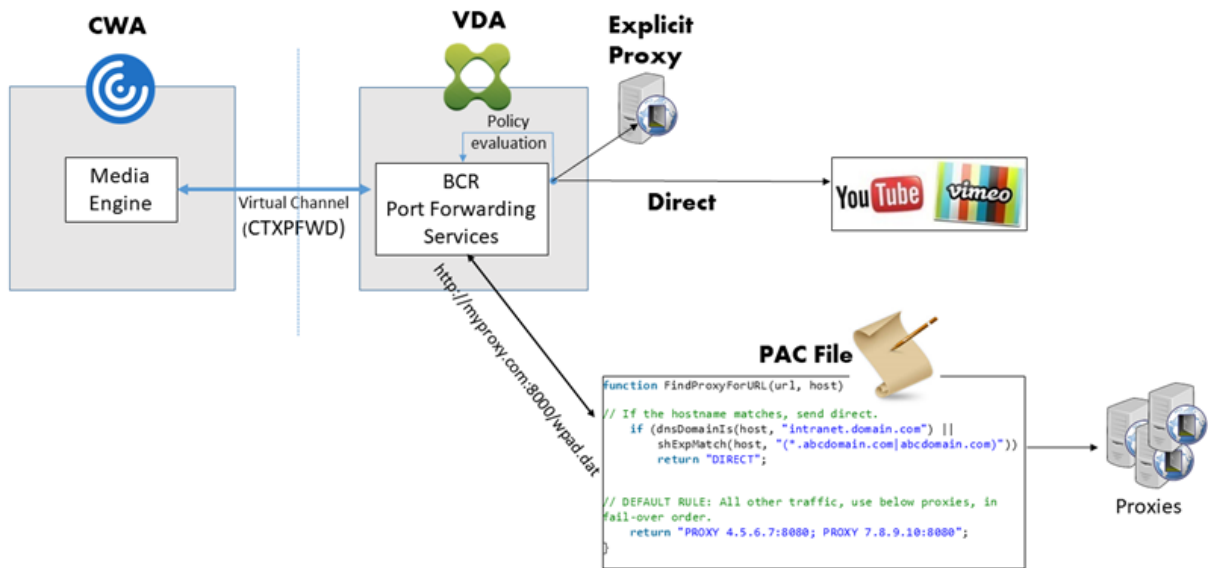
Here are scenarios of how Citrix Workspace app fetches content:

- **Server fetch and server render:** There is no redirection because you didn't add the site to the allow list or the redirection failed. We fall back to rendering the webpage on the VDA and use Thinwire to remote the graphics. Use policies to control the fallback behavior. High CPU, RAM, and bandwidth consumption on the VDA.
- **Server fetch and client render:** Citrix Workspace app contacts and fetches content from the web server through the VDA using a virtual channel (CTXPFW). This option is useful when the client doesn't have internet access (for example, thin clients). Low CPU and RAM consumption on the VDA, but bandwidth is consumed on the ICA virtual channel.

There are three modes of operation for this scenario. The term proxy refers to a proxy device that the VDA accesses to gain Internet access.

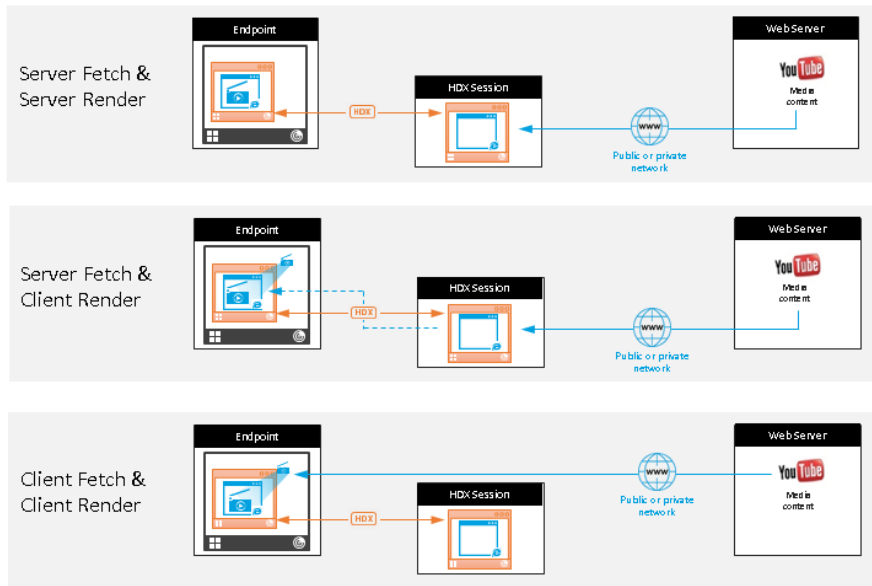
Which policy option to choose:

- Explicit Proxy - If you have a single explicit proxy in your Datacenter.
- Direct or Transparent - If you do not have proxies, or if you use transparent proxies.
- PAC files - If you rely on PAC files so browsers in the VDA can automatically choose the appropriate proxy server for fetching a specified URL.



- **Client fetch and client render:** Because Citrix Workspace app contacts the web server directly, it requires internet access. This scenario offloads all the network, CPU, and RAM usage from your XenApp and XenDesktop Site.

## Redirection scenarios



### Benefits:

- Better end user experience (Adaptive Bit Rate (ABR))
- Reduced VDA resource usage (CPU/RAM/IO)
- Reduced bandwidth consumption

### Fallback mechanism:

There might be times when client redirection fails. For example, if the client machine does not have direct internet access, an error response might go back to the VDA. In such cases, the browser on the VDA can then reload and render the page on the server.

You can suppress server rendering of video elements by using the existing **Windows media fallback prevention** policy. Set this policy to **Play all content only on client** or **Play only client-accessible content on client**. These settings block video elements from playing on the server if there are failures in client redirection. This policy takes effect only when you enable browser content redirection and the **Access Control List** policy contains the URL that falls back. The URL can't be in the block list policy.

### System requirements:

Windows endpoints:

- Windows 7, 8.x, or 10
- Citrix Workspace app 1809 for Windows or later

#### Note:

Citrix Workspace app 1912 LTSR for Windows as well as all cumulative updates to the Citrix Workspace app 1912 LTSR do not support browser content redirection.

Linux endpoints:

- Citrix Workspace app 1808 for Linux or later
- Citrix Receiver for Linux 13.9 or later
- Thin client terminals must include WebKitGTK+

Citrix Virtual Apps and Desktops 7 1808 and XenApp and XenDesktop 7.15 CU5, 7.18, 7.17, 7.16:

- VDA operating system: Windows 10 (minimum version 1607), Windows Server 2012 R2, Windows Server 2016
- Browser on the VDA:
  - Google Chrome v66 or higher (Chrome requires Citrix Workspace app 1809 for Windows on the user endpoint, Citrix Virtual Apps and Desktops 7 1808 VDA, and the browser content redirection extension)
  - Internet Explorer 11 and configure these options:
    - \* Clear **Enhanced Protected Mode** under: **Internet Options > Advanced > Security**
    - \* Check **Enable third-party browser extensions** under: **Internet Options > Advanced > Browsing**

### Troubleshooting

For troubleshooting information, see the Knowledge Center article <https://support.citrix.com/article/CTX230052>

## Browser content redirection Chrome extension

To use browser content redirection with Chrome, add the browser content redirection extension from the Chrome Web Store. Click **Add to Chrome** in the Citrix Virtual Apps and Desktops environment.


The extension is **not** required on the user's client machine – only in the VDA.

### System requirements

- Chrome v66 or higher
- Browser content redirection extension
- Citrix Virtual Apps and Desktops 7 1808 or higher
- Citrix Workspace app 1809 for Windows or higher

#### Note:

Citrix Workspace app 1912 LTSR for Windows as well as all cumulative updates to the Citrix Workspace app 1912 LTSR do not support browser content redirection.

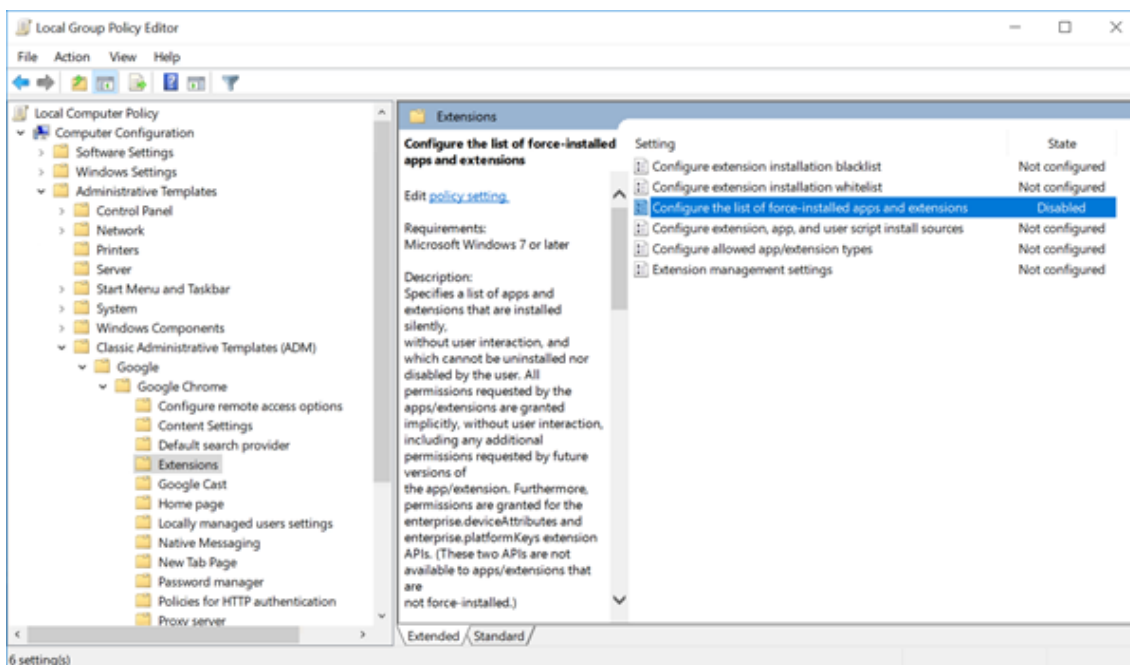


This method works for individual users. To deploy the extension to a large group of users in your organization, deploy the extension using Group Policy.

### Deploy the extension using Group Policy


1. Import the Google Chrome ADMX files into your environment. For information about downloading policy templates and installing and configuring the templates into your Group Policy Editor, see [Set Chrome Browser policies on managed PCs](#).

2. Open your Group Policy Management console and go to **User Configuration \ Administrative Templates\Classic Administrative Templates (ADM) \ Google\ Google Chrome \ Extensions**. Enable the **Configure the list of force-installed apps and extensions** setting.



3. Click **Show** and type the following string, which corresponds to the extension ID. Update the URL for the browser content redirection extension.

hdppkjifljbdpckfajcmlblbchhledln; <https://clients2.google.com/service/update2/crx>



4. Apply the setting and after a **gpupdate** refresh, the user automatically receives the extension. If you launch the Chrome browser in the user's session, the extension is already applied and they cannot remove it.

Any updates to the extension are automatically installed on the users' machines through the update URL that you specified in the setting.

If the **Configure the list of force-installed apps and extensions** setting is set to **Disabled**, the extension is automatically removed from Chrome for all users.


### Browser content redirection Edge Chromium extension

To install the browser content redirection extension in Edge, make sure you have version **83.0.478.37** or higher of the Edge browser installed.

1. Click the **Extensions** option in the menu and turn on **Allow extensions from other stores**.



2. Click the **Chrome Web Store** link and the extension appears at the bar on the top right.  
For more info on Microsoft Edge extensions, see [Extensions](#).



## Browser content redirection and DPI

When using browser content redirection with the DPI (scaling) set to anything over 100% on the user's machine, the redirected browser content screen displays incorrectly. To avoid this issue, do not set the DPI when using browser content redirection. Another way to avoid the issue is by disabling browser content redirection GPU acceleration for Chrome by creating the registry key on the user's machine. For information, see [Browser content redirection and DPI](#) in the list of features managed through the registry.

## User-agent request header

The user-agent header helps identify HTTP requests sent from browser content redirection. This setting can be useful when you configure proxy and firewall rules. For example, if the server blocks the requests sent from browser content redirection, you can create a rule that contains the user-agent header to bypass certain requirements.

Only Windows devices support the user-agent request header.

By default, the user-agent request header string is disabled. To enable the user-agent header for client-rendered content, use the Registry editor. For information, see [User-agent request header](#) in the list of features managed through the registry.

## HDX video conferencing and webcam video compression

June 18, 2021

### Warning:

Editing the registry incorrectly can cause serious problems that might require you to reinstall

your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

Webcams can be used by applications running within the virtual session by using HDX webcam video compression or HDX plug-n-play generic USB redirection. Use **Citrix Workspace app > Preferences > Devices** to switch between modes. Citrix recommends you always use HDX webcam video compression if possible. HDX generic USB redirection is recommended only when there are application compatibility issues with HDX video compression or when you require advanced native functionalities of the webcam. For better performance, Citrix recommends the Virtual Delivery Agent to have at least two virtual CPUs.

To prevent users from switching from HDX webcam video compression, disable USB device redirection by using the policy settings under **ICA policy settings > USB Devices policy** settings. Citrix Workspace app users can override the default behavior by choosing the Desktop Viewer Mic & Webcam setting **Don't use my microphone or webcam**.


### **HDX webcam video compression**

HDX webcam video compression is also called **Optimized** webcam mode. This type of webcam video compression sends the H.264 video directly to the video conferencing application running in the virtual session. To optimize VDA resources, HDX webcam compression doesn't encode, transcode, and decode webcam video. This feature is enabled by default.

To disable direct video streaming from the server to the video conferencing app, set the registry key to 0 on the VDA. For information, see [Webcam video compression](#) in the list of features managed through the registry.

If you disable the default functionality for streaming video resources, HDX webcam video compression uses the multimedia framework technology that is part of the client operating system to intercept video from capture devices, transcode, and compress it. Manufacturers of capture devices supply the drivers that plug into the OS kernel streaming architecture.

The client handles communication with the webcam. The client then sends the video only to the server that can display it properly. The server doesn't deal directly with the webcam, but its integration gives you the same experience in your desktop. Workspace app compresses the video to save bandwidth and provide better resiliency on WAN scenarios.



HDX webcam video compression requires that the following policy settings be enabled (all are enabled by default).

- Multimedia conferencing
- Windows Media Redirection

If a webcam supports hardware encoding, HDX video compression uses the hardware encoding by default. Hardware encoding might consume more bandwidth than software encoding. To force software compression, edit the registry key on the client. For information, see [Webcam software compression](#) in the list of features managed through the registry.

### HDX webcam video compression requirements

HDX webcam video compression supports the following versions of Citrix Workspace app:

Platform	Processor
Citrix Workspace app for Windows	Citrix Workspace app for Windows supports webcam video compression for 32-bit and 64-bit apps on XenApp and XenDesktop 7.17 and later. On earlier versions, Citrix Workspace app for Windows supports only 32-bit apps.
Citrix Workspace app for Mac	Citrix Workspace app for Mac 2006 or later supports webcam video compression for 64-bit apps on XenApp and XenDesktop 7.17 and later. On earlier versions, Citrix Workspace app for Mac supports only 32-bit apps.
Citrix Workspace app for Linux	Citrix Workspace app for Linux supports only 32-bit apps on the virtual desktop.

---

Platform	Processor
Citrix Workspace app for Chrome	Because some ARM Chromebooks don't support H.264 encoding, only 32-bit apps can use the optimized HDX webcam video compression.

---

Media foundation-based video applications support HDX webcam video compression on Windows 8.x or higher and Windows Server 2012 R2 and higher. For more information, see Knowledge Center article [CTX132764](#).

Other user device requirements:

- Appropriate hardware to produce sound.
- DirectShow-compatible webcam (use the webcam default settings). Webcams that are hardware encoding capable reduce client-side CPU usage.
- For HDX webcam video compression, install webcam drivers on the client, obtained from the camera manufacturer, if possible. Installation of the device drivers isn't required on the server.

Different webcams offer different frame rates and have different levels of brightness and contrast. Adjusting the contrast of the webcam can reduce upstream traffic significantly. Citrix uses the following webcams for initial feature validation:

- Microsoft LifeCam VX models (2000, 3000, 5000, 7000)
- Creative Live! Cam Optia Pro
- Logitech QuickCam Messenger
- Logitech C600, C920
- HP Deluxe Webcam

To adjust the preferred video frame rate, edit the registry key on the client. For information, see [Webcam video compression frame rate](#) in the list of features managed through the registry.

## High-definition webcam streaming

The video conferencing application on the server selects the webcam format and resolution based on the supported format types. When a session starts, the client sends the webcam information to the server. Choose a webcam from the application. When the webcam and the video conferencing application support high-definition rendering, the application uses high-definition resolution. We support webcam resolutions up to 1920x1080.

This feature requires the Citrix Workspace app for Windows, minimum version 1808 or Citrix Receiver for Windows, minimum version 4.10.

You can use a registry key to disable and enable the feature. For information, see [High-definition webcam streaming](#) in the list of features managed through the registry.

If the media type negotiation fails, HDX falls back to the default resolution of 352x288 CIF. You can use registry keys on the client to configure the default resolution. Ensure that the camera supports the specified resolution. For information, see [High-definition webcam resolution](#) in the list of features managed through the registry.

HDX webcam video compression uses significantly less bandwidth compared to plug-n-play generic USB redirection and works well over WAN connections. To adjust the bandwidth, set the registry key on the client. For information, see [High-definition webcam bandwidth](#) in the list of features managed through the registry.

Enter a value in bits per second. If you don't specify the bandwidth, the video conferencing applications use 350000 bps by default.

## **HDX plug-n-play generic USB redirection**

HDX plug-n-play generic USB redirection (isochronous) is also called **Generic** webcam mode. The benefit of HDX plug-n-play generic USB redirection is that you don't have to install drivers on your thin client/endpoint. The USB stack is virtualized such that anything you plug into the local client is sent to the remote VM. The remote desktop acts as if you plugged it in natively. The Windows desktop handles all the interaction with the hardware and runs through the plug-n-play logic to find the correct drivers. Most webcams work if the drivers exist on the server and can work over ICA. Generic webcam mode uses significantly more bandwidth (many Megabits per second) because you are sending uncompressed video down with USB protocol over the network.

## **HTML5 multimedia redirection**

August 21, 2020

HTML5 multimedia redirection extends the multimedia redirection features of HDX MediaStream to include HTML5 audio and video. Because of growth in online distribution of multimedia content, especially to mobile devices, the browser industry has developed more efficient ways to present audio and video.

Flash has been the standard, but it requires a plug-in, doesn't work on all devices, and has higher battery usage in mobile devices. Companies like YouTube, Netflix.com, and newer browsers versions of Mozilla, Google, and Microsoft are moving to HTML5 making it the new standard.

HTML5-based multimedia has many advantages over proprietary plug-ins, including:

- Company-independent standards (W3C)

- Simplified digital rights management (DRM) workflow
- Better performance without the security issues raised by plug-ins

## HTTP progressive downloads

HTTP progressive download is an HTTP-based pseudo-streaming method that supports HTML5. In a progressive download, the browser plays back a single file (encoded at a single quality) while it is being downloaded from an HTTP web server. The video is stored on the drive as it's received and is played from the drive. If you rewatch the video, the browser can load the video from cache.

For an example of a progressive download, see the [HTML5 video redirection test page](#). To inspect the video elements in the webpage and find the sources (mp4 container format) in HTML5 video tags, use the developer tools in your browser:

## Comparing HTML5 and Flash

Feature	HTML5	Flash
Requires a proprietary player	No	Yes
Runs on mobile devices	Yes	Some
Running speed on different platforms	High	Slow
Supported by iOS	Yes	No
Resource usage	Less	More
Load faster	Yes	No

## Requirements

We support only redirection for progressive downloads in mp4 format. We don't support WebM and Adaptive bitrate streaming technologies like DASH/HLS.

We support the following, and use policies to control them. For more information, see [Multimedia policy settings](#).

- Server side render
- Server fetch client render
- Client side fetching and rendering

Minimum versions of Citrix Workspace app and Citrix Receiver:

- Citrix Workspace app 1808 for Windows

- Citrix Receiver for Windows 4.5
- Citrix Workspace app 1808 for Linux
- Citrix Receiver for Linux 13.5

Minimum VDA browser version	Windows OS version/build/SP
Internet Explorer 11.0	Windows 10 x86 (1607 RS1) and x64 (1607 RS1); Windows 7 x86 and x64; Windows Server 2016 RTM 14393 (1607); Windows Server 2012 R2
Firefox 47 Manually add the certificates to the Firefox certificate store or configure Firefox to search for certificates from a Windows trusted certificate store. For more information, see <a href="https://wiki.mozilla.org/CA:AddRootToFirefox">https://wiki.mozilla.org/CA:AddRootToFirefox</a>	Windows 10 x86 (1607 RS1) and x64 (1607 RS1); Windows 7 x86 and x64; Windows Server 2016 RTM 14393 (1607); Windows Server 2012 R2
Chrome 51	Windows 10 x86 (1607 RS1) and x64 (1607 RS1); Windows 7 x86 and x64; Windows Server 2016 RTM 14393 (1607); Windows Server 2012 R2

## Components of the HTML5 video redirection solution

- **HdxVideo.js** - JavaScript hook intercepting video commands on the website. HdxVideo.js communicates with WebSocketService using Secure WebSockets (SSL/TLS).
- **WebSocket SSL Certificates**
  - For the CA (root): **Citrix XenApp/XenDesktop HDX In-Product CA** (C = US; S = Florida; L = Fort Lauderdale; O = Citrix Systems, Inc.; OU = XenApp/XenDesktop Engineering; CN = Citrix XenApp/XenDesktop HDX In-Product CA)  
Location: Certificates (Local Computer) > Trusted Root Certification Authorities > Certificates.
  - For the end-entity (leaf): **Citrix XenApp/XenDesktop HDX Service** (C = US; S = Florida; L = Fort Lauderdale; O = Citrix Systems, Inc.; OU = XenApp/XenDesktop Engineering; CN = Citrix XenApp/XenDesktop HDX Service)  
Location: Certificates (Local Computer) > Personal > Certificates.
- **WebSocketService.exe** - Runs on the local system and performs SSL termination and user session mapping. TLS Secure WebSocket listening on 127.0.0.1 port 9001.
- **WebSocketAgent.exe** - Runs on the user session and renders the video as instructed from WebSocketService commands.

## How do I enable HTML5 video redirection?


In this release, this feature is available for controlled webpages only. It requires the addition of the HdxVideo.js JavaScript (included in the Citrix Virtual Apps and Desktops Installation media) to the webpages where the HTML5 multimedia content is available. For example, videos on an internal training site.

Websites like youtube.com, which are based on Adaptive Bitrate technologies (for example, HTTP Live Streaming (HLS) and Dynamic Adaptive Streaming over HTTP (DASH)), are not supported.

For more information, see [Multimedia policy settings](#).

## Troubleshooting Tips

Errors might occur when the webpage tries to run HdxVideo.js. If the JavaScript fails to load, the HTML5 redirection mechanism fails. Ensure that there are no errors related to HdxVideo.js by inspecting the console in the developers tool windows of your browser. For example:



## Optimization for Microsoft Teams

September 2, 2021

### Important:

Optimization for Microsoft Teams requires a minimum of Microsoft Teams version 1.2.00.31357.

Citrix delivers optimization for desktop-based Microsoft Teams using Citrix Virtual Apps and Desktops and Citrix Workspace app. By default, we bundle all the necessary components into the Citrix Workspace app and the Virtual Delivery Agent (VDA).

Our optimization for Microsoft Teams contains VDA-side HDX services and an API to interface with the Microsoft Teams hosted app to receive commands. These components open a control virtual channel (CTXMTOP) to the Citrix Workspace app-side media engine. The endpoint decodes and renders the multimedia locally, moving the Citrix Workspace app window back into the hosted Microsoft Teams app.


Authentication and signaling occur natively on the Microsoft Teams-hosted app, just like the other Microsoft Teams services (for example chat or collaboration). Audio/video redirection doesn't affect them.



CTXMTOP is a command and control virtual channel. That means that media isn't exchanged between the Citrix Workspace app and the VDA.

Only client-fetch/client-render is available.

This video demo gives you an idea of how Microsoft Teams works in a Citrix virtual environment.



## Microsoft Teams installation

### Note:

We recommend installing the VDA before installing Microsoft Teams in the golden image. This installation order is needed for the **ALLUSER=1** flag to take effect. If you installed Microsoft Teams in the virtual machine before installing the VDA, uninstall and reinstall Microsoft Teams. If you're using App Layering, see [For App Layering](#) for more details.

We recommend that you follow the [Microsoft Teams machine-wide installation guidelines](#) and avoid using the .exe installer that installs Microsoft Teams in *AppData*. Instead, install in *C:\Program Files (x86)\Microsoft\Teams* by using the **ALLUSER=1** flag from the command line.

```
msiexec /i <path_to_msi> /l*v <install_logfile_name> ALLUSER=1 ALLUSERS=1
```

This example also uses the **ALLUSERS=1** parameter. When you set this parameter, the Microsoft Teams Machine-Wide Installer appears in **Programs and Features** in the Control Panel and in **Apps & features** in Windows Settings for all users of the computer. All users can then uninstall Microsoft Teams if they have administrator credentials.

It's important to understand the difference between **ALLUSERS=1** and **ALLUSER=1**. You can use the **ALLUSERS=1** parameter in non-VDI and VDI environments. Use the **ALLUSER=1** parameter only in VDI environments to specify a per-machine installation.

In **ALLUSER=1** mode, the Microsoft Teams application doesn't auto-update whenever there's a new version. We recommend this mode for non-persistent environments, such as hosted shared apps or desktops out of a Windows Server or Windows 10 random/pooled catalogs. For more information, see [Install Microsoft Teams using MSI](#) (VDI Installation section).

Suppose you have Windows 10 dedicated persistent VDI environments. You want the Microsoft Teams application to auto-update and prefer Microsoft Teams to install per-user under `Appdata/Local`. In this case, use the `.exe` installer or the MSI without **ALLUSER=1**.

## For App Layering

If using Citrix App Layering to manage VDA and Microsoft Teams installations in different layers, deploy this registry key on Windows before installing Microsoft Teams with **ALLUSER=1**. For information, see [Optimization for Microsoft Teams with Citrix App Layering](#) in the list of features managed through the registry.

## Profile Management recommendations

We recommend using the machine-wide installer for Windows Server and Pooled VDI Windows 10 environments.

When the **ALLUSER=1** flag is passed to the MSI from the command line (the machine-wide installer), the Microsoft Teams app installs under `C:\Program Files (x86)` (~300 MB). The app uses `AppData\Local\Microsoft\TeamsMeetingAddin` for logs and `AppData\Roaming\Microsoft\Teams` (~600–700 MB) for user specific configurations, caching of elements in the user interface, and so forth.

### Important:

If you don't pass the **ALLUSER=1** flag, the MSI places the Teams.exe installer and `setup.json` under `C:\Program Files (x86)\Teams Installer`.

A registry key (TeamsMachineInstaller) is added under:

```
HKEY_LOCAL_MACHINE \SOFTWARE\WOW6432Node\Microsoft\Windows\CurrentVersion
\Run
```

The next user logon triggers the final installation in **AppData**.

## Machine-wide installer

The following is an example of folders, desktop shortcuts, and registries created by installing Microsoft Teams machine-wide installer on a Windows Server 2016 64-bit VM:

*Folder:*

- `C:\Program Files (x86)\Microsoft\Teams`
- `C:\Users\\AppData\Roaming\Microsoft\Teams`

*Desktop shortcut:*

`C:\Program Files (x86)\Microsoft\Teams\current\Teams.exe`

*Registry:*

- `HKEY_LOCAL_MACHINE \SOFTWARE\WOW6432Node\Microsoft\Windows\CurrentVersion\Run`
- `HKEY_LOCAL_MACHINE \SOFTWARE\Microsoft\Windows\CurrentVersion\Run`
- `HKEY_CURRENT_USER \SOFTWARE\Microsoft\Windows\CurrentVersion\Run`

## Recommendations

- We recommend disabling auto-start by deleting the Microsoft Teams registry keys. Doing so prevents many logons that occur at the same time (for example, at the beginning of your work day) from spiking up the VM's CPU.
- If the Virtual Desktop does not have a GPU/vGPU, we recommend setting **Disable GPU hardware acceleration** in the Microsoft Teams **Settings** to improve performance. This setting ("`disableGpu`": `true`) is stored in `%Appdata%\Microsoft\Teams` in `desktop-config.json`. You can use a logon script to edit that file and set the value to `true`.
- If using Citrix Workspace Environment Management (WEM), enable **CPU Spikes Protection** to manage processor consumption for Microsoft Teams.

## Per-user installer

When using the `.exe` installer, the installation process differs. All the files are placed in AppData.

*Folder:*

- `C:\Users\\AppData\Local\Microsoft\Teams`
- `C:\Users\\AppData\Local\Microsoft\TeamsPresenceAddin`
- `C:\Users\\AppData\Local\Microsoft\TeamsMeetingAddin`
- `C:\Users\\AppData\Local\SquirrelTemp`
- `C:\Users\\AppData\Roaming\Microsoft\Teams`

*Desktop shortcut:*

`C:\Users\\AppData\Local\Microsoft\Teams\Update.exe --processStart "Teams.exe"`

*Registry:*

`HKEY_CURRENT_USER \SOFTWARE\Microsoft\Windows\CurrentVersion\Run`

## Best Practices

The best practice recommendations are based on the use-case scenarios.

Using Microsoft Teams with a non-persistent setup requires a profile caching manager for efficient Microsoft Teams runtime data synchronization. With a profile caching manager, the appropriate user-specific information (for example, user data, profile, and settings) is cached during the user session. Synchronize the data in these two folders:

- C:\Users\\AppData\Local\Microsoft\IdentityCache
- C:\Users\\AppData\Roaming\Microsoft\Teams

## Microsoft Teams cached content exclusion list for non-persistent setup

Exclude the following items from the Microsoft Teams caching folder, %AppData%/Microsoft/Teams. Excluding these items helps reduce the user caching size to further optimize your non-persistent setup.

Exclusion list – files

- Roaming\Microsoft\Teams\\* .txt

Exclusion list – directories

- Roaming\Microsoft\Teams\Logs
- Roaming\Microsoft\Teams\media-stack
- Roaming\Microsoft\Teams\Service Worker\CacheStorage
- Roaming\Microsoft\Teams\Application Cache
- Roaming\Microsoft\Teams\Cache
- Roaming\Microsoft\Teams\GPUCache
- Roaming\Microsoft\Teams\meeting-addin\Cache (Critical for issues where the Add-in is missing in Outlook)

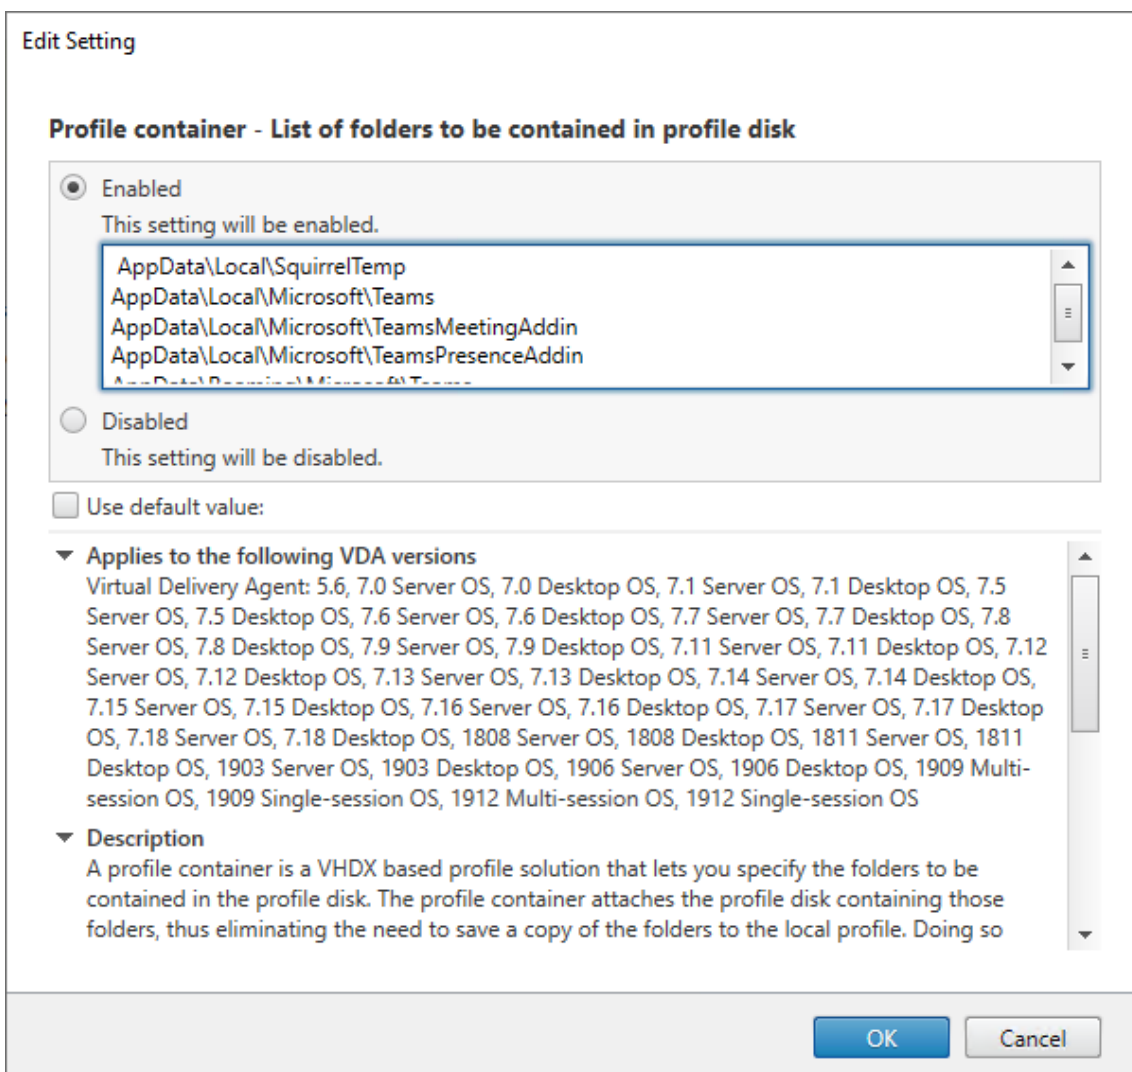
## Use case: single-session scenario

In this scenario, the end user uses Microsoft Teams in one location at a time. They don't need to run Microsoft Teams in two Windows sessions at the same time. For instance, in a common virtual desktop deployment, each user is assigned to one desktop, and Microsoft Teams is deployed in the virtual desktop as one application.

We recommend enabling the Citrix Profile container and redirecting per-user directories listed in Per-user installer into the container.

1. Deploy the Microsoft Teams machine wide installer (**ALLUSER=1**) in the golden image.
2. Enable Citrix Profile Management and set up the user profile store with the proper permissions.

3. Enable the following Profile Management policy setting: **File system > Synchronization > Profile container – List of folders to be contained in profile disk.**



List all the per-user directories into this configuration. Or, you can also configure these settings using the Citrix Workspace Environment Management (WEM) service.

4. Apply the settings to the correct delivery group.
5. Log in to validate the deployment.

## System requirements

### Minimum recommended version - Delivery Controller (DDCs) 1906.2

If you're using an earlier version, see [Enable optimization of Microsoft Teams](#).

Supported operating systems:

- Windows Server 2019, 2016, 2012R2 Standard and Datacenter Editions, and with the Server Core option

### **Minimum version - Virtual Delivery Agents (VDAs) 1906.2**

Supported operating systems:

- Windows 10 64-bit, versions 1607 and higher. (VM hosted apps aren't supported.)
- Windows Server 2019, 2016, and 2012 R2 (Standard and Datacenter Editions).

Requirements:

- BCR\_x64.msi - the MSI that contains the Microsoft Teams optimization code and starts automatically from the GUI. If you're using the command line interface for the VDA installation, don't exclude it.

### **Recommended version - Citrix Workspace app 2009 for Windows and Minimum version - Citrix Workspace app 1907 for Windows**

- Windows 8 and 10 (32-bit and 64-bit editions, including Embedded editions) - Support for Windows 7 ended at Version 2006.
- Windows 10 IoT Enterprise 2016 LTSB (v1607) and 2019 LTSC (v1809).
- Processor (CPU) architectures supported: x86 and x64 (ARM isn't supported).
- Endpoint requirement: Approximately 2.2–2.4 GHz dual core CPU that can support 720p HD resolution during a peer-to-peer video conference call.
- Dual or quad-core CPUs with lower base speeds (~1.5 GHz) equipped with Intel Turbo Boost or AMD Turbo Core that can boost up to at least 2.4 GHz.
- HP Thin Clients verified: t630/t640, t730/t740, mt44/mt45.
- Dell Thin Clients verified: 5070, 5470 Mobile TC.
- 10ZiG Thin Clients verified: 4510 and 5810q.
- For a complete list of verified endpoints, see [Thin Clients](#).
- Citrix Workspace app requires at least 600 MB free disk space and 1 GB RAM.
- Microsoft .NET Framework minimum requirement is version 4.6.2. Citrix Workspace app automatically downloads and installs .NET Framework if it isn't present in the system.

### **Minimum version - Citrix Workspace app 2006 for Linux**

For more information, see *Optimization for Microsoft Teams* in [What's new in 2006](#).

Software:

- GStreamer 1.0 or later or Cairo 2
- libc++-9.0 or later

- libgdk 3.22 or later
- OpenSSL 1.1.1d
- x64 Linux distribution

Hardware:

- Minimum 1.8 GHz dual-core CPU that can support 720p HD resolution during a peer-to-peer video conference call
- Dual or quad-core CPU with a base speed of 1.8 GHz and a high Intel Turbo Boost speed of at least 2.9 GHz

For more information, see [Prerequisites to install Citrix Workspace app](#).

### Minimum version - Citrix Workspace app 2012 for Mac

Supported operating systems:

- macOS Catalina (10.15).
- macOS Big Sur Beta 8 in test environments only. Don't use in production environments.

Features supported:

- Audio
- Video
- Screen sharing optimization (incoming and outgoing)

Microsoft Teams optimization works by default if the user has Citrix Workspace app 2012 or later and macOS 10.15.

If you want to disable Microsoft Teams optimization, run this command in a terminal and restart Workspace app:

```
defaults write com.citrix.receiver.nomas mtopEnabled -bool NO
```

### Enable optimization of Microsoft Teams


To enable optimization for Microsoft Teams, use the Manage console policy described in [Microsoft Teams redirection policy](#). It's **ON** by default. In addition to this policy being enabled, HDX checks to verify that the version of the Citrix Workspace app is at least the minimum required version. If you enabled the policy and the Citrix Workspace app version is supported, the **HKEY\_CURRENT\_USER\Software\Citrix\HDXMediaStream\MSTeamsRedirSupport** registry key is set to **1** automatically on the VDA. Microsoft Teams reads the key to load in VDI mode.

**Note:**


If you're using version 1906.2 VDAs or higher with older controller versions (for example, version

7.15) that don't have the policy available in the Manage console (Studio), your VDA can still be optimized. HDX optimization for Microsoft Teams is enabled by default in the VDA.


If you click **About > Version**, the **Citrix HDX Optimized** legend displays:



If you see **Citrix HDX Not Connected**, the Citrix API is loaded in Microsoft Teams. Loading the API is the first step toward redirection. But there's an error in later parts of the stack. The error is most likely in the VDA services or the Citrix Workspace app.



If you don't see any legend, Microsoft Teams failed to load the Citrix API. Exit Microsoft Teams by right-clicking the notification area icon and restarting. Make sure that the Manage console policy isn't set to **Prohibited** and that the Citrix Workspace app version is supported.



## Network requirements

Microsoft Teams relies on Media Processor servers in Office 365 for meetings or multiparty calls. Microsoft Teams relies on Office 365 Transport Relays for these scenarios:

- Two peers in a point-to-point call don't have direct connectivity.
- A participant does not have direct connectivity to the media processor.

So the network health between the peer and the Office 365 cloud determines the performance of the call.

We recommend evaluating your environment to identify any risks and requirements that can influence your overall cloud voice and video deployment.

Use the [Skype for Business Network Assessment Tool](#) to test if your network is ready for Microsoft Teams. For support information, see [Support](#).

## Summary of key network recommendations for Real Time Protocol (RTP) traffic

- Connect to the Office 365 network as directly as possible from the branch office.
- If you must use any of the following at the branch office, make sure that RTP/UDP Teams traffic is unhindered. HdxTeams.exe doesn't honor explicit proxies configured on the endpoint.
  - Bypass proxy servers
  - Network SSL intercept
  - Deep packet inspection devices



- VPN hairpins (use split tunneling if possible)
- Plan for and provide sufficient bandwidth.
- Check each branch office for network connectivity and quality.

The WebRTC media engine in the Workspace app (HdxTeams.exe) uses the Secure Real-time Transport Protocol (SRTP) for multimedia streams that are offloaded to the client. SRTP provides confidentiality and authentication to RTP by using symmetric keys (128 bit) to encrypt media and control messages and uses the AES encryption cipher in counter mode.

The following metrics are recommended for a positive user experience:

Metric	Endpoint to Office 365
Latency (one way)	< 50 msec
Latency (RTT)	< 100 msec
Packet Loss	<1% during any 15s interval
Packet inter-arrival jitter	<30ms during any 15s interval

For more information, see [Prepare your organization's network for Microsoft Teams](#).

In terms of bandwidth requirements, optimization for Microsoft Teams can use a wide variety of codecs for audio (OPUS/G.722/PCM G711) and video (H264).

The peers negotiate these codecs during the call establishment process using the Session Description Protocol (SDP) Offer/Answer.

Citrix minimum recommendations per user are:

Type	Bandwidth	Codec
Audio (each way)	~ 90 kbps	G.722
Audio (each way)	~ 60 kbps	Opus*
Video (each way)	~ 700 kbps	H264 360p @ 30 fps 16:9
Screen sharing	~ 300 kbps	H264 1080p @ 15 fps

\* Opus supports constant and variable bitrate encoding from 6 kbps up to 510 kbps.

Opus is the preferred codec for peer-to-peer calls between two optimized VDI users.

G.722 and H264 are the preferred codecs for a VDI user joining a meeting.

## Proxy servers

Depending on the location of the proxy, consider the following:

- Proxy configuration on the VDA:

If you configure an explicit proxy server in the VDA and route connections to localhost through a proxy, redirection fails. To configure the proxy correctly, select the **Bypass proxy servers for local address** setting in **Internet Options > Connections > LAN Settings > Proxy Servers** and make sure `127.0.0.1:9002` is bypassed.

If you use a PAC file, your VDA proxy configuration script from the PAC file must return **DIRECT** for `wss://127.0.0.1:9002`. If not, optimization fails. To make sure that the script returns **DIRECT**, use `shExpMatch(url, "wss://127.0.0.1:9002/*")`.

- Proxy configuration on Citrix Workspace app:

If the branch office is configured to access the internet through a proxy, these apps support proxy servers:

- Citrix Workspace app for Windows version 2012 (Negotiate/Kerberos, NTLM, Basic, and Digest)
- Citrix Workspace app for Linux version 2101 (anonymous authentication)
- Citrix Workspace app for Mac version 2104 (anonymous authentication)

Client devices with earlier releases of Citrix Workspace app can't read proxy configurations. These devices send traffic directly to Office 365 TURN servers.

### Important:

Verify that the client device can connect to the DNS server to perform DNS resolutions. A client device must be able to resolve three Microsoft Teams TURN server's FQDNs: `worldaz.turn.teams.microsoft.com`, `usaz.turn.teams.microsoft.com`, and `euaz.turn.teams.microsoft.com`.

## Call establishment and media flow paths

When possible, the HDX media engine in the Citrix Workspace app (`HdxTeams.exe`) tries to establish a direct network Secure Real-time Transport Protocol (SRTP) connection over User Datagram Protocol (UDP) in a peer-to-peer call. If the UDP ports are blocked, the media engine falls back to TCP 443.

The HDX media engine supports ICE, Session Traversal Utilities for NAT (STUN), and Traversal Using Relays around NAT (TURN) for candidate discovery and establishing connection.

Suppose that there is no direct path between the two peers or between a peer and a conference server, say if the user is joining a multi-party call or meeting. `HdxTeams.exe` uses a Microsoft Teams transport relay server in Office 365 to reach the other peer or the media processor, where meetings are hosted.

The user's client machine must have access to two Office 365 subnet IP address ranges and 4 UDP ports. For more information, see the Architecture diagram in the Call setup and [Office 365 URLs and IP address ranges ID 11](#).

ID	Category	Addresses	Destination Ports
11	Optimize required	13.107.64.0/18, 52.112.0.0/14, 52.120.0.0/14	<b>UDP:</b> 3478, 3479, 3480, 3481, <b>TCP:</b> 443 (fallback)

These ranges contain both Transport Relays and media processors.

The Microsoft Teams Transport Relays provide STUN and TURN functionality, but they are not ICE endpoints. Also, the Microsoft Teams Transport Relays don't terminate media or perform any transcoding. They can bridge TCP (if HdxTeams.exe uses TCP) to UDP when they forward traffic to other peers or media processors.

HdxTeams.exe contacts the closest Microsoft Teams Transport Relay in the Office 365 cloud. HdxTeams.exe uses anycast IP and port 3478–3481 UDP (different UDP ports per workload, though multiplexing can happen) or 443 TCP TLSv1.2 for fallbacks. Call quality depends on the underlying network protocol. Because UDP is always recommended over TCP, we advise you to design your networks to accommodate UDP traffic in the branch office.

If Microsoft Teams loaded in optimized mode and HdxTeams.exe is running on the endpoint, ICE failures might cause a call setup failure or one-way-only audio/video. When a call can't be completed or media streams aren't full duplex, check the **Wireshark trace** on the endpoint first. For more information about the ICE candidate gathering process, see "Collecting logs" in the [Support](#) section.

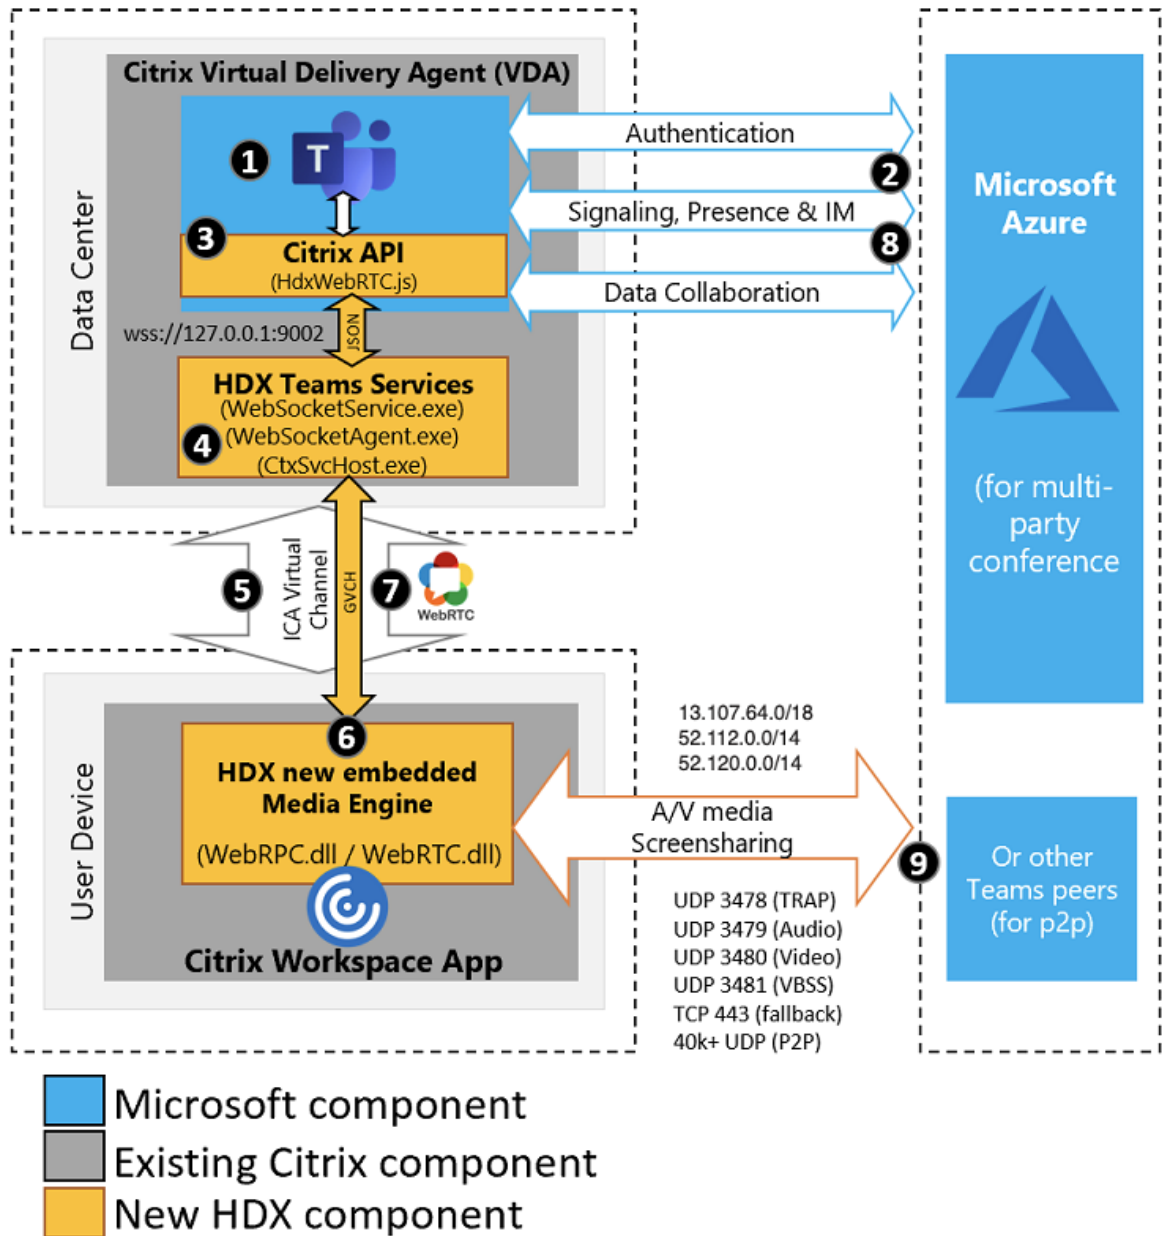
**Note:**

If the endpoints don't have internet access, the users might still be able to make a peer-to-peer call if they are both on the same LAN. Meetings fail. In this case, there's a 30-second timeout before the call setup begins.

### Call setup

Use this architecture diagram as a visual reference for the call flow sequence. The corresponding steps are indicated in the diagram.

### Architecture



1. Start Microsoft Teams.
2. Microsoft Teams authenticates to O365. Tenant policies are pushed down to the Microsoft Teams client, and relevant TURN and signaling channel information is relayed to the app.
3. Microsoft Teams detects that it's running in a VDA and makes API calls to the Citrix JavaScript API.
4. Citrix JavaScript in Microsoft Teams opens a secure WebSocket connection to WebSocketService.exe running on the VDA (127.0.0.1:9002), which spawns WebSocketAgent.exe inside the user session.
5. WebSocketAgent.exe instantiates a generic virtual channel by calling into the Citrix HDX Teams

Redirection Service (CtxSvcHost.exe).

6. Citrix Workspace app's wfica32.exe (HDX engine) spawns a new process called HdxTeams.exe, which is the new WebRTC engine used for Microsoft Teams optimization.
7. HdxTeams.exe and Teams.exe have a 2-way virtual channel path and can start processing multimedia requests.  
---User calls---
8. **Peer A** clicks the **call** button. Teams.exe communicates with the Microsoft Teams services in Office 365, establishing an end-to-end signaling path with **Peer B**. Microsoft Teams asks HdxTeams for a series of supported call parameters (codecs, resolutions, and so forth, which is known as a Session Description Protocol (SDP) offer). These call parameters are then relayed using the signaling path to the Microsoft Teams services in Office 365 and from there to the other peer.
9. The SDP offer/answer (single-pass negotiation) takes place through the signaling channel, and the ICE connectivity checks (NAT and Firewall traversal using STUN bind requests) complete. Then, Secure Real-time Transport Protocol (SRTP) media flows directly between HdxTeams.exe and the other peer (or Office 365 conference servers if it's a meeting).

## Microsoft Phone System


Phone System is Microsoft's technology that enables call control and PBX in the Office 365 cloud with Microsoft Teams. Optimization for Microsoft Teams supports Phone System using Office 365 Calling Plans or Direct Routing. With Direct Routing, you connect your own supported session border controller to the Microsoft Phone System directly without any additional on-premises software.

## Firewall considerations

When users start an optimized call using the Microsoft Teams client for the first time, they might notice a warning with the **Windows firewall** settings. The warning asks for users to allow communication for HdxTeams.exe (HDX Overlay Teams).



The following four entries are added under **Inbound Rules** in the **Windows Defender Firewall > Advanced Security** console. You can apply more restrictive rules if you want.



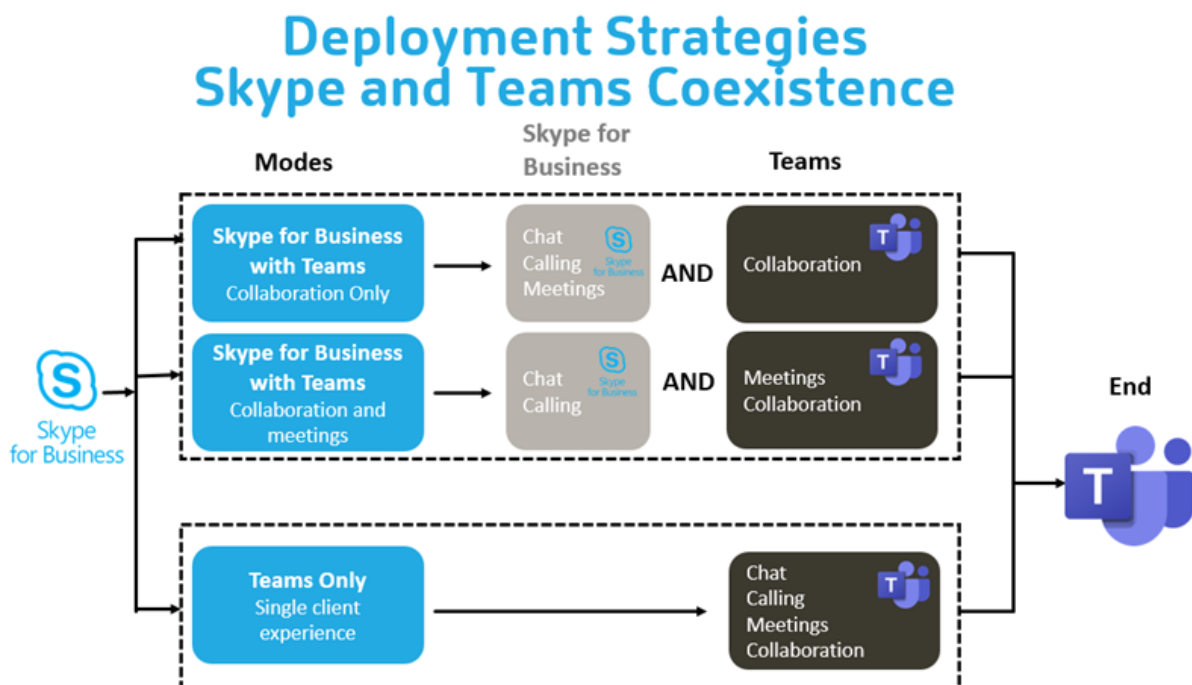
### Microsoft Teams and Skype for Business Coexistence

You can deploy Microsoft Teams and Skype for Business side by side as two separate solutions with overlapping capabilities.

For more information, see [Understand Microsoft Teams and Skype for Business coexistence and interoperability](#).

Citrix RealTime Optimization Pack and HDX optimization for Microsoft Teams multimedia engines then honor the configuration set in your environment. Examples include island modes, Skype for Business with Microsoft Teams collaboration, and Skype for Business with Microsoft Teams collaboration and meetings.

Peripheral access can be granted only to a single application at the time. For example, webcam access by the RealTime Media Engine during a call locks the imaging device. When the device is released, it becomes available for Microsoft Teams.



## Citrix SD-WAN: optimized network connectivity for Microsoft Teams

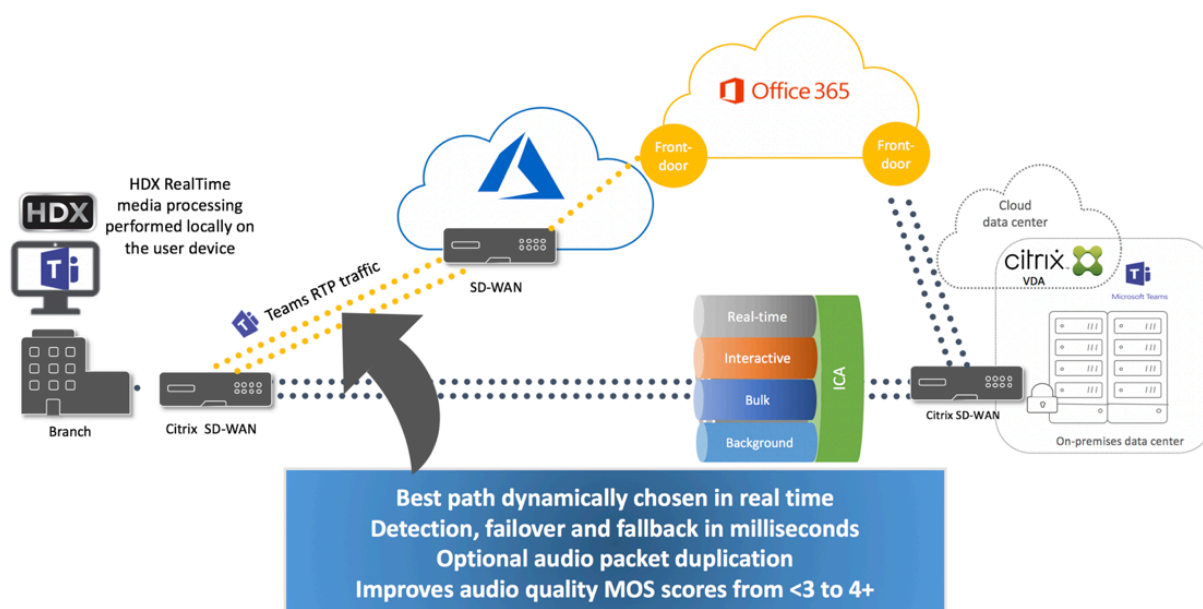
Optimal audio and video quality require a network connection to the Office 365 cloud that has low latency, low jitter, and low packet loss. Backhauling of Microsoft Teams audio-video RTP traffic from Citrix Workspace app users at branch office locations to a data center before going to the internet can add excessive latency. It might also cause congestion on WAN links. Citrix SD-WAN optimizes connectivity for Microsoft Teams following Microsoft Office 365 network connectivity principles. Citrix SD-WAN uses the Microsoft REST-based Office 365 IP address and web service and proximate DNS to identify, categorize, and steer Microsoft Teams traffic.

Business broadband internet connections in many areas suffer from intermittent packet loss, periods of excessive jitter, and outages.

Citrix SD-WAN offers two solutions to preserve Microsoft Teams audio-video quality when network health is variable or degraded.

- If you use Microsoft Azure, a Citrix SD-WAN virtual appliance (VPX) deployed in the Azure VNET provides advanced connectivity optimizations. These optimizations include seamless link failover and audio packet racing.
- Citrix SD-WAN customers can connect to Office 365 through the Citrix Cloud Direct service. This service provides reliable and secure delivery for all internet-bound traffic.

If the quality of the branch office internet connection isn't a concern, it might be enough to minimize latency by steering Microsoft Teams traffic directly from the Citrix SD-WAN branch appliance to the nearest Office 365 front door. For more information, see [Citrix SD-WAN Office 365 optimization](#).



## Gallery view and active speakers in Microsoft Teams

Microsoft Teams supports **Gallery**, **Large gallery**, and **Together mode** layouts.


Microsoft Teams displays a 2x2 grid with video streams of four participants (known as **Gallery**). In this case, Microsoft Teams sends four video streams to the client device for decoding. When there are more than four participants sharing video, only the last four most active speakers appear on the screen.

Microsoft Teams also provides the large gallery view with a grid up to 7x7. As a result, the Microsoft Teams conference server composites a single video feed and sends it to the client device for decoding, resulting in lower CPU consumption. This single, matrix-style feed might include users' self-preview video as well.

Lastly, Microsoft Teams supports **Together mode**, which is part of the new meeting experience. Using AI segmentation technology to digitally place participants in a shared background, Microsoft Teams puts all participants in the same auditorium.

The user can control these modes during a conference call by selecting **Gallery**, **Large gallery**, or **Together mode** layouts in the ellipses menu.





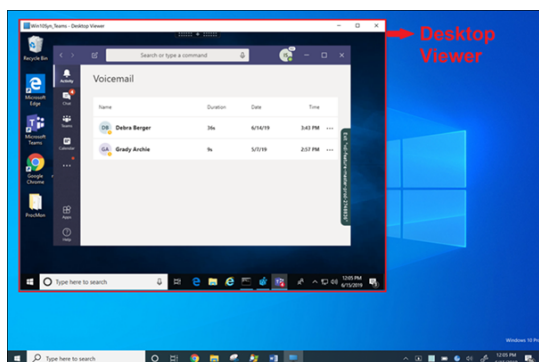
### Screen sharing in Microsoft Teams

Microsoft Teams relies on video-based screen sharing (VBSS), effectively encoding the desktop being shared with video codecs like H264 and creating a high-definition stream. With HDX optimization, incoming screen sharing is treated as a video stream. So if you are in the middle of a video call and the other peer starts to share the desktop, the original camera video feed is paused. Instead, the screen sharing video feed shows. The peer must then manually resume the camera sharing.

Outgoing screen sharing is also optimized and offloaded to Citrix Workspace app (version 1907 or higher). In this case, HdxTeams.exe captures and transmits only the Citrix Desktop Viewer (CD-

Viewer.exe) window. If you want to share a local application running in your client machine, you can overlay it on CDViewer and it's also captured.

Multi-monitor: In cases where CDViewer is in full screen mode and spanning across multi-monitor setups, only the primary monitor is shared. Users must drag the application in the virtual desktop to the primary monitor for the other peer on the call to see it.

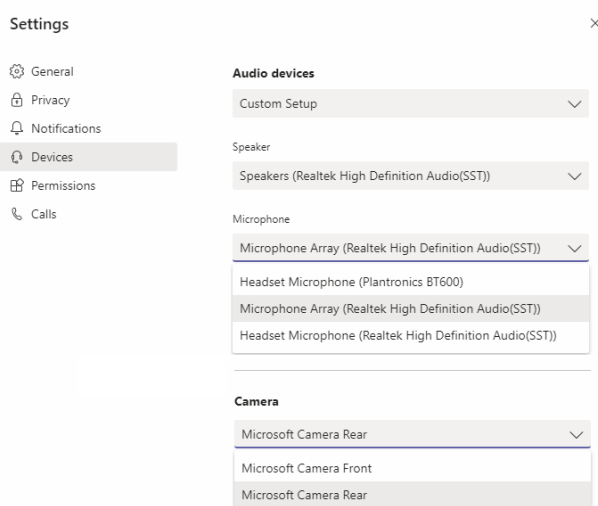


### Note:

If you're publishing Microsoft Teams as a standalone, seamless application, screen sharing captures the local desktop of your physical endpoint in Citrix Workspace app minimum version 1909.

## Peripherals in Microsoft Teams

When optimization for Microsoft Teams is active, the Citrix Workspace app accesses the peripherals (headset, microphone, cameras, speakers, and so forth). Then the peripherals are properly listed in the Microsoft Teams UI (**Settings > Devices**).



Microsoft Teams does not access the devices directly. Instead, it relies on HdxTeams.exe for acquiring, capturing, and processing the media. Microsoft Teams lists the devices for the user to select.

### Recommendations:

- [Microsoft Teams certified headsets](#) with built-in echo cancellation. In setups with multiple peripherals, where microphone and speakers are on separate devices, there might be an echo. An example is a webcam with a built-in microphone and a monitor with speakers. When using external speakers, place them as far as possible from the microphone and from any surface that might refract the sound into the microphone.
- [Microsoft Teams certified cameras](#), although [Skype for Business certified peripherals](#) are compatible with Microsoft Teams.
- HdxTeams.exe can't take advantage of CPU offloading with webcams that perform on-board H.264 encoding -UVC 1.1 and 1.5.

**Note:**

HdxTeams.exe supports only these specific audio device formats (channels, bit depth, and sample rate):

- Playback Devices: up to 2 channels, 16 bit, frequencies up to 96,000 Hz
- Recording Devices: up to 4 channels, 16 bit, frequencies up to 96,000 Hz

Even if one speaker or microphone does not match the expected settings, device enumeration in Microsoft Teams fails and **None** displays under **Settings > Devices**.

**Webrpc** logs in **HdxTeams.exe** show this type of information:

```
Mar 27 20:58:22.885 webrtcapi.WebRTCEngine Info: init. initializing...
```

```
Mar 27 20:58:23.190 webrtcapi.WebRTCEngine Error: init. couldn't create audio module!
```

As a workaround, disable the specific device or:

1. Open the **Sound Control Panel** (mmsys.cpl).
2. Select the playback or recording device.
3. Go to **Properties > Advanced** and change the settings to a supported mode.

## Fallback mode

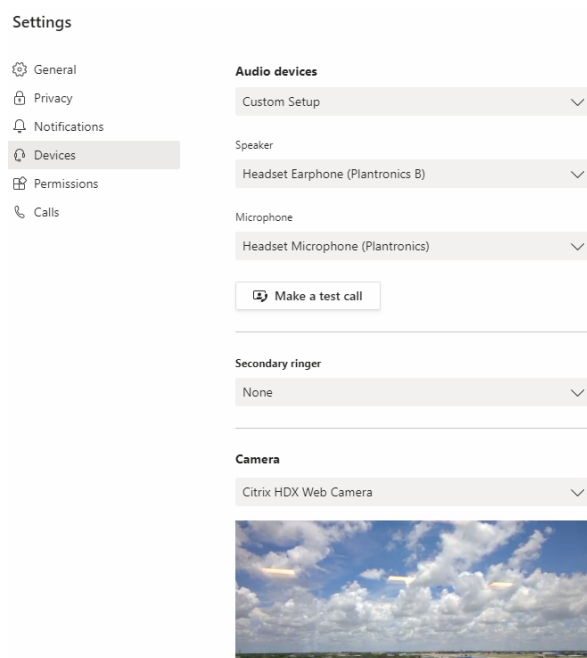
If Microsoft Teams fails to load in optimized VDI mode, the VDA falls back to legacy HDX technologies like webcam redirection and client audio and microphone redirection. In the unoptimized mode, the peripherals are mapped to the VDA. The peripherals appear to the Microsoft Teams app as if they were locally attached to the virtual desktop.

You can now granularly control the fallback mechanism by setting the registry keys in the VDA. For information, see [Microsoft Teams fallback mode](#) in the list of features managed through the registry.

This feature requires Microsoft Teams version 1.3.0.13565 or later.

To determine if you are in optimized or unoptimized mode when looking at the **Settings > Devices**

tab in Microsoft Teams, the main difference is the camera name. If Microsoft Teams loaded in unoptimized mode, legacy HDX technologies launch. The webcam name has the **Citrix HDX** suffix as shown in the following graphic. The speaker and microphone device names might be slightly different (or truncated) when compared to the optimized mode.



When legacy HDX technologies are used, Microsoft Teams doesn't offload audio, video, and screen sharing processing to the endpoint's Citrix Workspace app WebRTC media engine. Instead, HDX technologies use server-side rendering. Expect high CPU consumption on the VDA when you turn on video. Real-time audio performance might not be optimal.

### Known limitations

#### Citrix limitations

Limitations on Citrix Workspace app:

- DTMF tones aren't supported.
- HID buttons - Answer and end call aren't supported. Volume up and down are supported.
- When doing screen sharing in multi-monitor setups, only the main monitor is shared.
- We support only one video stream from an incoming camera or screen share stream. When there's an incoming screen share, that screen share is shown instead of the video of the dominant speaker.
- Secondary ringer (**Teams > Settings > Devices**) isn't supported.
- QoS settings in the Admin Center for Microsoft Teams don't apply for VDI users.
- App protection add-on feature for the Citrix Workspace app prevents outgoing screen sharing.
- The zoom in and zoom out function in Microsoft Teams isn't supported.

#### Limitation on the VDA:

- When you configure the Citrix Workspace app High DPI setting to **Yes** or to **No, use the native resolution**, the redirected video window appears out of place when the monitor's DPI scaling factor is set to anything above 100%.

#### Limitations on Citrix Workspace app and the VDA:

- Outgoing screen sharing: Application sharing isn't supported.
- You can only control the volume of an optimized call using the volume bar on the client machine – not on the VDA.

#### Microsoft limitations

- The options to blur or customize the background aren't supported.
- A 3x3 gallery view isn't supported. Microsoft Teams dependency – contact Microsoft for when to expect a 3x3 grid.
- Interoperability with Skype for Business is limited to audio calls, no video modality.
- Incoming and outgoing video stream maximum resolution is 720p. Microsoft Teams dependency – contact Microsoft for when to expect 1080p.
- PSTN call ringback tone isn't supported.
- Media bypass for Direct Routing isn't supported.

#### Citrix and Microsoft limitations

- When doing screen sharing, the option **include system audio** isn't available.
- Pop out chat (also known as multi-window chat or the new meeting experience) isn't supported.
- Breakout rooms are supported for VDI participants. Microsoft Teams doesn't support breakout rooms if the organizer is a VDI user.
- Give control and take control: Not supported during a desktop screen sharing or application sharing session. Supported only during a [PowerPoint sharing session](#).
- E911 and Location-Based Routing are not supported.

#### Additional information

- [Monitor, troubleshoot, and support Microsoft Teams](#)
- [Deploy the Teams desktop app to the VM](#)
- [Install Microsoft Teams using MSI \(VDI Installation section\)](#)
- [Thin clients](#)
- [Skype for Business Network Assessment Tool](#)
- [Understand Microsoft Teams and Skype for Business coexistence and interoperability](#)

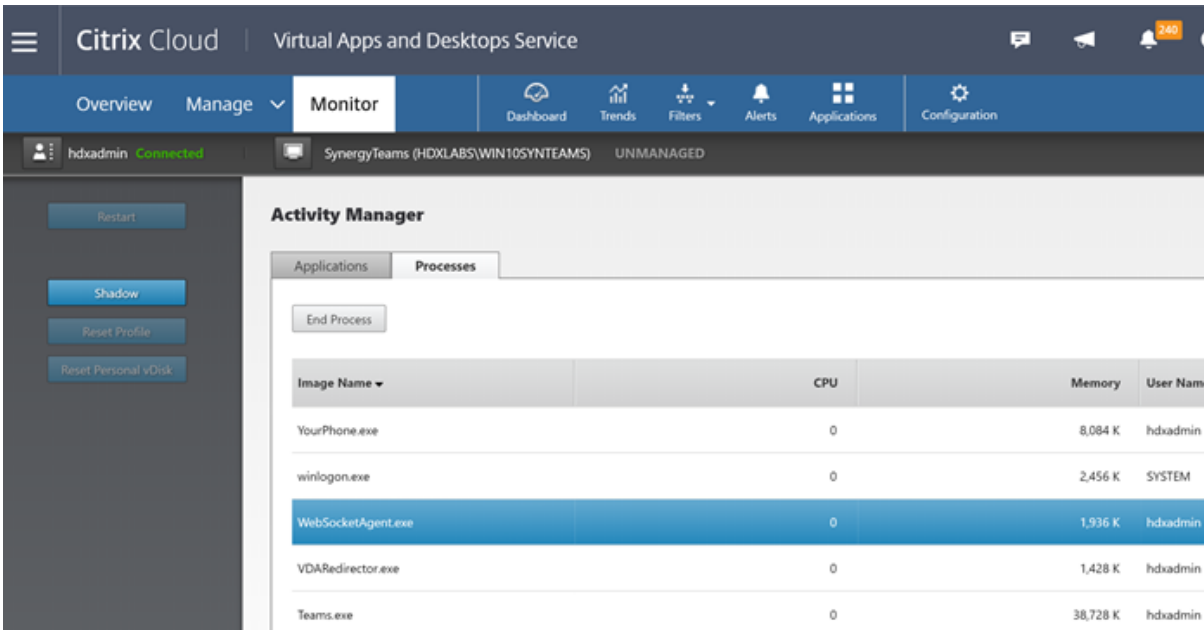
## Monitor, troubleshoot, and support Microsoft Teams

March 8, 2021

### Monitor Teams

This section provides guidelines for monitoring Microsoft Teams optimization with HDX.

If the user is running in optimized mode and `HdxTeams.exe` or `HdxRtcEngine.exe` is running on the client machine, there is a process on the VDA called `WebSocketAgent.exe` running in the session. Use the **Activity Manager** in Director to see the application.



The screenshot shows the Citrix Cloud interface for the Virtual Apps and Desktops Service. The 'Monitor' tab is active, displaying the 'Activity Manager' for a session named 'SynergyTeams (HDXLABS\WIN10SYNTEAMS)'. The 'Processes' tab is selected, showing a table of active processes. The 'WebSocketAgent.exe' process is highlighted in blue.

Image Name	CPU	Memory	User Name
YourPhone.exe	0	8,084 K	hdadmin
winlogon.exe	0	2,456 K	SYSTEM
WebSocketAgent.exe	0	1,936 K	hdadmin
VDARedirector.exe	0	1,428 K	hdadmin
Teams.exe	0	38,728 K	hdadmin

With the VDA minimum version 1912, you can monitor active Teams calls using the Citrix HDX Monitor (minimum version 3.11). The Citrix Virtual Apps and Desktops product ISO contains the latest `hdxmonitor.msi` in the folder `layout\image-full\Support\HDX Monitor`.


For more information, see *Monitoring* in the Knowledge Center article [CTX253754](#).

### Troubleshoot

This section provides troubleshooting tips for issues that you might encounter when using optimization for Microsoft Teams. For more information, see [CTX253754](#).

### On the Virtual Delivery Agent

There are four services installed by `BCR_x64.msi`. Only two are responsible for Microsoft Teams redirection in the VDA.




- **Citrix HDX Teams Redirection Service** establishes the virtual channel used in Microsoft Teams. The service relies on CtxSvcHost.exe.
- **Citrix HDX HTML5 Video Redirection Service** runs as WebSocketService.exe listening on 127.0.0.1:9002 TCP. WebSocketService.exe performs two main functions:

i. **TLS termination for secure WebSockets** receives a secure WebSocket connection from vdiCitrixPeerConnection.js, which is a component inside the Microsoft Teams app. You can track it with the Process Monitor. For more information about certificates, see the section “TLS and HTML5 video redirection, and browser content redirection” under [Communication between Controller and VDA](#).

Some antivirus and desktop security software interferes with the proper functioning of `WebSocketService.exe` and its certificates. While the Citrix HDX HTML5 Video Redirection service might be running in the `services.msc` console, the localhost 127.0.0.1:9002 TCP socket is never in listening mode as seen in netstat. Trying to restart the service causes it to hang (“Stopping...”). Ensure you apply the proper exclusions for the `WebSocketService.exe` process.




ii. **User session mapping**. When the Microsoft Teams application starts, `WebSocketService.exe` starts the `WebSocketAgent.exe` process in the user’s session in the VDA. `WebSocketService.exe` runs in Session 0 as a LocalSystem account.



You can use `netstat` to check if the `WebSocketService.exe` service is in an active listening state in the VDA.

Run `netstat -anob -p tcp` from an elevated command prompt window:



On a successful connection, the state changes to ESTABLISHED:

```

TCP    127.0.0.1:9002          127.0.0.1:58069        ESTABLISHED    8096
[WebSocketService.exe]
TCP    127.0.0.1:58069       127.0.0.1:9002        ESTABLISHED    748
[Teams.exe]

```

**Important:**

WebSocketService.exe listens in two TCP sockets, 127.0.0.1:9001 and 127.0.0.1:9002. Port 9001 is used for browser content redirection and HTML5 video redirection. Port 9002 is used for Microsoft Teams redirection. Ensure that you don't have any proxy configurations in the Windows OS of the VDA that can prevent a direct communication between Teams.exe and WebSocketService.exe. Sometimes, when you configure an explicit proxy in Internet Explorer 11 (**Internet Options > Connections > LAN settings > Proxy Server**), connections might flow through an assigned proxy server. Verify that **Bypass proxy server for local addresses** is checked when using a manual and explicit proxy setting.

**Services locations and descriptions**

Service	Path to executable in Windows Server OS	Log on as	Description
Citrix HTML5 Video Redirection Service	"C:\Program Files (x86)\Citrix\System32\VideoRedirectionService.exe"	Local System account	Provides multiple HDX Multimedia services with the initial framework required to perform media redirection between the virtual desktop and the endpoint device.
Citrix HDX Browser Redirection Service	"C:\Program Files (x86)\Citrix\System32\BrowserRedirSvc.exe"	This account (local system)	Provides browser content redirection between the endpoint device and the virtual desktop.
Citrix Port Forwarding Service	"C:\Program Files (x86)\Citrix\System32\PortFwdSvc.exe"	This account (local service)	Provides port forwarding between the endpoint device and the virtual desktop for browser content redirection.



Service	Path to executable in Windows Server OS	Log on as	Description
Citrix HDX Teams Redirection Service	“C:\Program Files (x86)\Citrix\System32\CtxSvcHost.exe” -g TeamsSvc	Local System account	Provides Microsoft Teams redirection between the endpoint device and the virtual desktop.

### Citrix Workspace app

On the user’s endpoint, the Citrix Workspace app for Windows instantiates a new service called HdxTeams.exe or HdxRtcEngine.exe. It does so when Microsoft Teams launches in the VDA and the user tries to call or access peripherals in self-preview. If you don’t see this service, check the following:

1. Ensure that you installed as a minimum the Workspace App version 1905 for Windows. Do you see HdxTeams.exe or HdxRtcEngine.exe and the webrpc.dll binaries in the Workspace app installation path?
2. If you validated step1, do the following to check if HdxTeams.exe or HdxRtcEngine.exe is getting launched.
  - a) Exit Microsoft Teams on the VDA.
  - b) Start services.msc on VDA.
  - c) Stop the Citrix HDX Teams Redirection Service.
  - d) Disconnect the ICA session.
  - e) Connect the ICA session.
  - f) Start the Citrix HDX Teams Redirection Service.
  - g) Restart the Citrix HDX HTML5 Video Redirection Service.
  - h) Launch Microsoft Teams on the VDA.
3. If you still don’t see HdxTeams.exe or HdxRtcEngine.exe being launched on the client endpoint, do the following:
  - a) Restart the VDA.
  - b) Restart the client endpoint.

### Support

Citrix and Microsoft jointly support the delivery of Microsoft Teams from Citrix Virtual Apps and Desktops using optimization for Microsoft Teams. This joint support is the result of close collaboration between the two companies. If you have valid support contracts and you experience an issue with this solution, open a support ticket with the vendor whose code you suspect to be causing the issue. That is, Microsoft for Teams or Citrix for the optimization components.

Citrix or Microsoft receives the ticket, triages the issue, and escalates as appropriate. There is no need for you to contact each company's support team.

When you have a problem, we recommend you click **Help > Report a Problem** in the Teams UI. VDA-side logs are automatically shared between Citrix and Microsoft to resolve technical issues faster.

### Collecting logs

HDX media engine logs can be found on the user's machine (not on the VDA). In case of any issues, make sure you attach logs to your support case.

#### Windows logs:

You can locate Windows logs at %TEMP% inside the **HDXTeams** folder (AppData/Local/Temp/HDX-Teams or AppData/Local/Temp/HdxRtcEngine). Look for a .txt file called webrpc\_Day\_Month\_timestamp\_Year.txt. If you are using newer versions of Citrix Workspace app, for example Citrix Workspace app 2009.5 or later, store the logs in AppData\Local\Temp\HdxRtcEngine.

Each session creates a separate folder for logs.

#### Mac logs:

1. VDWEBRTC log - records the execution of the virtual channel.

Location: `/Users/<User Name>/Library/Logs/Citrix Workspace/CitrixViewer_<Y_M_D_H_M_S>.txt`

2. HdxRtcEngine log - records the execution of the processes on HdxRtcEngine.

Location: `$TMPDIR/hdxrtcengine/<W_M_D_H_M_S_Y>/hdxrtcengine.log`

HdxRtcEngine log is enabled by default.

#### Linux logs:

You can locate Linux logs in the `/tmp/webrpc/<current date>/` and `/tmp/hdxrtcengine/<current date>/` directory.

When establishing a call, these four ICE phases are required:

- Candidate gathering
- candidate exchange
- Connectivity checks (STUN bind requests)
- Candidate promotion

In the HdxTeams.exe/HdxRtcEngine.exe logs, the following entries are the relevant Interactive Connectivity Establishment (ICE) entries. These entries must be there for a call set-up to succeed (see this sample snippet for the gathering stage):

```
1  RPCStubs Info: -> device id = \\?\display#int3470#4&1835d135&0&uid13424
   #{
2  65e8773d-8f56-11d0-a3b9-00a0c9223196 }
3  \{
4  bf89b5a5-61f7-4127-a279-e187013d7caf }
5  label = Microsoft Camera Front groupId =
6
7  webrtcapi.RTCPeerConnection Info: createOffer. audio = 1 video = 1
8  webrtcapi.RTCPeerConnection Info: setLocalDescription.
9  >>> begin:sdp
10 [ ... ]
11
12 webrtcapi.RTCPeerConnection Info: OnSignalingChange. signaling state =
    HaveLocalOffer
13
14 webrtcapi.RTCPeerConnection Info: OnIceGatheringChange. state =
    Gathering
15
16 [ ... ]
17 >>> begin:sdp
18 candidate:840548147 1 udp 2122194687 10.108.124.215 56927 typ host
    generation 0 ufrag oVk6 network-id 1
19 <<< end:sdp
20 [ ... ]
21 >>> begin:sdp
22 candidate:1938109490 1 udp 24911871 52.114.xxx.xxx 52786 typ relay
    raddr 73.205.xxx.x rport 25651 generation 0 ufrag dDML network-id 1
    network-cost 10
23 <<< end:sdp
24 [ ... ]
25 >>> begin:sdp
26 candidate:4271145120 1 udp 1685987071 66.xxx.xxx.xxx 55839 typ srflx
    raddr 10.108.124.215 rport 55839 generation 0 ufrag uAVH network-id
    1
27 <<< end:sdp
28 [ ... ]
29
30 webrtcapi.RTCPeerConnection Info: OnIceGatheringChange. state =
    Complete webrtcapi.RTCPeerConnection Info: setRemoteDescription.
31 >>> begin:sdp
32 [ ... ]
33
34 webrtcapi.RTCPeerConnection Info: OnSignalingChange. signaling state =
    HaveRemoteOffer
```

```
35
36 <!--NeedCopy-->
```

If there are multiple ICE candidates, the order of preference is:

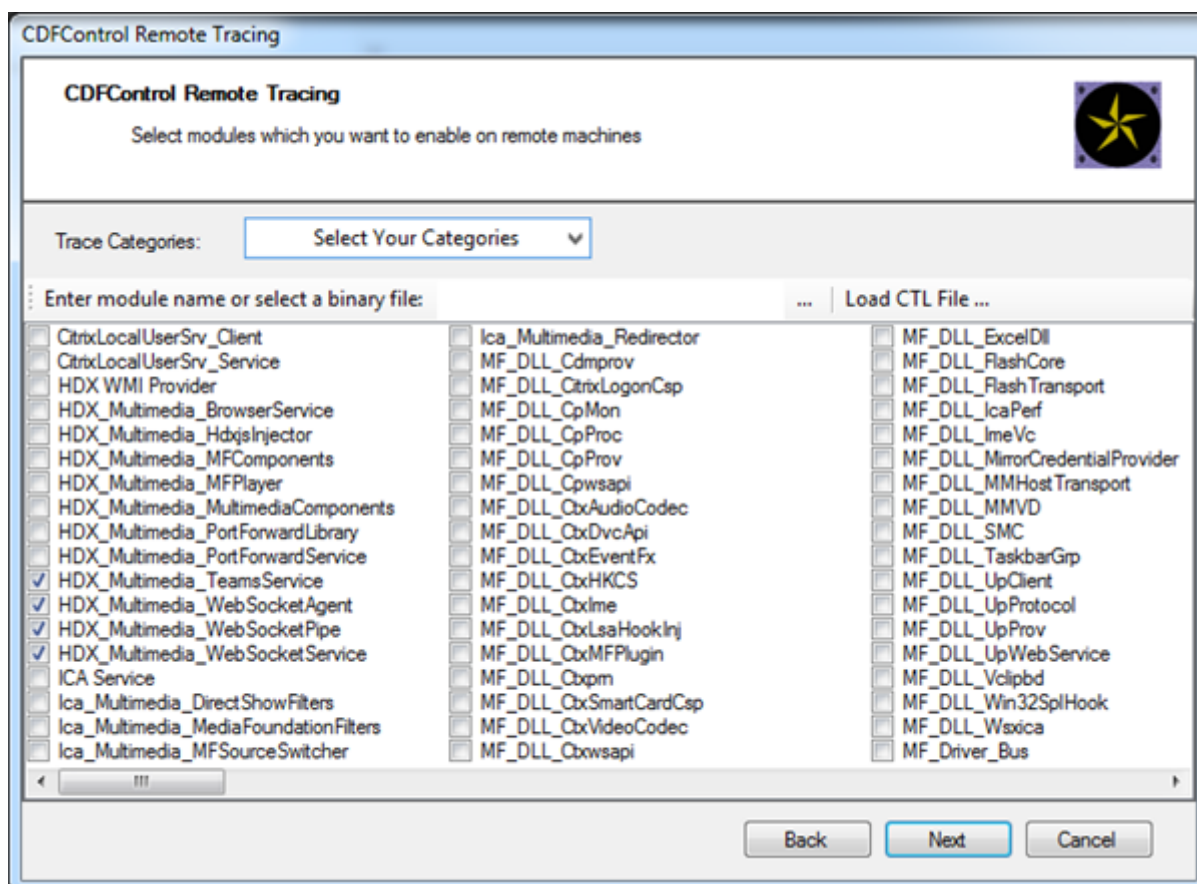
1. host
2. peer reflexive
3. server reflexive
4. transport relay

If you encounter an issue and can reproduce it consistently, we recommend clicking **Help > Report a problem** in Teams. Logs are shared between Citrix and Microsoft to resolve technical issues if you opened a case with Microsoft.


Capturing CDF traces before contacting Citrix Support is also beneficial. For more information, see the Knowledge Center article [CDFcontrol](#).

For recommendations for collecting CDF Traces, see the Knowledge Center article [Recommendations for Collecting the CDF Traces](#).

**VDA side CDF traces - Enable the following CDF trace providers:**



**Workspace app side CDF traces - Enable the following CDF trace providers:**



- IcaClient\_DriversVd\_TeamsRedir (optional)
- IcaClient\_Multimedia\_HdxTeams (requires Citrix Workspace app 2012 or later)

## Windows Media redirection

April 7, 2020

Windows Media redirection controls and optimizes the way servers deliver streaming audio and video to users. By playing the media run-time files on the client device rather than the server, Windows Media redirection reduces the bandwidth requirements for playing multimedia files. Windows Media redirection improves the performance of Windows Media Player and compatible players running on virtual Windows desktops.

If the requirements for Windows Media client-side content fetching are not met, media delivery automatically uses server-side fetching. This method is transparent to users. You can use the Citrix Scout to perform a Citrix Diagnosis Facility (CDF) trace from HostMMTransport.dll to determine the method used. For more information see, [Citrix Scout](#).

Windows Media redirection intercepts the media pipeline at the host server, captures the media data in its native compressed format, and redirects the content to the client device. The client device then recreates the media pipeline to decompress and render the media data received from the host server. Windows Media redirection works well on client devices running a Windows operating system. Those devices have the multimedia framework required to rebuild the media pipeline as it existed on the host server. Linux clients use similar open-source media frameworks to rebuild the media pipeline.

The policy setting **Windows Media Redirection** controls this feature and is **Allowed** by default. Usually, this setting increases audio and video quality rendered from the server to a level that is comparable to content played locally on a client device. In the rare cases, media playing using Windows Media redirection appears worse than media rendered using basic ICA compression and regular audio. You can disable this feature by adding the **Windows Media Redirection** setting to a policy and setting its value to **Prohibited**.

For more information about the policy settings, see [Multimedia policy settings](#).

### **Limitation:**

When you're using Windows Media Player and Remote Audio & Video Extensions (RAVE) enabled inside a session, a black screen might appear. This black screen might appear if you right-click on the video content and select **Always show Now Playing on top**.

## **General content redirection**

April 7, 2020

Content redirection allows you to control whether users access information by using applications published on servers or by using applications running locally on user devices.

### [Client folder redirection](#)

Client folder redirection changes the way client-side files are accessible on the host-side session.

- When you enable only client drive mapping on the server, client-side full volumes are automatically mapped to the sessions as Universal Naming Convention (UNC) links.
- When you enable client folder redirection on the server and the user configures it on the Windows desktop device, the portion of the local volume specified by the user is redirected.

### [Host to client redirection](#)

Consider using host to client redirection for specific uncommon use cases. Normally, other forms of content redirection might be better. We support this type of redirection only on Multi-session OS VDAs and not on Single-session OS VDAs.

### [Local App Access and URL redirection](#)

Local App Access seamlessly integrates locally installed Windows applications in to a hosted desktop environment. It does so without changing from one computer to another.

HDX technology provides **generic USB redirection** for specialty devices that don't have any optimized support or where it is unsuitable.

## **Client folder redirection**

June 18, 2021

Client folder redirection changes the way client-side files are accessible on the host-side session. If you enable only client drive mapping on the server, client-side full volumes are automatically mapped as Universal Naming Convention (UNC) links to the sessions. When you enable client folder redirection

on the server and the user configures it on the user device, the portion of the local volume specified by the user is redirected.

Only the user-specified folders appear as UNC links inside sessions. That is, instead of the complete file system on the user device. If you disable UNC links through the registry, client folders appear as mapped drives inside the session.

Client folder redirection is supported on Windows Single-session OS machines only.

Client folder redirection for an external USB drive is not saved on detaching and reattaching the device.

Enable client folder direction on the server. Then, on the client device, specify which folders to redirect. The application you use to specify the client folder options is included with the Citrix Workspace app supplied with this release.

### **Requirements:**

For servers:

- Windows Server 2019, Standard and Datacenter Editions
- Windows Server 2016, Standard and Datacenter Editions
- Windows Server 2012 R2, Standard and Datacenter Editions

For clients:

- Windows 10, 32-bit and 64-bit editions (minimum version 1607)
- Windows 8.1, 32-bit and 64-bit editions (including Embedded edition)
- Windows 7, 32-bit and 64-bit editions (including Embedded edition)

To enable client folder redirection on the server, see [Client folder redirection](#) in the list of features managed through the registry.

On the user device, specify which folders to redirect:

1. Ensure that the latest version of Citrix Workspace app is installed.
2. From the Citrix Workspace app installation directory, start CtxCFRUI.exe.
3. Choose the **Custom** radio button and add, edit, or remove folders.
4. Disconnect and reconnect your sessions for the setting to take effect.

## **Host to client redirection**

July 20, 2021

Host to client redirection allows URLs, embedded as hyperlinks in applications running on a Citrix session, to open using the corresponding application on the user endpoint device. Some common use cases for host to client redirection include:

- Redirection of websites in cases where the Citrix server doesn't have Internet or network access to the source.
- Redirection of websites when running a web browser inside the Citrix session is not desired for security, performance, compatibility, or scalability reasons.
- Redirection of specific URL types in cases where the required applications to open the URL are not installed on the Citrix server.

Host to client redirection is not intended for URLs that you access on a webpage or type in the address bar of the web browser running in the Citrix session. For redirection of URLs in web browsers, see [Bidirectional URL redirection](#) or [Browser content redirection](#).

### System requirements

- Multi-session OS VDA
- Supported clients:
  - Citrix Workspace app for Windows
  - Citrix Workspace app for Mac
  - Citrix Workspace app for Linux
  - Citrix Workspace app for HTML5
  - Citrix Workspace app for Chrome

The client device must have an application installed and configured for handling the redirection of the URL types.

### Configuration

Use the [Host to client redirection](#) Citrix policy to enable this functionality. **Host to client redirection** is disabled by default. After you enable the Host to client redirection policy, the Citrix Launcher application registers with the Windows server to ensure that it can intercept URLs and send them to the client device.

Then you must configure the Windows Group Policy to use Citrix Launcher as the default application for the required URL types. On the Citrix server VDA, create the `ServerFTAdefaultPolicy.xml` file and insert the following XML code.

```
1 <?xml version="1.0" encoding="UTF-8"?>
2
3 <DefaultAssociations>
4
5 <Association Identifier="http" ProgId="ServerFTAHTML" ApplicationName="
   ServerFTA" />
6
```



```
7 <Association Identifier="https" ProgId="ServerFTAHTML" ApplicationName=
  "ServerFTA" />
8
9 </DefaultAssociations>
10 <!--NeedCopy-->
```

From the Group Policy management console, go to **Computer configuration > Administrative Templates > Windows Components > File Explorer > Set a default associations configuration file**, and save your ServerFTAdefaultPolicy.xml file.

**Note:**

If a Citrix server doesn't have the Group Policy settings, Windows prompts users to select an application for opening URLs.

By default, we support redirection of the following URL types:

- HTTP
- HTTPS
- RTSP
- RTSPU
- PNM
- MMS

To include additional standard or custom URL types on the list for redirection, create a new **Association Identifier** line in the ServerFTAdefaultPolicy.xml file referenced earlier. For example:

```
<Association Identifier="ftp"ProgId="ServerFTAHTML"ApplicationName="ServerFTA" />
```

```
<Association Identifier="mailto"ProgId="ServerFTAHTML"ApplicationName="ServerFTA" />
```

```
<Association Identifier="customtype1"ProgId="ServerFTAHTML"ApplicationName="ServerFTA" />
```

```
<Association Identifier="customtype2"ProgId="ServerFTAHTML"ApplicationName="ServerFTA" />
```

Adding URL types to the list also requires client configuration. Create the following registry key and values on the Windows client.

**Note:**

Editing the registry incorrectly can cause serious problems that might require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the

registry before you edit it.

- Key: HKEY\_LOCAL\_MACHINE\SOFTWARE\WOW6432Node\Policies\Citrix\ICA Client\SFTA
- Value name: ExtraURLProtocols
- Value type: REG\_SZ
- Value data: Specify the required URL types separated by semicolon. Include everything before the authority portion of the URL. For example:

```
ftp://;mailto;;customtype1://;customtype2://
```

You can add URL types only for Windows clients. Clients missing the registry settings above reject redirection back to the Citrix session. Client must have an application installed and configured to handle the specified URL types.

To remove URL types from the default redirection list, create the following registry key and values on the server VDA.

- Key: HKEY\_LOCAL\_MACHINE\SOFTWARE\Wow6432Node\Citrix\SFTA
- Value name: DisableServerFTA
- Value type: DWORD
- Value data: 1
- Value name: NoRedirectClasses
- Value type: REG\_MULTI\_SZ
- Value data: Specify any combination of the values: [http](#), [https](#), [rtsp](#), [rtspu](#), [pnm](#), or [mms](#). Type multiple values on separate lines. For example:

```
http
```

```
https
```

```
rtsp
```

To enable host to client redirection for a specific set of websites, create a registry key and values on the server VDA.

- Key: HKEY\_LOCAL\_MACHINE\SOFTWARE\Wow6432Node\Citrix\SFTA
- Value name: ValidSites
- Value type: REG\_MULTI\_SZ
- Value data: Specify any combination of fully qualified domain names (FQDNs). Type multiple FQDNs on separate lines. Include the FQDN only, without protocols ([http://](#) or [https://](#)). An FQDN can include an asterisk (\*) as a wildcard character in the leftmost position only. This wildcard matches a single level of domain, which is consistent with the rules in RFC 6125. For example:

[www.example.com](http://www.example.com)

\*.example.com

**Note:**

You cannot use the **ValidSites** key in combination with the **DisableServerFTA** and **NoRedirect-Classes** keys.

## Server VDA default browser configuration

Enabling host to client redirection as referenced in this section supersedes any previous default browser configuration on the server VDA. If a web URL is not redirected, the Citrix Launcher passes the URL to the browser configured in the `command_backup` registry key. The key points to Internet Explorer by default, but you can modify it to include the path to a different browser. For more information, see [Server VDA default browser configuration](#) in the list of features managed through the registry.

## Local App Access and URL redirection

July 8, 2021

### Introduction

Local App Access seamlessly integrates locally installed Windows applications into a hosted desktop environment without switching from one desktop to another. With Local App Access, you can:

- Access applications installed locally on a physical laptop, PC, or other device directly from the virtual desktop.
- Provide a flexible application delivery solution. If users have local applications that you cannot virtualize or that IT does not maintain, those applications still behave as though they are installed on a virtual desktop.
- Eliminate the double-hop latency when applications are hosted separately from the virtual desktop. Do so by putting a shortcut to the published application on the user's Windows device.
- Use applications such as:
  - Video conferencing software such as GoToMeeting.
  - Specialty or niche applications that are not yet virtualized.
  - Applications and peripherals that would otherwise transfer large amounts of data from a user device to a server and back to the user device. For example, DVD burners and TV tuners.

In Citrix Virtual Apps and Desktops, hosted desktop sessions use URL redirection to start Local App Access applications. URL redirection makes the application available under more than one URL address. It launches a local browser (based on the browser's URL block list) by selecting embedded links within a browser in a desktop session. If you navigate to a URL that is not present in the block list, the URL is opened in the desktop session again.

URL redirection works only for desktop sessions, not application sessions. The only redirection feature you can use for application sessions is host-to-client content redirection, which is a type of server FTA (File Type Association) redirection. This FTA redirects certain protocols to the client, such as HTTP, HTTPS, RTSP, or MMS. For example, if you only open embedded links with HTTP, the links directly open with the client application. There is no URL block list or allow list support.

When Local App Access is enabled, URLs that are displayed to users as links from locally running applications, from user-hosted applications, or as shortcuts on the desktop are redirected in one of the following ways:

- From the user's computer to the hosted desktop
- From the Citrix Virtual Apps and Desktops server to the user's computer
- Rendered in the environment in which they are started (not redirected)

To specify the redirection path of content from specific websites, configure the URL allow list and URL block list on the Virtual Delivery Agent. Those lists contain multi-string registry keys that specify the URL redirection policy settings. For more information, see the [Local App Access policy settings](#).

URLs can be rendered on the VDA with the following exceptions:

- Geo/Locale information — Websites that require locale information, such as msn.com or news.google.com (opens a country specific page based on the Geo). For example, if the VDA is provisioned from a data center in the UK and the client is connecting from India, the user expects to see in.msn.com. Instead, the user sees uk.msn.com.
- Multimedia content — Websites containing rich media content, when rendered on the client device, give the end users a native experience and also save bandwidth even in high latency networks. This feature redirects sites with other media types such as Silverlight. This process is in a secure environment. That is, the URLs that the administrator approves are run on the client while the rest of the URLs are redirected to the VDA.

In addition to URL redirection, you can use FTA redirection. FTA starts local applications when a file is encountered in the session. If the local app is started, the local app must have access to the file to open it. Therefore, you can only open files that reside on network shares or on client drives (using client drive mapping) using local applications. For example, when opening a PDF file, if a PDF reader is a local app, then the file opens using that PDF reader. Because the local app can access the file directly, there is no network transfer of the file through ICA to open the file.

## Requirements, considerations, and limitations

We support Local App Access on the valid operating systems for VDAs for Windows Multi-session OS and for VDAs for Windows Single-session OS. Local App Access requires Citrix Workspace app for Windows version 4.1 (minimum). The following browsers are supported:

- Internet Explorer 11. You can use Internet Explorer 8, 9, or 10, but Microsoft supports (and Citrix recommends using) version 11.
- Firefox 3.5 through 21.0
- Chrome 10

Review the following considerations and limitations when using Local App Access and URL redirection.

- Local App Access is designed for full-screen, virtual desktops spanning all monitors:
  - The user experience can be confusing if you use Local App Access with a virtual desktop that runs in windowed mode or does not cover all monitors.
  - Multiple monitors — When one monitor is maximized, it becomes the default desktop for all applications started in that session. This default occurs even if the subsequent applications typically start on another monitor.
  - The feature supports one VDA. There is no integration with multiple concurrent VDAs.
- Some applications can behave unexpectedly, affecting users:
  - The drive letters might confuse users, such as local C: rather than virtual desktop C: drive.
  - Available printers in the virtual desktop are not available to local applications.
  - Applications that require elevated permissions cannot be started as client-hosted applications.
  - There is no special handling for single-instance applications (such as Windows Media Player).
  - Local applications appear with the Windows theme of the local machine.
  - Full-screen applications are not supported. These applications include applications that open to a full screen, such as PowerPoint slide shows or photo viewers that cover the entire desktop.
  - Local App Access copies the properties of the local application (such as the shortcuts on the client's desktop and Start menu) on the VDA. However, it does not copy other properties such as shortcut keys and read-only attributes.
  - Applications that customize how overlapping window order is handled can have unpredictable results. For example, some windows might be hidden.
  - Shortcuts are not supported, including My Computer, Recycle Bin, Control Panel, Network Drive shortcuts, and folder shortcuts.
  - The following file types and files are not supported: custom file types, files with no associated programs, zip files, and hidden files.
  - Taskbar grouping is not supported for mixed 32-bit and 64-bit client-hosted or VDA applications. That is, grouping 32-bit local applications with 64-bit VDA applications.

- Applications cannot be started using COM. For example, if you click an embedded Office document from within an Office application, the process start cannot be detected, and the local application integration fails.
- Double-hop scenarios, where a user is starting a virtual desktop from within another virtual desktop session, are not supported.
- URL redirection supports only explicit URLs (that is, URLs appearing in the browser's address bar or found using the in-browser navigation, depending on the browser).
- URL redirection works only with desktop sessions, not with application sessions.
- The local desktop folder in a VDA session does not allow users to create files.
- Multiple instances of a locally running application behave according to the taskbar settings established for the virtual desktop. However, shortcuts to locally running applications are not grouped with running instances of those applications. They are also not grouped with running instances of hosted applications or pinned shortcuts to hosted applications. Users can close only windows of locally running applications from the Taskbar. Although users can pin local application windows to the desktop Taskbar and Start menu, the applications might not start consistently when using these shortcuts.
- If you set the **Allow local app access** policy setting to **Enabled**, browser content redirection isn't supported.

## Interaction with Windows

The Local App Access interaction with Windows includes the following behaviors.

- Windows 8 and Windows Server 2012 shortcut behavior
  - Windows Store applications installed on the client are not enumerated as part of Local App Access shortcuts.
  - Image and video files are opened by default using Windows store applications. However, Local App Access enumerates the Windows store applications and opens shortcuts with desktop applications.
- Local Programs
  - For Windows 7, the folder is available in the Start menu.
  - For Windows 8, Local Programs is available only when the user chooses **All Apps** as a category from the Start screen. Not all subfolders are displayed in Local Programs.
- Windows 8 graphics features for applications
  - Desktop applications are restricted to the desktop area and are covered by the Start screen and Windows 8 style applications.
  - Local App Access applications do not behave like desktop applications in multi-monitor mode. In multi-monitor mode, the Start screen and the desktop display on different monitors.
- Windows 8 and Local App Access URL Redirection

- Because Windows 8 Internet Explorer has no add-ons enabled, use desktop Internet Explorer to enable URL redirection.
- In Windows Server 2012, Internet Explorer disables add-ons by default. To implement URL Redirection, disable the Internet Explorer enhanced configuration. Then reset the Internet Explorer options and restart to ensure that add-ons are enabled for standard users.

## Configure Local App Access and URL redirection

To use Local App Access and URL redirection with Citrix Workspace app:

- Install Citrix Workspace app on the local client machine. You can enable both features during the Citrix Workspace app installation or you can enable Local App Access template using the Group Policy editor.
- Set the **Allow local app access** policy setting to **Enabled**. You can also configure URL allow list and block list policy settings for URL redirection. For more information, see [Local App Access policy settings](#).

## Enable Local App Access and URL redirection

To enable Local App Access for all local applications, follow these steps:

1. From **Manage > Full Configuration**, select **Policies** in the left pane.
2. Select **Create Policy** in the action bar.
3. In the Create Policy window, type “Allow Local App Access” in the search box and then click **Select**.
4. In the Edit Setting window, select **Allowed**. By default, the **Allow local app access** policy is prohibited. When this setting is allowed, the VDA allows the end-user to decide whether published applications and Local App Access shortcuts are enabled in the session. (When this setting is prohibited, both published applications and Local App Access shortcuts do not work for the VDA.) This policy setting applies to the entire machine and the URL redirection policy.
5. In the Create Policy window, type “URL redirection allow list” in the search box and then click **Select**. The URL redirection allow list specifies URLs to open in the default browser of the remote session.
6. In the Edit Setting window, click **Add** to add the URLs and then click **OK**.
7. In the Create Policy window, type “URL redirection block list” in the search box and then click **Select**. The URL redirection block list specifies URLs that are redirected to the default browser running on the endpoint.
8. In the Edit Setting window, click **Add** to add the URLs and then click **OK**.
9. On the Settings page, click **Next**.
10. On the Users and Machines page, assign the policy to the applicable Delivery Groups and then click **Next**.

11. On the Summary page, review the settings and then click **Finish**.

To enable URL redirection for all local applications during Citrix Workspace app installation, follow the steps below:

1. Enable URL redirection when you install Citrix Workspace app for all users on a machine. Doing so also registers the browser add-ons required for URL redirection.
2. From the command prompt, run the appropriate command to install the Citrix Workspace app using one of the following options:
  - For CitrixReceiver.exe, use `/ALLOW_CLIENHOSTEDAPPSURL=1`.
  - For CitrixReceiverWeb.exe, use `/ALLOW_CLIENHOSTEDAPPSURL=1`.

### Enable the Local App Access template using the Group Policy editor

#### Note:

- Before you enable the Local App Access template using the Group Policy editor, add the `receiver.admx/adml` template files to the local GPO. For more information, see [Configuring the Group Policy Object administrative template](#).
- Citrix Workspace app for Windows template files are available in the local GPO in **Administrative Templates > Citrix Components > Citrix Workspace** folder only when you add the `CitrixBase.admx/CitrixBase.adml` to the `%systemroot%\policyDefinitions` folder.

To enable the Local App Access template using the Group Policy editor, follow these steps:

1. Run **gpedit.msc**.
2. Go to **Computer Configuration > Administrative Templates > Classic Administrative Templates (ADM) > Citrix Components > Citrix Workspace > User Experience**.
3. Click **Local App Access settings**.
4. Select **Enabled** and then select **Allow URL Redirection**. For URL redirection, register browser add-ons using the command line described in the *Register browser add-ons* section further down in this article.

### Provide access only to published applications

You can provide access to published applications using the Registry Editor or the PowerShell SDK.

To the Registry Editor, see [The Local App Access for published applications](#) in the list of features managed through the registry.

To use the PowerShell SDK:

1. Open PowerShell on the machine where the Delivery Controller is running.
2. Enter the following command: `set-configsitemetadata -name "studio_clientHostedAppsEnabled" -value "true"`.



To have access to **Add Local App Access Application** in a Cloud service deployment, use the Citrix Virtual Apps and Desktops Remote PowerShell SDK. For more information, see [Citrix Virtual Apps and Desktops Remote PowerShell SDK](#).

1. Download the installer:

<https://download.apps.cloud.com/CitrixPoshSdk.exe>

2. Run these commands:

- a) `asnp citrix.*`
- b) `Get-XdAuthentication`

3. Enter the following command: `set-configsitemetadata -name "studio_clientHostedAppsEnabled" -value "true"`.

After you complete the applicable preceding steps, follow these steps to continue.

1. From **Manage > Full Configuration**, select **Applications** in the left pane.
2. In the upper middle pane, right-click the blank area and select **Add Local App Access Application** from the menu. You can also click **Add Local App Access Application** in the Actions pane. To display the Add Local App Access Application option in the Actions pane, click **Refresh**.
3. Publish Local App Access application.
  - The Local Application Access wizard launches with an Introduction page, which you can remove from future launches of the wizard.
  - The wizard guides you through the Groups, Location, Identification, Delivery, and Summary pages described below. When you are finished with each page, click **Next** until you reach the Summary page.
  - On the Groups page, select one or more Delivery Groups where the new applications will be added, and then click **Next**.
  - On the Location page, type the full executable path of the application on the user's local machine, and type the path to the folder where the application is located. Citrix recommends that you use the system environment variable path; for example, `%ProgramFiles(x86)%\Internet Explorer\iexplore.exe`.
  - On the Identification page, accept the default values or type the information that you want and then click **Next**.
  - On the Delivery page, configure how this application is delivered to users and then click **Next**. You can specify the icon for the selected application. You can also specify whether the shortcut to the local application on the virtual desktop will be visible on the Start menu, the desktop, or both.
  - On the Summary page, review the settings and then click **Finish** to exit the Local Application Access wizard.

## Register browser add-ons

### Note:

The browser add-ons required for URL redirection are registered automatically when you install Citrix Workspace app from the command line using the `/ALLOW_CLIENTHOSTEDAPPSURL=1` option.

You can use the following commands to register and unregister one or all add-ons:

- To register add-ons on a client device: `<client-installation-folder>\redirector.exe /reg<browser>`
- To unregister add-ons on a client device: `<client-installation-folder>\redirector.exe /unreg<browser>`
- To register add-ons on a VDA: `<VDInstallation-folder>\VDARedirector.exe /reg<browser>`
- To unregister add-ons on a VDA: `<VDInstallation-folder>\VDARedirector.exe /unreg<browser>`

Where `<browser>` is Internet Explorer, Firefox, Chrome, or All.

For example, the following command registers Internet Explorer add-ons on a device running Citrix Workspace app.

```
C:\Program Files\Citrix\ICA Client\redirector.exe/regIE
```

The following command registers all add-ons on a Windows Multi-session OS VDA.

```
C:\Program Files (x86)\Citrix\HDX\bin\VDARedirector.exe /regAll
```

## URL interception across browsers

- By default, Internet Explorer redirects the specified URL. If the URL is not in the block list but the browser or website redirects it to another URL, the final URL is not redirected. It is not redirected even if it is on the block list.

For URL redirection to work correctly, enable the add-on when prompted by the browser. If the add-ons that are using Internet options or the add-ons in the prompt are disabled, URL redirection does not work correctly.

- The Firefox add-ons always redirect the URLs.

When an add-on is installed, Firefox prompts to allow or prevent installing the add-on on a new tab page. Allow the add-on for the feature to work.

- The Chrome add-on always redirects the final URL that is navigated, and not the entered URLs.

The extensions have been installed externally. When you disable the extension, the URL redirection feature does not work in Chrome. If the URL redirection is required in Incognito mode, allow the extension to run in that mode in the browser settings.

## Configure local application behavior on logoff and disconnect

### Note:

If you do not follow these steps to configure the settings, by default, local applications continue to run when a user logs off or disconnects from the virtual desktop. After reconnection, local applications are reintegrated if they are available on the virtual desktop.

To configure local application behavior on logoff and disconnect, see [Local application behavior on logoff and disconnect](#) in the list of features managed through the registry.

## Generic USB redirection and client drive considerations

August 4, 2021

HDX technology provides **optimized support** for most popular USB devices. Optimized support offers an improved user experience with better performance and bandwidth efficiency over a WAN. Optimized support is usually the best option, especially in high latency or security-sensitive environments.

HDX technology provides **generic USB redirection** for specialty devices that don't have optimized support or where it is unsuitable, for example:

- The USB device has more advanced features that are not part of optimized support, such as a mouse or webcam having more buttons.
- Users need functions which are not part of optimized support.
- The USB device is a specialized device, such as test and measurement equipment or an industrial controller.
- An application requires direct access to the device as a USB device.
- The USB device only has a Windows driver available. For example, a smart card reader might not have a driver available for Citrix Workspace app for Android.
- The version of Citrix Workspace app does not provide any optimized support for this type of USB device.

With generic USB redirection:

- Users do not need to install device drivers on the user device.
- USB client drivers are installed on the VDA machine.

### Important:

- Generic USB redirection can be used together with optimized support. If you enable generic USB redirection, configure Citrix [USB devices policy settings](#) for both generic USB redirection and optimized support.
- The Citrix policy setting in [Client USB device optimization rules](#) is a specific setting for

generic USB redirection, for a particular USB device. It doesn't apply to optimized support as described here.

- When brokering a session using Citrix software to an Azure Virtual Machine, Citrix provides best effort support for USB redirection to the Azure Virtual Machine. We support fixing a Citrix software problem, but we do not support the underlying Azure Virtual Machine.
- CD/DVD devices with disc burning capabilities can be redirected, but the burning capabilities of these devices cannot be used. This is due to the buffer limits of a session.

### **Performance considerations for USB devices**

Network latency and bandwidth can affect user experience and USB device operation when using generic USB redirection for some types of USB devices. For example, timing-sensitive devices might not operate correctly over high-latency low-bandwidth links. Use optimized support instead where possible.

Some USB devices require high bandwidth to be usable, for example a 3D mouse (used with 3D apps that also typically require high bandwidth). If bandwidth cannot be increased, you might be able to mitigate the issue by tuning bandwidth usage of other components using the bandwidth policy settings. For more information, see [Bandwidth policy settings](#) for Client USB device redirection, and [Multi-stream connection policy settings](#).

### **Security considerations for USB devices**

Some USB devices are security-sensitive by nature, for example, smart card readers, fingerprint readers, and signature pads. Other USB devices such as USB storage devices can be used to transmit data that might be sensitive.

USB devices are often used to distribute malware. Configuration of Citrix Workspace app and Citrix Virtual Apps and Desktops can reduce, but not eliminate, risk from these USB devices. This situation applies whether generic USB redirection or optimized support is used.

#### **Important:**

For security-sensitive devices and data, always secure the HDX connection using either [TLS](#) or IPsec.

Only enable support for the USB devices that you need. Configure both generic USB redirection and optimized support to meet this need.

Provide guidance to users for safe use of USB devices:

- Use only USB devices that have been obtained from a trustworthy source.
- Don't leave USB devices unattended in open environments - for example, a flash drive in an internet cafe.

- Explain the risks of using a USB device on more than one computer.

## Compatibility with generic USB redirection

Generic USB redirection is supported for USB 2.0 and earlier devices. Generic USB redirection is also supported for USB 3.0 devices connected to a USB 2.0 or USB 3.0 port. Generic USB redirection does not support USB features introduced in USB 3.0, such as super speed.

These Citrix Workspace apps support generic USB redirection:

- Citrix Workspace app for Windows, see [Configuring application delivery](#).
- Citrix Workspace app for Mac, see [Citrix Workspace app for Mac](#).
- Citrix Workspace app for Linux, see [Optimize](#).
- Citrix Workspace app for Chrome OS, see [Citrix Workspace app for Chrome](#).

For Citrix Workspace app versions, see the [Citrix Workspace app feature matrix](#).

If you are using earlier versions of Citrix Workspace app, see the Citrix Workspace app documentation to confirm that generic USB redirection is supported. See Citrix Workspace app documentation for any restrictions on USB device types that are supported.

Generic USB redirection is supported for desktop sessions from VDA for Single-session OS version 7.6 through current.

Generic USB redirection is supported for desktop sessions from VDA for Multi-session OS version 7.6 through current, with these restrictions:

- The VDA must be running Windows Server 2012 R2 or Windows Server 2016.
- The USB device drivers must be fully compatible with Remote Desktop Session Host (RDSH) for the VDA OS (Windows 2012 R2), including full virtualization support.

Some types of USB devices are not supported for generic USB redirection because it would not be useful to redirect them:

- USB modems.
- USB network adapters.
- USB hubs. The USB devices connected to USB hubs are handled individually.
- USB virtual COM ports. Use COM port redirection rather than generic USB Redirection.

For information on USB devices that have been tested with generic USB redirection, see [Citrix Ready Marketplace](#). Some USB devices do not operate correctly with generic USB redirection.

## Configure generic USB redirection

You can control, and separately configure, which types of USB devices use generic USB redirection:

- On the VDA, using Citrix policy settings. For more information, see [Redirection of client drives and user devices](#) and [USB devices policy settings](#) in the Policy settings reference
- In Citrix Workspace app, using Citrix Workspace app-dependent mechanisms. For example, an Administrative Template controls registry settings that configure Citrix Workspace app for Windows. By default, USB redirection is allowed for certain classes of USB devices and denied for others. For more information, see [Configure](#) in the Citrix Workspace app for Windows documentation.

This separate configuration provides flexibility. For example:


- If two different organizations or departments are responsible for Citrix Workspace app and VDA, they can enforce control separately. This configuration applies when a user in one organization accesses an application in another organization.
- Citrix policy settings can control USB devices that are allowed only for certain users or for users connecting only over a LAN (rather than by using Citrix Gateway).

### **Enable generic USB redirection**

To enable generic USB Redirection, and not require manual redirection by the user, configure both Citrix policy settings and Citrix Workspace app connections preferences.

In Citrix policy settings:


1. Add the [Client USB device redirection](#) to a policy and set its value to **Allowed**.



2. (Optional) To update the list of USB devices available for redirection, add the [Client USB device redirection rules](#) setting to a policy and specify the USB policy rules.

In Citrix Workspace app:

3. Specify that devices are connected automatically without manual redirection. You can do this using an Administrative template or in Citrix Workspace app for Windows > Preferences > Connections.



If you specified USB policy rules for the VDA in the previous step, specify those same policy rules for Citrix Workspace app.


For thin clients, consult the manufacturer for details of USB support and any required configuration.

### Configuring the types of USB devices available for generic USB redirection

USB devices are automatically redirected when USB support is enabled and the USB user preference settings are set to connect USB devices automatically. USB devices are also automatically redirected when the connection bar is not present.

Users can explicitly redirect devices that are not automatically redirected by selecting the devices from the USB device list. For more information, the Citrix Workspace app for Windows user help article, [Display your devices in the Desktop Viewer](#).





To use generic USB redirection rather than optimized support, you can either:

- In Citrix Workspace app, manually select the USB device to use generic USB redirection, choose **Switch to generic** from the Devices tab of the Preferences dialog box.
- Automatically select the USB device to use generic USB redirection, by configuring auto-redirection for the USB device type (for example, `AutoRedirectStorage=1`) and set USB user preference settings to automatically connect USB devices. For more information, see [Configure automatic redirection of USB devices](#).

**Note:**

Only configure generic USB redirection for use with a webcam if the webcam is found to be incompatible with HDX multimedia redirection.

To prevent USB devices from ever being listed or redirected, you can specify device rules for Citrix Workspace app and the VDA.

For generic USB redirection, you need to know at least the USB device class and subclass. Not all USB devices use their obvious USB device class and subclass. For example:

- Pens use the mouse device class.
- Smart card readers can use the vendor-defined or HID device class.

For more precise control, you need to know the Vendor ID, Product ID, and Release ID. You can get this information from the device vendor.

**Important:**

Malicious USB devices might present USB device characteristics that do not match their intended usage. Device rules are not intended to prevent this behavior.

You control the USB devices available for generic USB redirection by specifying USB device redirection rules for both VDA and Citrix Workspace app, to override the default USB policy rules.

For the VDA:

- Edit the administrator override rules for the Multi-session OS machines through group policy rules. The Group Policy Management Console is included on the installation media:
  - For x64: `dvd root \os\lang\x64\Citrix Policy\CitrixGroupPolicyManagement_x64.msi`
  - For x86: `dvd root \os\lang\x86\Citrix Policy\CitrixGroupPolicyManagement_x86.msi`

At Citrix Workspace app for Windows:

- Edit the user device registry. An Administrative template (ADM file) is included on the installation media so you can change the user device through Active Directory Group Policy:  
`dvd root \os\lang\Support\Configuration\icaclient_usb.adm`

**Warning:**

Editing the registry incorrectly can cause serious problems that might require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

The product default rules are stored in `HKLM\SOFTWARE\Citrix\PortICA\GenericUSB\DeviceRules`. Do not edit these product default rules. Instead, use them as a guide for creating administrator override rules, which is explained later in this article. The GPO overrides are evaluated before the product default rules.

The administrator override rules are stored in `HKLM\SOFTWARE\Policies\Citrix\PortICA\GenericUSB\DeviceRules`. GPO policy rules take the format **{Allow: | Deny:}** followed by a set of `tag=value` expressions separated by white space.

The following tags are supported:

Tag	Description
VID	Vendor ID from the device descriptor
PID	Product ID from the device descriptor
REL	Release ID from the device descriptor
Class	Class from either the device descriptor or an interface descriptor; see the USB website at <a href="http://www.usb.org/">http://www.usb.org/</a> for available USB Class Codes
SubClass	Subclass from either the device descriptor or an interface descriptor
Prot	Protocol from either the device descriptor or an interface descriptor

When creating policy rules, note the following:

- Rules are case-insensitive.
- Rules can have an optional comment at the end, introduced by `##`. A delimiter is not required, and the comment is ignored for matching purposes.
- Blank and pure comment lines are ignored.
- White space is used as a separator, but cannot appear in the middle of a number or identifier. For example, `Deny: Class=08 SubClass=05` is a valid rule, but `Deny: Class=0 Sub Class=05` is not.
- Tags must use the matching operator `=`. For example, `VID=1230`.
- Each rule must start on a new line or form part of a semicolon-separated list.

**Note:**

If you are using the ADM template file, you must create rules on a single line, as a semicolon-separated list.

Examples:

- The following example shows an administrator-defined USB policy rule for vendor and product identifiers:

```
Allow: VID=046D PID=C626 ## Allow Logitech SpaceNavigator 3D Mouse Deny
: VID=046D ## Deny all Logitech products
```

- The following example shows an administrator-defined USB policy rule for a defined class, subclass, and protocol:

```
Deny: Class=EF SubClass=01 Prot=01 ## Deny MS Active Sync devices Allow
: Class=EF SubClass=01 ## Allow Sync devices Allow: Class=EF ## Allow
```

all USB-Miscellaneous devices

## Use and remove USB devices

Users can connect a USB device before or after starting a virtual session.

When using Citrix Workspace app for Windows, the following apply:

- Devices connected after a session begins appear immediately in the USB menu of the Desktop Viewer.
- If a USB device is not redirecting properly, you can try to resolve the problem by waiting to connect the device until after the virtual session starts.
- To avoid data loss, use the Windows “Safely Remove Hardware” icon before removing the USB device.

## Security controls for USB mass storage devices

Optimized support is provided for USB mass storage devices. This support is part of Citrix Virtual Apps and Desktops client drive mapping. Drives on the user device are automatically mapped to drive letters on the virtual desktop when users log on. The drives are displayed as shared folders that have mapped drive letters. To configure client drive mapping, use the **Client removable drives** setting. This setting is in the [File Redirection policy settings](#) section of the ICA policy settings.

With USB mass storage devices you can use either Client drive mapping or generic USB redirection, or both. Control them using Citrix policies. The main differences are:

Feature	Client drive mapping	Generic USB redirection
Enabled by default	Yes	No
Read-only access configurable	Yes	No
Encrypted device access	Yes, if encryption is unlocked before the device is accessed	Yes
BitLocker To Go devices	No	No
Safe to delete device during a session	No	Yes, provided users follow operating system recommendations for safe removal

If both generic USB redirection and the client drive mapping policies are enabled and a mass storage

device is inserted either before or after a session starts, it is redirected using client drive mapping. When both generic USB redirection and the client drive mapping policies are enabled and a device is configured for automatic redirection and a mass storage device is inserted either before or after a session starts, it is redirected using generic USB redirection. For more information, see Knowledge Center article [CTX123015](#).

**Note:**

USB redirection is supported over lower bandwidth connections, for example 50 Kbps. However, copying large files doesn't work.

## Control file access with client drive mapping

You can control whether users can copy files from their virtual environments to their user devices. By default, files and folders on mapped client-drives are available in read/write mode from within the session.

To prevent users from adding or changing files and folders on mapped client-devices, enable the **Read-only client drive access** policy setting. When adding this setting to a policy, ensure that the Client drive redirection setting is set to **Allowed** and is also added to the policy.

## Policies

August 29, 2018

Policies are a collection of settings that define how sessions, bandwidth, and security are managed for a group of users, devices, or connection types.

You can apply policy settings to physical and virtual machines or to users. You can apply settings to individual users at the local level or in security groups in Active Directory. The configurations define specific criteria and rules, and if you do not specifically assign the policies, the settings are applied to all connections.

For complete information about Citrix policies, begin with [Policies](#). From that article, you can move on to:

- [Work with policies](#)
- [Policy templates](#)
- [Create policies](#)
- [Compare, prioritize, model, and troubleshoot policies](#)
- [Default policy settings](#)
- [Policy settings reference](#)

## Manage

August 19, 2021

Citrix manages Citrix Virtual Apps and Desktops service deployments by installing and maintaining the core components and features in Citrix Cloud.

You take care of the machines (VDAs) in resource locations that deliver apps and desktops. You also manage connections to those resource locations, plus the apps, desktops, and users.

- **Autoscale:** A consistent, high-performance solution to proactively power manage your machines.
- **Applications:** Manage applications in delivery groups.
- **Virtual IP and virtual loopback:** The Microsoft virtual IP address feature provides a published application with a unique dynamically assigned IP address for each session. With Citrix virtual loopback, you can configure applications that depend on communications with localhost (127.0.0.1 by default) to use a unique virtual loopback address in the localhost range (127.\*).
- **VDA registration:** Before a VDA can facilitate delivery of apps and desktops, it must register (establish communication) with a Cloud Connector. You can specify Cloud Connector addresses using several methods, which are described in this article. As you add Cloud Connectors, VDAs must have current information.
- **Sessions:** Maintaining session activity is critical to providing the best user experience. Several features can optimize the reliability of sessions, reduce inconvenience, downtime, and loss of productivity.
- **Using Search:** To view information about machines, sessions, machine catalogs, applications, or delivery groups in the Full Configuration management interface, use the flexible search feature.
- **IPv4/IPv6 support:** Citrix Virtual Apps and Desktops supports pure IPv4, pure IPv6, and dual-stack deployments that use overlapping IPv4 and IPv6 networks. This article describes and illustrates these deployments. It also describes the Citrix policy settings that control the use of IPv4 or IPv6.
- **Profile management:** Citrix Profile Management can be installed when you install a VDA. If you use this user profile solution, review its documentation.
- **Citrix Insight Services:** Citrix Insight Services (CIS) is a Citrix platform for instrumentation, telemetry, and business insight generation. Analytics and diagnostics are collected when you install a VDA.
- **Local Host Cache:** Local Host Cache enables connection brokering operations to continue when a Cloud Connector in a resource location cannot communicate with Citrix Cloud. [Scale](#),

[size, and other configuration considerations](#) are also provided.

- **Delegated administration:** With delegated administration, you can configure the access permissions that all of your administrators need, according to their role in your organization.
- **Configuration logging:** Configuration logging tracks configuration changes and administrative activities.
- **Event logs:** Services within Citrix Virtual Apps and Desktops log the events that occur. Event logs can be used to monitor and troubleshoot operations.
- **Licenses:** You can view Citrix license usage information for this service from the Citrix Cloud console.
- **Load balance machines:** You can control how to load balance machines.

## Autoscale

August 31, 2021

Autoscale is a feature exclusive to Citrix Virtual Apps and Desktops service that provides a consistent, high-performance solution to proactively power manage your machines. It aims to balance costs and user experience. Autoscale incorporates the deprecated Smart Scale technology into the **Manage** console's power management solution.

Autoscale enables proactive power management of all registered single-session and multi-session OS machines in a delivery group.

### Supported VDA hosting platforms

Autoscale supports all the platforms that Virtual Apps and Desktops service supports. This includes various infrastructure platforms including Citrix Hypervisor, Amazon Web Services, Google Cloud Platform, Microsoft Azure Resource Manager, VMware vSphere, and many more. For a complete list of supported platforms, see [System requirements](#) for Citrix Virtual Apps and Desktops Service.

### Supported workloads

Autoscale supports both multi-session OS and single-session OS delivery groups. There are three user interfaces to be aware of:

- Autoscale user interface for multi-session OS delivery groups (formerly RDS delivery groups)
- Autoscale user interface for single-session OS random (pooled) delivery groups (formerly pooled VDI delivery groups)

- Autoscale user interface for single-session OS static delivery groups (formerly static VDI delivery groups)

For more information about the user interfaces for different delivery groups, see [Autoscale user interfaces](#).

## Benefits

The Autoscale feature delivers the following benefits:

- Provide you with a single, consistent mechanism to power manage machines in a delivery group.
- Ensure availability and control costs by powering machines with load-based or schedule-based power management, or a combination of both.
- To monitor metrics such as cost savings and capacity utilization, and to enable notifications, use [Director](#), available on the **Monitor** tab.

## Watch a 2-minute video

The following video provides a quick tour of Autoscale.

[This is an embedded video. Click the link to watch the video](#)

## Autoscale features

Autoscale features include:

- [Schedule-based and load-based scaling settings](#)
- [Autoscale restriction](#)
- [Dynamic machine provisioning](#)
- [Force user logoff](#)

## Autoscale user interfaces

There are three types of Autoscale user interfaces to be aware of.

Autoscale user interface for *single-session OS static delivery groups*:



**Manage Autoscale** Enabled

General

Schedule and Peak Ti...

Load-based Settings

ADVANCED

Restrict Autoscale

### Getting Started with Autoscale

Autoscale helps you deliver cloud use cases and save costs by optimizing resource utilization. It provides various basic (schedule-based and load-based) and advanced features for your use. You can enable one or more. To monitor metrics such as cost savings and capacity utilization, and to enable notifications, use Director, available on the Monitor tab. [Learn more](#)

Enable Autoscale ?

**Power-off delay** ?

Delay powering off machines by:  minutes

**Machine cost** ?

Visualize cost savings achieved by Autoscale on the Monitor tab.

Machine instance cost per hour (\$):

Save Cancel Apply

Autoscale user interface for *single-session OS random delivery groups*:

The screenshot shows the 'Manage Autoscale' configuration page. At the top, the title 'Manage Autoscale' is followed by a status indicator 'Enabled'. A left-hand navigation pane lists several categories: 'General', 'Schedule and Peak Ti...', 'Load-based Settings', 'ADVANCED', and 'Restrict Autoscale'. The main content area is titled 'Getting Started with Autoscale' and contains the following text: 'Autoscale helps you deliver cloud use cases and save costs by optimizing resource utilization. It provides various basic (schedule-based and load-based) and advanced features for your use. You can enable one or more. To monitor metrics such as cost savings and capacity utilization, and to enable notifications, use Director, available on the Monitor tab. [Learn more](#)'. Below this text are three settings: 1. 'Enable Autoscale' with a checked checkbox and a help icon. 2. 'Power-off delay' with a help icon, followed by the text 'Delay powering off machines by:' and a text input field containing '30' followed by the word 'minutes'. 3. 'Machine cost' with a help icon, followed by the text 'Visualize cost savings achieved by Autoscale on the Monitor tab.' and 'Machine instance cost per hour (\$):' followed by a text input field containing '1'. At the bottom right of the page are three buttons: 'Save', 'Cancel', and 'Apply'.

Autoscale user interface for *multi-session OS delivery groups*:

## Manage Autoscale Disabled

- General
- Schedule and Peak Ti...
- Load-based Settings
- ADVANCED
- Dynamic Session Tim...
- Force User Logoff
- Restrict Autoscale

### Getting Started with Autoscale

Autoscale helps you deliver cloud use cases and save costs by optimizing resource utilization. It provides various basic (schedule-based and load-based) and advanced features for your use. You can enable one or more. To monitor metrics such as cost savings and capacity utilization, and to enable notifications, use Director, available on the Monitor tab. [Learn more](#)

Enable Autoscale ?

Power-off delay ?

Delay powering off machines by:  minutes

**Machine cost** ?

Visualize cost savings achieved by Autoscale on the Monitor tab.

Machine instance cost per hour (\$):

Save Cancel Apply

## Enable or disable Autoscale for a delivery group

### Note:

In the *legacy console*, Autoscale is enabled by default when you create a delivery group.  
In the *web-based console*, Autoscale is disabled by default when you create a delivery group.

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane.
2. Select the delivery group you want to manage and then click **Manage Autoscale**.

Delivery Group	Delivering	Machine Count	Session in Use
Basic Desktop Multi-session OS   Managed by Citrix Cloud	Applications and Desktops	Total: 1 Unregistered: 1	Total: 0 Disconnected: 0
Education Team Multi-session OS   Managed by Citrix Cloud	Applications and Desktops	Total: 1 Unregistered: 1	Total: 2 Disconnected: 2
Facilities Team Multi-session OS   Managed by Citrix Cloud	Applications and Desktops	Total: 1 Unregistered: 1	Total: 0 Disconnected: 0
Faculty Team Multi-session OS   Managed by Citrix Cloud	Applications and Desktops	Total: 1 Unregistered: 1	Total: 1 Disconnected: 1
Office Productivity Multi-session OS   Managed by Citrix Cloud	Applications and Desktops	Total: 1 Unregistered: 1	Total: 1 Disconnected: 1
Students Multi-session OS   Managed by Citrix Cloud	Applications and Desktops	Total: 1 Unregistered: 1	Total: 21 Disconnected: 21
Windows 10 Dedicated Single-session OS   Managed by Citrix Cloud	Desktops (Static machine assignment)	Total: 5 Unregistered: 5	Total: 0 Disconnected: 0

- On the **Manage Autoscale** page, select the **Enable Autoscale** check box to enable Autoscale. After you enable Autoscale, the options on the page are enabled.

**Manage Autoscale** Enabled

**Getting Started with Autoscale**  
Autoscale helps you deliver cloud use cases and save costs by optimizing resource utilization. It provides various basic (schedule-based and load-based) and advanced features for your use. You can enable one or more. To monitor metrics such as cost savings and capacity utilization, and to enable notifications, use Director, available on the Monitor tab. [Learn more](#)

**Enable Autoscale** ?

**Power-off delay** ?  
Delay powering off machines by:  minutes

**Machine cost** ?  
Visualize cost savings achieved by Autoscale on the Monitor tab.  
Machine instance cost per hour (\$):

Save Cancel Apply

- To disable Autoscale, clear the **Autoscale** check box. The options on the page turn gray to indicate that Autoscale is disabled for the selected delivery group.

**Important:**

- If you disable Autoscale, all machines managed by Autoscale remain in the state they are in at the time of disabling.

- After you disable Autoscale, the machines in drain state are taken out of drain state. For more information about drain state, see [Drain state](#).

## Monitoring metrics

You can monitor the following metrics of Autoscale-managed machines from the **Monitor** tab.

- Machine usage
- Estimated savings
- Alert notifications for machines and sessions
- Machine status
- Load evaluation trends

For more information about the metrics, see [Monitor Autoscale-managed machines](#).

## Broker PowerShell SDK commands

You can configure Autoscale for delivery groups using the Broker PowerShell SDK. To configure Autoscale using PowerShell commands, you must use Remote PowerShell SDK version 7.21.0.12 or later. For more information about the Remote PowerShell SDK, see [SDKs and APIs](#).

### Set-BrokerDesktopGroup

Disables or enables an existing BrokerDesktopGroup or alters its settings. For more information about this cmdlet, see <https://citrix.github.io/delivery-controller-sdk/Broker/Set-BrokerDesktopGroup/>.

### New-BrokerPowerTimeScheme

Creates a BrokerPowerTimeScheme for a delivery group. For more information about this cmdlet, see <https://citrix.github.io/delivery-controller-sdk/Broker/New-BrokerPowerTimeScheme/>.

## Examples

See the following examples for details about how to use the PowerShell cmdlets.

Enable Autoscale

- Suppose you want to enable Autoscale for the delivery group whose name is “MyDesktop.” Use the `Set-BrokerDesktopGroup` PowerShell command. For example:
  - `PS C:\> Set-BrokerDesktopGroup "MyDesktop"-AutoscalingEnabled $true`

Configure the capacity buffer separately for peak and off-peak times

- Suppose you want to set the capacity buffer to 20% for peak times and 10% for off-peak times for a delivery group whose name is “MyDesktop.” Use the `Set-BrokerDesktopGroup` PowerShell command. For example:

```
- PS C:\> Set-BrokerDesktopGroup "MyDesktop"-PeakBufferSizePercent 20  
-OffPeakBufferSizePercent 10
```

#### Configure the when disconnected timeout

- Suppose you want to set the when disconnected timeout to 60 minutes for peak times and 30 minutes for off-peak times for a delivery group whose name is “MyDesktop.” Use the `Set-BrokerDesktopGroup` PowerShell command. For example:

```
- PS C:\> Set-BrokerDesktopGroup "MyDesktop"-PeakDisconnectTimeout 60  
-OffPeakDisconnectTimeout 30
```

#### Configure the when logged off timeout

- Suppose you want to set the when logged off timeout to 60 minutes for peak times and 30 minutes for off-peak times for a delivery group whose name is “MyDesktop.” Use the `Set-BrokerDesktopGroup` PowerShell command. For example:

```
- PS C:\> Set-BrokerDesktopGroup "MyDesktop"-PeakLogOffTimeout 60 -  
OffPeakLogOffTimeout 30
```

#### Configure the power-off delay

- Suppose you want to set the power-off delay to 15 minutes for a delivery group whose name is “MyDesktop.” Use the `Set-BrokerDesktopGroup` PowerShell command. For example:

```
- PS C:\> Set-BrokerDesktopGroup "MyDesktop"-PowerOffDelay 15
```

#### Configure a time period during which the power-off delay does not take effect

- Suppose you want the power-off delay to take effect until 30 minutes have elapsed for a delivery group whose name is “MyDesktop.” Use the `Set-BrokerDesktopGroup` PowerShell command. For example:

```
- C:\PS> Set-BrokerDesktopGroup "MyDesktop"-SettlementPeriodBeforeAutoShutdown  
30.
```

#### Configure the machine instance cost

- Suppose you want to set the machine instance cost per hour to 0.2 dollars for a delivery group whose name is “MyDesktop.” Use the `Set-BrokerDesktopGroup` PowerShell command. For example:

```
- PS C:\> Set-BrokerDesktopGroup "MyDesktop"-MachineCost 0.2
```

#### Create a power time scheme

- Suppose you want to create a power time scheme for a delivery group whose UID value is 3. The new scheme covers weekend, Monday, and Tuesday. The 8:00 AM to 6:30 PM time slot is

defined as peak times for the days included in the scheme. For peak times, the pool size (the number of machines kept powered on) is 20. For off-peak times, it is 5. You can use the `Set-BrokerDesktopGroup` PowerShell command. For example:

```
- PS C:\> $ps48=(0..47 | %{ if ($_ -lt 16 -or $_ -gt 37){ 5 } else { 20 } } )
- PS C:\> $pt48=(0..47 | %{ if ($_ -lt 16 -or $_ -gt 37){ $false } else { $true } } )
- PS C:\> New-BrokerPowerTimeScheme -Name 'First Half Week'-DaysOfWeek Weekend,Monday,Tuesday -DesktopGroupUid 3 -PeakHalfHours $pt48 - PoolSize $ps48
```

## Drain state

Autoscale always attempts to scale down the number of powered-on machines in the delivery group to the configured pool size and capacity buffer. It does so by putting the excess machines with the fewest sessions into “drain state” and powering them off when all sessions are logged off. This occurs when session demand lessens and the schedule requires fewer machines than are powered on.

Autoscale puts excess machines into “drain state” one by one. If two or more machines have the same number of active sessions, Autoscale drains the machine that has been powered on for the specified power-off delay. Doing so avoids putting recently powered-on machines into drain state because those machines are more likely to have the fewest sessions. If two or more machines have been powered on for the specified power-off delay, Autoscale drains those machines one by one at random.

Machines in drain state no longer host new session launches and are waiting for the existing sessions to be logged off. A machine becomes a candidate for shutdown only when all sessions are logged off. However, if there are no machines immediately available for session launches, Autoscale prefers directing the session launches to a machine in drain state over powering on a machine.

A machine is taken out of drain state when one of the following conditions is met:

- The machine is powered off.
- Autoscale is disabled for the delivery group to which the machine belongs.
- Autoscale utilizes the machine to meet schedule or load demand requirements. This case occurs when the schedule (schedule-based scaling) or the current demand (load-based scaling) requires more machines than the number of machines that are currently powered on.

### Important:

If no machines are immediately available for session launches, Autoscale prefers directing session launches to a machine in drain state over powering on a machine. A machine in drain state that hosts a session launch remains in drain state.

To find out which machines are in drain state, use the `Get-BrokerMachine` PowerShell command.

For example: `Get-BrokerMachine -DrainingUntilShutdown $true`. Alternatively, you can use the Manage console. See Display machines in drain state.

## Display machines in drain state

### Note:

This feature is available only in the web-based console and applies only to multi-session machines.

In **Manage > Full Configuration**, you can display machines that are in drain state, letting you know which machines are about to shut down. Complete the following steps:

1. Navigate to the **Search** node and then click **Columns to Display**.
2. In the **Columns to Display** window, select the check box next to **Drain State**.
3. Click Save to exit the **Columns to Display** window.

The **Drain State** column can display the following information:

- **Draining until shutdown.** Appears when machines are in drain state until they are shut down.
- **Not draining.** Appears when machines are not yet in drain state.

Name ↓	Machine Catalog	Delivery Group	Maintenance Mode	User Change Per...	Power State	Registration State	Sessio...	Drain State
318zjh001.xd.local	zjh-mul	zjh-mul	Off	Discard	On	Registered	-	Draining until shutdown
318zjh002.xd.local	zjh-mul	zjh-mul	Off	Discard	On	Registered	1	Not draining
318zjh003.xd.local	zjh-mul	zjh-mul	Off	Discard	On	Registered	1	Not draining

## Load index

### Important:

Load index applies only to multi-session delivery groups.

The load index value ranges from 0 to 10,000, which is calculated using the Citrix Load Management policy settings configured for concurrent logon, session, CPU, disk, and memory use. The digit “0” indicates an unloaded machine. A machine with a load index value of 0 is at a baseline load. The digit



“10,000” indicates a fully loaded machine that cannot run any more sessions. The load index metric determines how likely a machine is to receive connections. By default, a machine is considered at full load when it is hosting 250 sessions.

### **Good to know**

Autoscale works at a delivery group level. It is configured on a per-delivery group basis. It power manages only the machines in the selected delivery group.

### **Capacity and machine registration**

To ensure that Autoscale has an accurate view of machines that can accept session requests, Autoscale includes only machines that are registered with the site when determining the capacity for a given delivery group. Powered-on machines that are unregistered cannot accept session requests. As a result, they are not included in the overall capacity of the delivery group.

### **Scaling across multiple machine catalogs**

In some sites, multiple machine catalogs might be associated with a single delivery group. Autoscale randomly powers on machines from each catalog to meet schedule or session demand requirements.

For example, a delivery group has two machine catalogs: Catalog A has three machines powered on and Catalog B has one machine powered on. If Autoscale needs to power on an extra machine, it might power on a machine from either Catalog A or Catalog B.

### **Machine provisioning and session demand**

The machine catalog associated with the delivery group must have enough machines to power on and off as demand increases and decreases. If session demand exceeds the total number of registered machines in the delivery group, Autoscale ensures that all registered machines are powered on.

**Autoscale does not provision additional machines.** To overcome this bottleneck, you can use a PowerShell script to create machines and delete them dynamically.

### **Availability of monitoring data**

Monitoring data is available when Autoscale is enabled for the delivery group. Monitoring data remains available if Autoscale is enabled and then disabled for the delivery group. Autoscale collects monitoring data at 5-minute intervals.

**Note:**

When you initially enable Autoscale for a delivery group, it might take a few minutes to display monitoring data for that delivery group.

### Instance size considerations

You can optimize your costs if you right size your instances in public clouds. Smaller instances host fewer user sessions than larger instances. Therefore, in the case of smaller instances, Autoscale puts machines into drain state much faster because it takes less time for the last user session to be logged off. As a result, Autoscale powers off smaller instances sooner, thereby reducing costs. We recommend that you provision smaller instances as long as they match your workload performance and capacity requirements.

### More information

For more information on Autoscale, see [Citrix Autoscale](#) in Tech Zone.

## Schedule-based and load-based settings

August 11, 2021

### How Autoscale power manages machines

Autoscale powers machines on and off based on the selected schedule. Autoscale lets you set multiple schedules that include specific days of the week and adjust the number of machines available during those times. If you expect a set of users to consume the machine resources at a specific time on specific days, Autoscale helps provide an optimized experience. Note that those machines will be powered on during the schedule, whether or not there are sessions running on them.

**Note:**

Autoscale does not support power managing manually provisioned machines and machines in dedicated machine catalogs.

The schedule is based on the **time zone** of the delivery group. To change the time zone, you can change user settings in a delivery group. For more information, see [Manage Delivery Groups](#).

Autoscale has two default schedules: *Weekdays* (Monday through Friday) and *Weekend* (Saturday and Sunday). By default, the **Weekdays** schedule keeps one machine powered on from 07:00 AM to 06:30

PM during peak times and none during off-peak times. The default capacity buffer is set to 10% during peak and off-peak times. By default, the **Weekend** schedule keeps no machines powered on.

**Note:**

Autoscale treats only those machines that are registered with the site as part of the available capacity in the calculations it makes. “Registered” means that the machine is available for use or already in use. Doing so ensures that only machines that can accept user sessions are included in the capacity for the delivery group.

**User interfaces**

There are three types of user interfaces to be aware of.

User interface for *single-session OS static delivery groups*:

The screenshot shows the 'Manage Autoscale' interface with the 'Enabled' status. The left sidebar contains navigation options: 'General', 'Schedule and Peak Ti...', 'Load-based Settings', 'ADVANCED', and 'Restrict Autoscale'. The main content area is titled 'Schedule and Peak Times' and includes a descriptive paragraph and a 'Set schedules' button. Below this is a 'Weekdays' section with a timeline showing peak times from 9:00 AM to 6:00 PM on Wednesday, Thursday, Friday, and Saturday. A 'Weekend' section is also visible below the Weekdays section. At the bottom right, there are 'Save', 'Cancel', and 'Apply' buttons.

## Manage Autoscale Enabled

---

General

Schedule and Peak Ti...

Load-based Settings

ADVANCED

Restrict Autoscale

### Load-based Settings

Capacity buffer is used to accommodate sudden spikes in demand by keeping a buffer of machines powered on based on the maximum load index of the delivery group. By default, the capacity buffer is 10% of the maximum supported load index of the delivery group. A lesser value decreases the cost. A greater value ensures an optimized user experience so that users do not have to wait for additional machines to power on when launching sessions. [Learn more](#)

	During peak times	During off-peak times
Capacity buffer (%):	<input type="text" value="10"/>	<input type="text" value="10"/>
When disconnected (minutes):	<input type="text" value="0"/> <input style="border: none; border-bottom: 1px solid #ccc; width: 80px;" type="text" value="No action"/>	<input type="text" value="0"/> <input style="border: none; border-bottom: 1px solid #ccc; width: 80px;" type="text" value="No action"/>
When logged off (minutes):	<input type="text" value="0"/> <input style="border: none; border-bottom: 1px solid #ccc; width: 80px;" type="text" value="No action"/>	<input type="text" value="0"/> <input style="border: none; border-bottom: 1px solid #ccc; width: 80px;" type="text" value="No action"/>

Autoscale user interface for *single-session OS random delivery groups*:

## Manage Autoscale Enabled

- General
- Schedule and Peak Ti...**
- Load-based Settings

ADVANCED

- Restrict Autoscale

### Schedule and Peak Times

If you expect a set of users to consume the machine resources at a specific time on specific days, use Autoscale to set multiple schedules that include specific days of the week. You can then adjust the number of machines available during those times. Designate peak usage times for optimized behavior and to control other settings such as capacity buffer under Load-based Settings. [Learn more](#)

[Set schedules](#)

▼ New schedule

Days applied:	Mon	Tue	Wed	Thu	Fri	Sat	Sun
Machines	<a href="#">Edit</a>						
	5	5	5	5	5	5	5
	4	4	4	4	4	4	4
	3	3	3	3	3	3	3
	2	2	2	2	2	2	2
	1	1	1	1	1	1	1
	0	0	0	0	0	0	0

12:00 AM 03:00 AM 06:00 AM 09:00 AM 12:00 PM 03:00 PM 06:00 PM 09:00 PM 12:00 AM

Peak times

> Weekdays

> Weekend

[Save](#) [Cancel](#) [Apply](#)

## Manage Autoscale Enabled

General

Schedule and Peak Ti...

Load-based Settings

ADVANCED

Restrict Autoscale

### Load-based Settings

Capacity buffer is used to accommodate sudden spikes in demand by keeping a buffer of machines powered on based on the maximum load index of the delivery group. By default, the capacity buffer is 10% of the maximum supported load index of the delivery group. A lesser value decreases the cost. A greater value ensures an optimized user experience so that users do not have to wait for additional machines to power on when launching sessions. [Learn more](#)

	During peak times	During off-peak times
Capacity buffer (%):	<input type="text" value="4"/>	<input type="text" value="10"/>
When disconnected (minutes):	<input type="text" value="2"/> <input type="text" value="Suspend"/>	<input type="text" value="3"/> <input type="text" value="Shut down"/>

Autoscale user interface for *multi-session OS delivery groups*:

## Manage Autoscale Enabled

General

Schedule and Peak Ti...

Load-based Settings

ADVANCED

Dynamic Session Tim...

Force User Logoff

Restrict Autoscale

### Schedule and Peak Times

If you expect a set of users to consume the machine resources at a specific time on specific days, use Autoscale to set multiple schedules that include specific days of the week. You can then adjust the number of machines available during those times. Designate peak usage times for optimized behavior and to control other settings such as capacity buffer under Load-based Settings. [Learn more](#)

[Set schedules](#)

▼ New schedule

Days applied:	Mon	Tue	Wed	Thu	Fri	Sat	Sun
Machines	<a href="#">Edit</a>						
	5	5	5	1	5	5	5

12:00 AM 03:00 AM 06:00 AM 09:00 AM 12:00 PM 03:00 PM 06:00 PM 09:00 PM 12:00 AM

Peak times

> Weekdays

> Weekend

Save Cancel Apply

## Manage Autoscale Enabled

- General
- Schedule and Peak Ti...
- Load-based Settings**
- ADVANCED
- Dynamic Session Tim...
- Force User Logoff
- Restrict Autoscale

### Load-based Settings

Capacity buffer is used to accommodate sudden spikes in demand by keeping a buffer of machines powered on based on the maximum load index of the delivery group. By default, the capacity buffer is 10% of the maximum supported load index of the delivery group. A lesser value decreases the cost. A greater value ensures an optimized user experience so that users do not have to wait for additional machines to power on when launching sessions. [Learn more](#)

	During peak times	During off-peak times
Capacity buffer (%):	<input type="text" value="11"/>	<input type="text" value="12"/>

## Schedule-based settings

**Autoscale schedule.** Lets you add, edit, select, and delete schedules.

**Days applied.** Highlights the days you applied to the selected schedule. The remaining days are grayed out.

**Edit.** Lets you assign the machines against each hour or each half hour. You can assign the machines by numbers and by percentages.

### Note:

- This option is available only in the Autoscale user interfaces for multi-session OS and single-session OS random delivery groups.
- The histogram next to **Edit** plots the number or percentage of machines that are running in different time slots.
- You can **assign machines** against each time slot by clicking **Edit** above **Peak times**. Depending on the option you selected from the menu in the **Machines to start** window, you



can assign the machines by numbers or by percentages.

- For multi-session OS delivery groups, you can set the minimum number of running machines separately in granular increments of 30 minutes during each day. For single-session OS random delivery groups, you can set the minimum number of running machines separately in granular increments of 60 minutes during each day.

To define your own schedules, follow these steps:

1. On the **Schedule and Peak Times** page of the **Manage Autoscale** window, click **Set schedules**.
2. In the **Edit Autoscale Schedules** window, select the days you want to apply to each schedule. You can also delete schedules as applicable.
3. Click **Done** to save the schedules and to return to the **Schedule and Peak Times** page.
4. Select the applicable schedule and configure it as needed.
5. Click **Apply** to exit the **Manage Autoscale** window or configure settings on other pages.

**Important:**

- Autoscale does not allow the same day to overlap in different schedules. For example, if you select Monday in schedule2 after selecting Monday in schedule1, Monday is automatically cleared in schedule1.
- A schedule name is not case sensitive.
- A schedule name must not be blank or contain only spaces.
- Autoscale allows blank spaces between characters.
- A schedule name must not contain the following characters: \ / ; : # . \* ? = < > | [ ] ( ) { } “ ‘ .
- Autoscale does not support duplicate schedule names. Enter a different name for each schedule.
- Autoscale does not support empty schedules. This means that schedules without days selected are not saved.

**Note:**

The days included in the selected schedule are highlighted, while those not included are grayed out.

## Load-based settings

**Peak times.** Lets you define the peak times for the days you applied in the selected schedule. You can do so by right-clicking the horizontal bar graph. After you define the peak times, the remaining, undefined times default to off-peak times. By **default**, the 7:00 AM to 7:00 PM time slot is defined as peak times for the days included in the selected schedule.

**Important:**

- For multi-session OS delivery groups, the peak times bar graph is used for the capacity buffer.
- For single-session OS delivery groups, the peak times bar graph is used for the capacity buffer and controls the actions to be triggered after logoff and/or disconnection.
- In the *web-based console*, you can define the peak times for the days included in a schedule at a granular level of 30 minutes for both multi-session OS and single-session OS delivery groups. In the *legacy console*, you can do that only for multi-session OS delivery groups. Alternatively, you can use the `New-BrokerPowerTimeScheme PowerShell` command instead. For more information, see [Broker PowerShell SDK commands](#).

**Capacity buffer.** Lets you keep a buffer of powered-on machines. A lesser value decreases the cost. A greater value ensures an optimized user experience so that when launching sessions, users do not have to wait for additional machines to power on. By default, the capacity buffer is 10% for peak and off-peak times. If you set the capacity buffer to 0 (zero), users might have to wait for additional machines to power on when launching sessions. Autoscale lets you determine the capacity buffer separately for peak and off-peak times.

### Miscellaneous settings

**Tip:**

- You can choose to configure the miscellaneous settings using the Broker PowerShell SDK. For more information, see [Broker PowerShell SDK commands](#).
- To understand the SDK commands associated with the when disconnected and when logged off settings, see [https://citrix.github.io/delivery-controller-sdk/Broker/about\\_Broker\\_PowerManagement/#power-policy](https://citrix.github.io/delivery-controller-sdk/Broker/about_Broker_PowerManagement/#power-policy).

**When disconnected.** Lets you specify how long a disconnected, locked machine remains powered on after session disconnection before it is suspended or shut down. If a time value is specified, the machine is suspended or shut down when the specified disconnection time elapses, depending on the action you configured. By default, no action is assigned to disconnected machines. You can define actions separately for peak and off-peak times. To do so, click the down arrow and then select one of the following options from the menu:

- **No action.** If selected, the machine after session disconnection remains powered on. Autoscale does not act on it.
- **Suspend.** If selected, Autoscale pauses the machine without shutting it down when the specified disconnection time elapses. The following option becomes available after you select **Suspend**.
  - **When no reconnection in (minutes).** Suspended machines remain available to discon-

nected users when they reconnect but are not available for new users. To make the machines available again to handle all workloads, shut them down. Specify the timeout, in minutes, after which Autoscale shuts them down.

- **Shut down.** If selected, Autoscale shuts down the machine when the specified disconnection time elapses.

**Note:**

This option is available only in the Autoscale user interfaces for single-session OS random and static delivery groups.

**When logged off.** Lets you specify how long a machine remains powered on after session logoff before it is suspended or shut down. If a time value is specified, the machine is suspended or shut down when the specified logoff time elapses, depending on the actions you configured. By default, no action is assigned to logged-off machines. You can define actions separately for peak and off-peak times. To do so, click the down arrow and then select one of the following options from the menu:

- **No action.** If selected, the machine after session logoff remains powered on. Autoscale does not act on it.
- **Suspend.** If selected, Autoscale pauses the machine without shutting it down when the specified logoff time elapses.
- **Shut down.** If selected, Autoscale shuts down the machine when the specified logoff time elapses.

**Note:**

This option is available only in the Autoscale user interface for single-session OS static delivery groups.

**Power-off delay.** Lets you specify the minimum number of minutes that must elapse after a machine is powered on before Autoscale powers it off. By default, the power-off delay is 30 minutes. You can set it in a range of 0–60 minutes. For more information, see [How power-off delay works](#).

**Machine instance cost per hour.** Lets you specify the machine instance cost per hour that matches your cost basis. Machine instance cost per hour is the cost per hour, in US\$, of the computing capacity being used. This setting is used to calculate the cost savings of the Autoscale settings above. To view the savings, go to **Monitor > Trends > Machine Usage**. For more information, see [Monitor Autoscale-managed machines](#).

**Note:**

Autoscale does not support changing the currency unit for your cost basis.

### **Power manage single-session OS machines transitioning to a different time period with disconnected sessions**

**Important:**

- This enhancement applies only to single-session OS machines with disconnected sessions. It does not apply to single-session OS machines with logged off sessions.
- For this enhancement to take effect, you need to enable Autoscale for the applicable delivery group. Otherwise, disconnect power policy actions are not triggered on period transition.

In earlier releases, a single-session OS machine transitioning to a time period where an action (disconnect action=“**Suspend**” or “**Shutdown**”) was required remained powered on. This scenario occurred if the machine disconnected during a time period (peak or off-peak times) where no action (disconnect action=“**Nothing**”) was required.

Starting with this release, Autoscale suspends or powers off the machine when the specified disconnection time elapses, depending on the disconnect action configured for the destination time period.

For example, you configure the following power policies for a single-session OS delivery group:

- Set `PeakDisconnectAction` to “Nothing”
- Set `OffPeakDisconnectAction` to “Shutdown”
- Set ‘`OffPeakDisconnectTimeout`’ to “10”

**Note:**

For more information about the disconnect action power policy, see [https://developer-docs.citrix.com/projects/delivery-controller-sdk/en/latest/Broker/about\\_Broker\\_PowerManagement/#power-policy](https://developer-docs.citrix.com/projects/delivery-controller-sdk/en/latest/Broker/about_Broker_PowerManagement/#power-policy) and <https://developer-docs.citrix.com/projects/delivery-controller-sdk/en/latest/Broker/Get-BrokerDesktopGroup/>.

In earlier releases, a single-session OS machine with a session disconnected during peak times remained powered on when it transitioned from peak to off-peak. Starting with this release, the `OffPeakDisconnectAction` and the `OffPeakDisconnectTimeout` policy actions are applied to the single-session OS machine on period transition. As a result, the machine is powered off 10 minutes after it transitions to off-peak.

In case you want to revert to the previous behavior (that is, take no action on machines that transition from peak to off-peak or off-peak to peak with disconnected sessions), do one of the following:

- Set the “LegacyPeakTransitionDisconnectedBehaviour” registry value to 1 (true; enables the previous behavior). By default, the value is 0 (false; triggers disconnect power policy actions on period transition).
  - Path: HKEY\_LOCAL\_MACHINE\SOFTWARE\Citrix\DesktopServer
  - Name: LegacyPeakTransitionDisconnectedBehaviour
  - Type: REG\_DWORD
  - Data: 0x00000001 (1)

- Configure the setting by using the `Set-BrokerServiceConfigurationData` PowerShell command. For example:
  - `PS C:\> Set-BrokerServiceConfigurationData HostingManagement.LegacyPeakTransitionDisconnectedBehaviour -SettingValue $true`

A machine must meet the following criteria before power policy actions can be applied to it on period transition:

- Has a disconnected session.
- Has no pending power actions.
- Belongs to a single-session OS delivery group that transitions to a different time period.
- Has a session that disconnects during a certain time period (peak or off-peak times) and transitions to a period where a power action is assigned.

### How capacity buffer works

Capacity buffer is used to add spare capacity to the current demand to account for dynamic load increases. There are two scenarios to be aware of:

- For multi-session OS delivery groups, the capacity buffer is defined as a percentage of the total capacity of the delivery group in terms of load index. For more information about load index, see [Load index](#).
- For single-session OS delivery groups, the capacity buffer is defined as a percentage of the total capacity of the delivery group in terms of the number of machines.

#### Note:

In scenarios where you restrict Autoscale to tagged machines, the capacity buffer is defined as a percentage of the total capacity of the tagged machines in the delivery group in terms of load index.

Autoscale lets you set the capacity buffer separately for peak and off-peak times. A lesser value in the capacity buffer field decreases the cost because Autoscale powers on less spare capacity. A greater value ensures an optimized user experience so that users do not have to wait for additional machines to power on when launching sessions. By default, the capacity buffer is 10%.

#### Important:

The capacity buffer results in machines being powered on when the total spare capacity drops to a level below “X” percent of the total capacity of the delivery group. Doing so reserves the required percentage of spare capacity.

## How power-off delay works

Use the power-off delay feature to specify the minimum number of minutes that must elapse after a machine is powered on before Autoscale powers it off. Doing so keeps machines from “flip-flopping” on and off during volatile session demands.

### Note:

- In the legacy console, the power-off delay is available only in the Autoscale user interface for multi-session OS delivery groups. You can configure the power-off delay for single-session OS static delivery groups by using the PowerShell SDK. For example: `C:\PS> Set-BrokerDesktopGroup "MyDesktop"-PowerOffDelay 15`.
- In the web-based console, it is available in the Autoscale user interface for both multi-session OS and single-session OS delivery groups.
- For single-session OS static delivery groups, the power-off delay applies to both assigned and unassigned machines.
- In some scenarios, you might want to configure a time period during which the power-off delay does not take effect to prevent Autoscale from powering off the relevant machines. For example, configure a time period to make sure that your logoff scripts can complete successfully before machines are powered off. You can use the PowerShell SDK. For example: `C:\PS> Set-BrokerDesktopGroup "MyDesktop"-SettlementPeriodBeforeAutoShutdown <TimeSpan>`. For more information, see [Broker PowerShell SDK commands](#).

## Multi-session OS delivery groups

### When are machines powered on?

#### Important:

If a schedule is selected, Autoscale powers on all machines configured to be powered on in the schedule. Autoscale keeps this specified number of machines powered on during the schedule, regardless of the load.

When the number of powered-on machines in the delivery group can no longer meet the buffer needed for honoring the buffer capacity in terms of load index, Autoscale powers on extra machines. For example, let's say your delivery group has 20 machines and 3 machines are scheduled to be powered on as part of schedule-based scaling with a capacity buffer of 20%. Eventually, 4 machines will be powered on when there is no load. This is because a 4 x 10k load index is needed as a buffer; therefore at least 4 machines need to be powered on. This case might occur during peak times, increased load on machines, new session launches, and when you add new machines to the delivery group. Note that Autoscale powers on only the machines that meet the following criteria:

- The machines are not in maintenance mode.

- The hypervisor on which the machines are running is not in maintenance mode.
- The machines are currently powered off.
- The machines have no pending power actions.

### When are machines powered off?

#### Important:

- If a schedule is selected, Autoscale powers off the machines based on the schedule.
- Autoscale does not power off the machines configured in the schedule to be powered on during the schedule.

When there are more than enough machines to support the targeted number of powered-on machines (including the buffer) for the delivery group, Autoscale powers off extra machines. This case might occur during off-peak times, decreased load on machines, and session logoffs, and when you remove machines from the delivery group. Autoscale powers off only the machines that meet the following criteria:

- The machines and the hypervisor on which the machines are running are not in maintenance mode.
- The machines are currently powered on.
- The machines are registered as available or waiting to register after start-up.
- The machines have no active sessions.
- The machines have no pending power actions.
- The machines satisfy the specified power-off delay. This means that the machines have been powered on for at least “X” minutes, where “X” is the power-off delay specified for the delivery group.

### Example scenario

Suppose you have the following scenario:

- **Delivery group configuration.** The delivery group that you want Autoscale to power manage contains 10 machines (M1 to M10).
- **Autoscale configuration**
  - Capacity buffer is set to 10%.
  - No machine is included in the selected schedule.

The scenario is executed in the following sequence:

1. No user logs on.

2. User sessions increase.
3. More user sessions start.
4. User session load decreases because of session termination.
5. User session load decreases further until the session load is handled only by on-premises resources.

See below for details about how Autoscale works in the scenario above.

- No user load (initial state)
  - One machine (for example, M1) is powered on. The machine is powered on because of the configured capacity buffer. In this case, 10 (number of machines) x 10,000 (load index) x 10% (configured capacity buffer) equals 10,000. Therefore, one machine is powered on.
  - The load index value of the powered-on machine (M1) is at a baseline load (load index equals 0).
- The first user logs on
  - The session is directed to be hosted on machine M1.
  - The load index of the powered-on machine M1 increases and machine M1 is no longer at a baseline load.
  - Autoscale starts to power on an additional machine (M2) to meet the demand because of the configured capacity buffer.
  - The load index value of machine M2 is at a baseline load.
- Users increase load
  - The sessions are load-balanced across machines M1 and M2. As a result, the load index of the powered-on machines (M1 and M2) increases.
  - The total spare capacity is still at a level above 10,000 in terms of load index.
  - The load index value of machine M2 is no longer at a baseline load.
- More user sessions start
  - The sessions are load-balanced across machines (M1 and M2). As a result, the load index of the powered-on machines (M1 and M2) increases further.
  - When the total spare capacity drops to a level below 10,000 in terms of load index, Autoscale starts to power on an additional machine (M3) to meet the demand because of the configured capacity buffer.
  - The load index value of machine M3 is at a baseline load.
- Even more user sessions start
  - The sessions are load-balanced across machines (M1 to M3). As a result, the load index of the powered-on machines (M1 to M3) increases.
  - The total spare capacity is at a level above 10,000 in terms of load index.
  - The load index value of machine M3 is no longer at a baseline load.
- User session load decreases because of session termination
  - After users log off from their sessions or idle sessions time out, the freed-up capacity on



machines M1 to M3 is reused to host sessions started by other users.

- When the total spare capacity increases to a level above 10,000 in terms of load index, Autoscale puts one of the machines (for example, M3) into drain state. As a result, sessions started by other users are no longer directed to that machine unless new changes occur. For example, end-user load increases again or other machines become least loaded.
- User session load continues to decrease
  - After all sessions on machine M3 are terminated and the specified power-off delay times out, Autoscale powers off machine M3.
  - After more users terminate their sessions, the freed-up capacity on powered-on machines (M1 and M2) is reused to host sessions started by other users.
  - When the total spare capacity increases to a level above 10,000 in terms of load index, Autoscale puts one of the machines (for example, M2) into drain state. As a result, sessions started by other users are no longer directed to that machine.
- User session load continues to decrease until there are no sessions
  - After all sessions on machine M2 are terminated and the specified power-off delay times out, Autoscale powers off machine M2.
  - The load index value of the powered-on machine (M1) is at a baseline load. Autoscale does not put machine M1 into drain state because of the configured capacity buffer.

**Note:**

For multi-session OS delivery groups, all changes to the desktop are lost when users log off sessions. However, if configured, user-specific settings are roamed along with the user profile.

### Single-session OS random delivery groups

Capacity buffer is used to accommodate sudden spikes in demand by keeping a buffer of machines powered on based on the total number of machines in the delivery group. By default, the capacity buffer is 10% of the total number of machines in the delivery group.

If the number of machines (including the capacity buffer) exceeds the total number of currently powered-on machines, additional machines are powered on to meet the demand. If the number of machines (including the capacity buffer) is less than the total number of currently powered-on machines, the excess machines are shut down or suspended, depending on the actions you configured.

### Example scenario

Suppose you have the following scenario:

- **Delivery group configuration.** The delivery group that you want Autoscale to power manage contains 10 machines (M1 to M10).

- **Autoscale configuration**

- Capacity buffer is set to 10%.
- No machine is included in the selected schedule.

The scenario is executed in the following sequence:

1. No user logs on.
2. User sessions increase.
3. More user sessions start.
4. User session load decreases because of session termination.
5. User session load decreases further until the session load is handled only by on-premises resources.

See below for details about how Autoscale works in the scenario above.

- No user load (initial state)
  - One machine (M1) is powered on. The machine is powered on because of the configured capacity buffer. In this case, 10 (number of machines) x 10% (configured capacity buffer) equals 1. Therefore, one machine is powered on.
- A first user logs on
  - The first time a user logs on to use a desktop, the user is assigned a desktop from a pool of desktops hosted on powered-on machines. In this case, the user is assigned a desktop from machine M1.
  - Autoscale starts to power on an additional machine (M2) to meet the demand because of the configured capacity buffer.
- A second user logs on
  - The user is assigned a desktop from machine M2.
  - Autoscale starts to power on an additional machine (M3) to meet the demand because of the configured capacity buffer.
- A third user logs on
  - The user is assigned a desktop from machine M3.
  - Autoscale starts to power on an additional machine (M4) to meet the demand because of the configured capacity buffer.
- A user logs off
  - After a user logs off or the user's desktop times out, the freed-up capacity (for example, M3) is available as buffer. As a result, Autoscale starts to power off machine M4 because the capacity buffer is configured as 10%.
- More users log off until there are no users
  - After more users log off, Autoscale powers off machines (for example, M2 or M3).
  - Even though there are no users left, Autoscale does not power off the remaining one machine (for example, M1) because that machine is reserved as a spare capacity.

**Note:**

For single-session OS random delivery groups, all changes to the desktop are lost when users log off sessions. However, if configured, user-specific settings are roamed along with the user profile.

### Single-session OS static delivery groups

Capacity buffer is used to accommodate sudden spikes in demand by keeping a buffer of unassigned machines powered on based on the total number of unassigned machines in the delivery group. By default, the capacity buffer is 10% of the total number of unassigned machines in the delivery group.

**Important:**

After all machines in the delivery group are assigned, the capacity buffer does not play a role in powering machines on or off.

If the number of machines (including the capacity buffer) exceeds the total number of currently powered-on machines, additional, unassigned machines are powered on to meet the demand. If the number of machines (including the capacity buffer) is less than the total number of currently powered-on machines, excess machines are powered off or suspended, depending on the actions you configured.

For single-session OS static delivery groups, Autoscale:

- Powers assigned machines on during peak times and off during off-peak times only when the `AutomaticPowerOnForAssigned` property of the applicable single-session OS delivery group is set to true.
- Automatically powers on a machine during peak times if it is powered off and the `AutomaticPowerOnForAssignedDuringPeak` property of the delivery group to which it belongs is set to true.

To understand how capacity buffer works with assigned machines, consider the following:

- The capacity buffer works only when the delivery group has one or more unassigned machines.
- If the delivery group has no unassigned machines (all machines in the delivery group are assigned), the capacity buffer does not play a role in powering machines on or off.
- The `AutomaticPowerOnForAssignedDuringPeak` property determines whether assigned machines are powered on during peak times. If it is set to true, Autoscale keeps the machines powered on during peak times. Autoscale will also power them on even if they are powered off.

### Example scenario

Suppose you have the following scenario:

- **Delivery group configuration.** The delivery group that you want Autoscale to power manage contains 10 machines (M1 to M10).
- **Autoscale configuration**
  - Machines M1 to M3 are assigned, and machines M4 to M10 are unassigned.
  - Capacity buffer set to 10% for peak and off-peak times.
  - According to the selected schedule, Autoscale power manages machines between 09:00 AM and 06:00 PM.

See below for details about how Autoscale works in the scenario above.

- Start of schedule – 09:00 AM
  - Autoscale powers on machines M1 to M3.
  - Autoscale powers on an additional machine (for example, M4) because of the configured capacity buffer. Machine M4 is unassigned.
- A first user logs on
  - The first time a user logs on to use a desktop, the user is assigned a desktop from a pool of desktops hosted on unassigned powered-on machines. In this case, the user is assigned a desktop from machine M4. Subsequent logons from that user connect to the same desktop that was assigned on first use.
  - Autoscale starts to power on an additional machine (for example, M5) to meet the demand because of the configured capacity buffer.
- A second user logs on
  - The user is assigned a desktop from the unassigned powered-on machines. In this case, the user is assigned a desktop from machine M5. Subsequent logons from that user connect to the same desktop that was assigned on first use.
  - Autoscale starts to power on an additional machine (for example, M6) to meet the demand because of the configured capacity buffer.
- Users log off
  - As users log off from their desktops or the desktops time out, Autoscale keeps the machines M1 to M5 powered on during 09:00 AM – 06:00 PM. When those users log on the next time, they connect to the same desktop that was assigned on first use.
  - The unassigned machine M6 is waiting to serve a desktop to an incoming, unassigned user.
- End of schedule – 06:00 PM
  - At 06:00 PM, Autoscale powers off machines M1 to M5.
  - Autoscale keeps the unassigned machine M6 powered on because of the configured capacity buffer. That machine is waiting to serve a desktop to an incoming, unassigned user.
  - In the delivery group, machines M6 to M10 are unassigned machines.

## Dynamic session timeouts

July 8, 2021

This feature lets you configure disconnected and idle session timeouts for your peak and off-peak usage times to achieve faster machine draining and cost savings. This feature applies only to multi-session OS machines. A multi-session VDA reports idle times for sessions that have been idle for more than 10 minutes, so dynamic session timeouts will not be able to disconnect idle sessions within 10 minutes of being idle. A lesser value removes lingering sessions sooner, thus reducing costs.

### Manage Autoscale Enabled

- General
- Schedule and Peak Ti...
- Load-based Settings
- ADVANCED
- Dynamic Session Tim...**
- Force User Logoff
- Restrict Autoscale

#### Dynamic Session Timeout

Configure dynamic timeouts for your peak and off-peak usage times to achieve faster VM draining and cost savings. Larger values can improve user experience and smaller values can achieve faster draining. [Learn more](#)

	During peak times	During off-peak times
Idle session timeout: ?	2 <input type="text"/> hour <input type="text"/>	3 <input type="text"/> min <input type="text"/>
Disconnected session timeout: ?	4 <input type="text"/> min <input type="text"/>	5 <input type="text"/> sec <input type="text"/>

**⚠** Autoscale dynamic timeouts are for cost savings. If used for security purposes, the configured timeouts might conflict with your GPO or Studio policies. When a conflict occurs, the shorter timeout prevails. [↗](#)

### Note:

Autoscale dynamic timeouts are for cost savings. If used for security purposes, the configured timeouts might conflict with your GPO or Manage console policies. When a conflict occurs, the shorter timeout prevails.

**Idle session timeout.** Enables or disables a timer that specifies how long an uninterrupted user connection is maintained if there is no user input. When the timer expires, the session is placed in the disconnected state and the **Disconnected session timeout** applies. If the **Disconnected session time-**

**out** is disabled, the session is not logged off.

**Important:**

If you specify a value less than or equal to 10 minutes (600 seconds), Autoscale disconnects the relevant sessions after they have been idle for 10 minutes. This is because Autoscale relies on session idle times that multi-session VDAs report. Multi-session VDAs report idle times only for sessions that have been idle for more than 10 minutes.

**Disconnected session timeout.** Enables or disables a timer that specifies how long a disconnected desktop remains locked before the session is logged off. If enabled, the disconnected session is logged off when the timer expires.

### Broker PowerShell SDK commands

The following Broker PowerShell SDK cmdlets were extended to provide configuration support for dynamic session timeouts.

- \*-BrokerDesktopGroup

The \*-BrokerDesktopGroup PowerShell SDK cmdlets have been extended by adding the following new parameters:

- **DisconnectPeakIdleSessionAfterSeconds** – Represents the time in seconds after which an idle session is disconnected during peak time. This property has a default value of 0, which indicates the disablement of its associated behavior during peak time. A value greater than 0 enables its behavior for the delivery group during peak time only.
- **DisconnectOffPeakIdleSessionAfterSeconds** - Represents the time in seconds after which an idle session is disconnected during off-peak hours. The default value of this property is 0, which indicates the disablement of its associated behavior during off-peak. A value greater than 0 enables its associated behavior for the delivery group during off-peak hours only.
- **LogoffPeakDisconnectedSessionAfterSeconds** - Represents the time in seconds after which a disconnected session is terminated during peak time. The default value of this property is 0, which indicates the disablement of its associated behavior during peak time. A value greater than 0 enables its associated behavior for the delivery group during peak time only.
- **LogoffOffPeakDisconnectedSessionAfterSeconds** - Represents the time in seconds after which a disconnected session is terminated during off-peak hours. The default value of this property is 0, which indicates the disablement of its associated behavior during off-peak. A value greater than 0 enables its associated behavior for the delivery group during off-peak hours only.

This SDK extension affects only the following \*-BrokerDesktopGroup cmdlets:

- Get-BrokerDesktopGroup

- New-BrokerDesktopGroup
- Set-BrokerDesktopGroup

Suppose you want to set the idle session timeout to 3,600 seconds during peak times for a delivery group whose name is “MyDesktop.” Use the Set-BrokerDesktopGroup PowerShell command. For example:

- `C:\PS> Set-BrokerDesktopGroup "MyDesktop"-DisconnectOffPeakIdleSessionAfterSeconds 3600`

Doing that disconnects sessions that have been idle for more than 1 hour in off-peak for the desktop group whose name is “MyDesktop.”

## Restrict Autoscale (cloud burst)

September 8, 2021

### Restrict Autoscale to certain machines in a delivery group

Autoscale provides the flexibility to power manage only a subset of machines in a delivery group. To achieve this, apply a tag to one or more machines and then configure Autoscale to power manage only tagged machines.

This feature can be useful in cloud bursting use cases, where you want to use on-premises resources (or reserved public cloud instances) to handle workloads before cloud-based resources address additional demand (that is, burst workloads). To let on-premises machines (or reserved instances) address workloads first, you must use tag restriction along with zone preference.

Tag restriction specifies machines to be power managed by Autoscale. Zone preference specifies machines in the preferred zone to handle user launch requests. For more information, see [Tags](#) and [Zone preference](#).

To restrict Autoscale to certain machines, you can use the Manage console or PowerShell.

### Use the Manage console to restrict Autoscale to certain machines

(This feature is not available in the Legacy Configuration console.)

To restrict Autoscale to certain machines, complete the following steps:

1. Create a tag and apply that tag to the applicable machines in the delivery group. For more information, see [Manage tags and tag restrictions](#).
2. Select the delivery group and then open the **Manage Autoscale** wizard.

3. On the **Restrict Autoscale** page, select **Enable tag restriction**, select a tag from the list, and then click **Apply** to save your changes.

User interface for single-session OS *static* and *random* delivery groups:

The screenshot shows the 'Manage Autoscale' interface. At the top, it says 'Manage Autoscale' with a status indicator 'Enabled'. On the left, there is a navigation menu with options: 'General', 'Schedule and Peak Ti...', 'Load-based Settings', 'ADVANCED', and 'Restrict Autoscale'. The 'Restrict Autoscale' option is selected. The main content area is titled 'Restrict Autoscale to Tagged Machines' and includes a descriptive paragraph: 'Use this feature if you want on-premises resources (or reserved public cloud instances) to handle workloads before cloud-based resources address additional demand (burst workloads). To have on-premises machines (or reserved instances) address workloads first, you must use tag restriction along with zone preference. [Learn more](#)'. Below this is a checkbox labeled 'Enable tag restriction:' which is currently unchecked. To the right of the checkbox is a dropdown menu with the text 'Select a tag' and a downward arrow. At the bottom right of the interface, there are three buttons: 'Save', 'Cancel', and 'Apply'.

User interface for *multi-session OS* delivery groups:



The screenshot shows the 'Manage Autoscale' configuration page. At the top, it says 'Manage Autoscale' with a status indicator 'Enabled'. On the left is a navigation menu with items: 'General', 'Schedule and Peak Ti...', 'Load-based Settings', 'ADVANCED', 'Dynamic Session Tim...', 'Force User Logoff', and 'Restrict Autoscale'. The main content area is titled 'Restrict Autoscale to Tagged Machines' and contains the following text: 'Use this feature if you want on-premises resources (or reserved public cloud instances) to handle workloads before cloud-based resources address additional demand (burst workloads). To have on-premises machines (or reserved instances) address workloads first, you must use tag restriction along with zone preference. [Learn more](#)'. Below this text is a checkbox labeled 'Enable tag restriction:' which is currently unchecked. To the right of the checkbox is a dropdown menu with the text 'Select a tag' and a downward arrow. At the bottom right of the page are three buttons: 'Save', 'Cancel', and 'Apply'.

**Warning:**

- Restricting Autoscale to machines with a specific tag might cause the histogram to update automatically to reflect the number of machines per the tag. On the **Schedule and Peak Times** page, you can manually assign machines against each time slot if needed.
- You cannot delete a tag that is being used to restrict Autoscale. To delete the tag, you must first remove the tag restriction.

After you apply the tag restriction, you might want to remove it from the delivery group later. To do so, go to the **Manage Autoscale > Restrict Autoscale** page and then clear **Enable tag restriction**.

**Warning:**

- If you remove the tag from the applicable machines without clearing **Enable tag restriction**, you might receive a warning when you open the **Manage Autoscale** wizard. Removing the tag from the machines can leave no machines for Autoscale to manage because the tag you specified in Autoscale has become invalid. To resolve the warning, go to the **Restrict Autoscale** page, remove the invalid tag, and then click **Apply** to save your changes.

## Use PowerShell to restrict Autoscale to certain machines

To use the PowerShell SDK directly, complete the following steps:

1. **Create a tag.** Use the `New-BrokerTag` PowerShell command to create a tag.
  - For example: `$managed = New-BrokerTag Managed`. In this case, the tag is named “Managed.” For more information about the `New-BrokerTag` PowerShell command, see <https://developer-docs.citrix.com/projects/delivery-controller-sdk/en/latest/Broker/New-BrokerTag/>.
2. **Apply the tag to machines.** Use the `Get-BrokerMachine` PowerShell command to apply the tag to machines in a catalog that you want Autoscale to power manage.
  - For example: `Get-BrokerMachine -CatalogName "cloud" | Add-BrokerTag $managed.Name`. In this case, the catalog is named “cloud.”
  - For more information about the `Get-BrokerMachine` PowerShell command, see <https://developer-docs.citrix.com/projects/delivery-controller-sdk/en/latest/Broker/Get-BrokerMachine/>.

### Note:

You might add new machines to the catalog after applying the tag. The tag is *NOT* automatically applied to those new machines.

3. **Add tagged machines to the delivery group that you want Autoscale to power manage.** Use the `Get-BrokerDesktopGroup` PowerShell command to add a tag restriction to the delivery group that contains the machines (in other words, “restrict launches to machines with tag X”).
  - For example: `Get-BrokerDesktopGroup -Uid 1 | Set-BrokerDesktopGroup -RestrictAutoscaleTagUid $managed.Uid`. In this case, the UID of the Delivery Group is 1.
  - For more information about the `Get-BrokerDesktopGroup` PowerShell command, see <https://developer-docs.citrix.com/projects/delivery-controller-sdk/en/latest/Broker/Get-BrokerDesktopGroup/>.

After you apply the tag restriction, you might want to remove it from the delivery group later. To do so, use the `Get-BrokerDesktopGroup` PowerShell command.

Example: `Get-BrokerDesktopGroup -Uid 1 | Set-BrokerDesktopGroup -RestrictAutoscaleTagUid $null`. In this case, the UID of the delivery group is 1.

### Note:

Untagged machines restart automatically after users power them off. This behavior ensures that they become available to handle workloads sooner. This can be enabled or disabled on a per desktop group using the `Set-BrokerDesktopGroup's AutomaticRestartForUntaggedMachines`

property. For more information, see <https://citrix.github.io/delivery-controller-sdk/Broker/Set-BrokerDesktopGroup/>.

## Example scenario

Suppose you have the following scenario:

- **Machine catalog configuration.** There are two machine catalogs (C1 and C2).
  - Catalog C1 contains 5 machines (M1 to M5) that are local in the on-premises deployments.
  - Catalog C2 contains 5 machines (M6 to M10) that are remote in the cloud deployments.
- **Tag restriction.** A tag named “Cloud” is created and applied to machines M6 to M10 in catalog C2.
- **Zone configuration.** Two zones (Z1 and Z2) are created.
  - Zone Z1 containing catalog C1 corresponds to the on-premises deployments.
  - Zone Z2 containing catalog C2 corresponds to the cloud deployments.
- **Delivery group configuration**
  - The delivery group contains 10 machines (M1 to M10), 5 machines from catalogs C1 (M1 to M5) and 5 from catalog C2 (M6 to M10).
  - Machines M1 to M5 are powered on manually and remain powered on throughout the schedule.
- **Autoscale configuration**
  - Capacity buffer is set to 10%.
  - Autoscale power manages only machines with the tag “Cloud.” In this case, Autoscale power manages cloud machines M6 to M10.
- **Published application or desktop configuration.** Zone preferences are configured for the published desktops (for example), where Zone Z1 is preferred over Zone Z2 for a user launch request.
  - Zone Z1 is configured as the preferred zone (home zone) for the published desktops.

The scenario is executed in the following sequence:

1. No user logs on.
2. User sessions increase.
3. User sessions increase further until all available on-premises machines are consumed.
4. More user sessions start.
5. User session decreases because of session termination.
6. User session decreases further until the session load is handled only by on-premises machines.

See below for details about how Autoscale works in the scenario above.

- No user load (initial state)
  - The on-premises machines M1 to M5 are all powered on.
  - One machine in the cloud (for example, M6) is powered on. The machine is powered on because of the configured capacity buffer. In this case, 10 (number of machines) x 10,000 (load index) x 10% (configured capacity buffer) equals 10,000. Therefore, one machine is powered on.
  - The load index value of all the powered-on machines (M1 to M6) is at a baseline load (load index equals 0).
- Users log on
  - The sessions are directed to be hosted on machines M1 to M5 through the configured zone preference and are load-balanced across these on-premises machines.
  - The load index value of the powered-on machines (M1 to M5) increases.
  - The load index value of the powered-on machine M6 is at a baseline load.
- Users increase load, consuming all on-premises resources
  - The sessions are directed to be hosted on machine M1 to M5 through the configured zone preference and are load-balanced across these on-premises machines.
  - The load index value of all the powered-on machines (M1 to M5) has reached 10,000.
  - The load index value of the powered-on machine M6 remains at a baseline load.
- One more user logs on
  - The session overflows the zone preference and is directed to be hosted on cloud machine M6.
  - The load index value of all the powered-on machines (M1 to M5) has reached 10,000.
  - The load index value of the powered-on machine M6 increases and is no longer at a baseline load. When the total spare capacity drops to a level below 10,000 in terms of load index, Autoscale starts to power on an additional machine (M7) to meet the demand because of the configured capacity buffer. Note that it might take some time to power on machine M7. So there might be a delay until machine M7 is ready.
- More users log on
  - The sessions are directed to be hosted on machine M6.
  - The load index value of all the powered-on machines (M1 to M5) has reached 10,000.
  - The load index value of the powered-on machine M6 increases further, but the total spare capacity is at a level above 10,000 in terms of load index.
  - The load index value of the powered-on machine M7 remains at a baseline load.
- Even more users log on
  - After machine M7 is ready, the sessions are directed to be hosted on machines M6 and M7 and are load-balanced across these machines.
  - The load index value of all the powered-on machines (M1 to M5) has reached 10,000.
  - The load index value of machine M7 is no longer at a baseline load.
  - The load index value of the powered-on machines (M6 and M7) increases.

- The total spare capacity is still at a level above 10,000 in terms of load index.
- User session load decreases because of session termination
  - After users log off from their sessions or idle sessions time out, the freed-up capacity on machines M1 to M7 is reused to host sessions started by other users.
  - When the total spare capacity increases to a level above 10,000 in terms of load index, Autoscale puts one of the cloud machines (M6 to M7) into drain state. As a result, sessions started by other users are no longer directed to that machine (for example, M7) unless new changes occur; for example, user load increases again or other cloud machines become least loaded.
- User session load decreases further until one or more cloud machines are no longer needed
  - After all sessions on machine M7 are terminated and the specified power-off delay times out, Autoscale powers off machine M7.
  - The load index value of all the powered-on machines (M1 to M5) might drop to a level below 10,000.
  - The load index value of the powered-on machine (M6) decreases.
- User session decreases further until no cloud machines are needed.
  - Even though there are no user sessions on machine M6, Autoscale does not power it off because it is reserved as a spare capacity.
  - Autoscale keeps the remaining cloud machine M6 powered on because of the configured capacity buffer. That machine is waiting to serve a desktop to an incoming user.
  - Sessions are not directed to be hosted on machine M6 as long as the on-premises machines have available capacity.

## Dynamic machine provisioning

February 11, 2021

### Dynamically provision machines with Autoscale

Autoscale provides the capability to create machines and delete them dynamically. You can leverage the capability by using a PowerShell script. The script helps you dynamically scale up or down the number of machines in the delivery group based on the current load conditions.

The script offers the following benefits (and more):

- **Reducing storage costs.** Different from Autoscale, which helps reduce your computing costs, the script provides a more cost-effective solution to provision machines.
- **Effectively handling load changes.** The script helps you handle load changes by automatically scaling up or down the number of machines based on the current delivery group load.

## Download the script

The PowerShell script is available at <https://github.com/citrix/Powershell-Scripts/tree/master/XAXD/AutoscaleMcs>.

## How the script works

### Important:

- You cannot specify a machine catalog in more than one delivery group that is to be managed by the script. In other words, if multiple delivery groups share the same machine catalog, the script does not work with any of those delivery groups.
- You cannot concurrently run the script for the same delivery group from multiple locations.

The script works at a delivery group level. It measures the load (in terms of [load index](#)) and then determines whether to create or delete machines.

Machines created through this script are uniquely tagged (through the `ScriptTag` parameter) so that they can be identified later. Creating or deleting machines is based on:

- **Maximum percentage load of a delivery group.** Specifies the maximum level at which to create machines for Autoscale to address extra loads. When this threshold is exceeded, machines are created in batches to ensure that the current load decreases to or below the threshold.
- **Minimum percentage load of a delivery group.** Specifies the minimum level at which to delete machines created through this script that have no active sessions. When this threshold is exceeded, machines created through this script that have no active sessions are deleted.

This script is intended to monitor across a delivery group and to create or delete machines when the trigger criterion is met. It executes on a per-run basis. This means that you need to run the script on a regular basis so that it can function as intended. We recommend that you run the script at a minimum interval of five minutes. Doing so improves overall responsiveness.

The script relies on the following parameters to work:

Parameter	Type	Default value	Description
DeliveryGroupName	String	X	Name of the delivery group to be monitored to determine the current load. You can provide a semicolon-separated list of names. For example: <code>Invoke-AutoscaleMachineCreation .ps1 - DeliveryGroupName 'dg1;dg2;dg3' - XdProfileName profile</code> .
XdProfileName	String	X	Name of the profile to use for authenticating to remote servers. For details about authenticating to remote servers using this parameter, see <a href="#">Authentication API</a> .
HighWatermark	Integer	80	Maximum percentage load (in terms of load index) at which to create machines for Autoscale to address extra loads.
LowWatermark	Integer	15	Minimum percentage load (in terms of load index) at which to delete machines created through this script that have no active sessions.

Parameter	Type	Default value	Description
MachineCatalogName	String	X	Name of the machine catalog where machines are to be created.
MaximumCreatedMachines	Integer	-1	Maximum amount of machines that can be created in a specified delivery group. If the value is equal to or less than 0, the script does not process this parameter.
ScriptTag	String	AutoscaledScripted	Tag that applies to machines created through the script.
EventLogSource	String	X	Source name that appears in Windows Event Viewer.

**Note:**

An “X” indicates that no default value is specified for that parameter.

By default, the script requires all parameters (except the `ScriptTag` parameter) the first time it runs. On subsequent runs, only the `DeliveryGroupName` and the `XdProfileName` parameters are required. Optionally, you can choose to update the minimum and maximum percentage loads.

Note that you must specify a single delivery group the first time you run the script. For example, the script does *not* work if you use the following PowerShell command to specify two delivery groups the first time you run the script:

- `Invoke-AutoscaleMachineCreations.ps1 -DeliveryGroupName 'dg1;dg2' -XdProfileName profile -LowWatermark 20 -HighWatermark 70 -MachineCatalogName 'cat1' \`

Instead, first specify a single delivery group (in this example, `dg1`) using the following command:

- `Invoke-AutoscaleMachineCreations.ps1 -DeliveryGroupName 'dg1' -XdProfileName profile -LowWatermark 20 -HighWatermark 70 -MachineCatalogName 'cat1' \`



Then, use the following command to run the script for the second delivery group (in this example, dg 2):

- `Invoke-AutoscaleMachineCreations.ps1 -DeliveryGroupName 'dg1;dg2' -XdProfileName profile`

## Prerequisites

To run the script, make sure that these prerequisites are met:

- The machine resides within the same domain where machines are being created.
- Remote PowerShell SDK is installed on that machine. For more information about the Remote PowerShell SDK, see [SDKs and APIs](#).
- Other prerequisites:
  - A delivery group to monitor
  - A machine catalog created through Machine Creation Services (MCS) that has an associated provisioning scheme (template)
  - An identity pool that is associated with the provisioning scheme
  - An event log source to be created so that the script can write information to the Windows Event Log
  - A secure client that allows you to authenticate to remote servers

## Permissions, recommendations, and notices

When you run the script, keep the following in mind:

- To authenticate to remote servers using the `XdProfileName` parameter, you need to define an authentication profile by using an API access secure client, created in the Citrix Cloud console. For details, see [Authentication API](#).
- You must have permissions to create and delete machine accounts in Active Directory.
- We recommend that you automate the PowerShell script with Windows Task Scheduler. For details, see [Create an automated task using Windows Task Scheduler](#).
- If you want the script to write information (for example, failures and actions) to the Windows Event Log, you need to first specify a source name using the `New-EventLog` cmdlet. For example, `New-EventLog -LogName Application -Source <sourceName>`. You can then view the events in the **Application** pane of Windows Event Viewer.
- If errors occurred during execution of the script, execute the script manually and then troubleshoot problems by performing script checks.

## Authentication API

Before you run the script, you need to define an authentication profile by using an API access secure client. You must create a secure client using the same account under which the script will run.

The secure client must have the following permissions:

- Create and delete machines using MCS.
- Edit machine catalogs (to add and remove machines).
- Edit delivery groups (to add and remove machines).

When you create a secure client, make sure that your account has the permissions above because the secure client automatically inherits the permissions from your current account.

To create a secure client, complete these steps:

1. Sign in to Citrix Cloud and then navigate to **Identity and Access Management > API Access**.
2. Type the name for your secure client and then click **Create Client**.

To authenticate to remote servers, use the `Set-XDCredentials` PowerShell command. For example:

- `Set-XDCredentials -APIKey <key_id> -CustomerId <customer_id> -SecretKey <secret_key> -StoreAs <name specified by the XdProfileName parameter>`

## Create an automated task using Windows Task Scheduler

You can automate the PowerShell script with Windows Task Scheduler. Doing so lets the script run automatically at certain intervals or when certain conditions are met. To execute this script with Windows Task Scheduler, make sure to select **Do not start a new instance** on the **Create Task > Settings** tab. Doing so prevents the Windows Task Scheduler from running a new instance of the script if the script is already running.

## Script execution example

See below for an example of executing the script. Note that the script file is invoked multiple times. In this example, to simulate the load, one session is launched and then terminated.

```
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName devtest -XdProfileName profile -MachineCatalogName autoscaled -ScriptTag "devtest"
[devtest]: Assuming default values for watermarks [15 : 80].
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName engtest -XdProfileName profile -MachineCatalogName autoscaled2 -ScriptTag "engtest"
[engtest]: Assuming default values for watermarks [15 : 80].
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
[engtest]: Provisioning more machines. Current Usage [99.99] >= High Watermark [80].
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
[engtest]: Began provisioning of [1] machines to [engtest]. Monitoring task [ca2b0cad-9c50-4e20-8f1d-9ff81307b201].
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
[engtest]: Provisioning task [ca2b0cad-9c50-4e20-8f1d-9ff81307b201] is complete. [1] created. [0] failed to create.
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
[engtest]: Added [1] machines to [engtest].
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
[engtest]: Removing extraneous machines: Current Usage [0] <= Low Watermark [15].
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
[engtest]: Removing [1] machines from [engtest]. Monitoring task [28c6c242-af81-4693-a2a8-0587f09689b4]
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
[engtest]: Machine deletion task [28c6c242-af81-4693-a2a8-0587f09689b4] is [Finished].
PS C:\Users\crisj\Desktop> .\Invoke-AutoscaleMachineCreation.ps1 -DeliveryGroupName "devtest;engtest" -XdProfileName profile
```

### Troubleshooting checklist for the script

The script writes information (for example, errors and actions) to the Windows Event Log. The information helps you troubleshoot issues you experience when executing the script. It might be helpful to keep the following troubleshooting checklist in mind:

- Failure to communicate with remote servers. Possible actions:
  - Verify your connection to the server.
  - Verify that the API key you use is valid.
- Failure to create machines. Possible actions:
  - Verify that the user account running the script has sufficient permissions to create user accounts in the domain.
  - Verify that the user who created the API key has sufficient permissions to use MCS to provision machines.
  - Verify the validity of the machine catalog (that is, its image still exists and is in good state).
- Failure to add machines to a machine catalog or a delivery group. Possible action:
  - Verify that the user who created the API key has sufficient permissions to add and remove machines to and from machine catalogs and delivery groups.

### Force user logoff

June 14, 2021

To better achieve cost savings, Autoscale allows you to force log off lingering sessions. It does so by allowing the administrator to send a custom notification to the users and a grace period after which the sessions are force logged off. This is done only for machines in the drain mode and not for all

powered-on machines.

## Manage Autoscale Enabled

General

Schedule and Peak Ti...

Load-based Settings

ADVANCED

Dynamic Session Tim...

Force User Logoff

Restrict Autoscale

### Force User Logoff

This feature allows you to shut machines down faster by removing lingering sessions from the machines in drain state. You can send a notification to users before logging them off after the specified time. [Learn more](#)

Enable force logoff during peak times

Time after which users are logged off from their sessions

min

Enable force logoff during off-peak times

Time after which users are logged off from their sessions

min

**Display a notification before logging users off**

Notification title:

Notification message: 

Example: Warning: To save costs, the machine will be shut down and you will be logged off from the session. Save your work and log back on to get a different machine.

Save

Cancel

Apply

You have the following two options:

- **Force user logoffs during peak times.** If selected, Autoscale logs off those users from their sessions during peak times when the specified time elapses.
- **Force user logoffs during off-peak times.** If selected, Autoscale logs off those users from their sessions during off-peak times when the specified time elapses.

#### **Important:**

This feature is available only in the web-based console. It is available only in the Autoscale user interface for multi-session app based delivery groups.

**Display a notification before logging users off.** Lets you send notifications to users before they are logged off.

- **Notification title.** Lets you specify a title of the notification to be sent to users. Example: A forced logoff has been initiated.
- **Notification message.** Lets you specify the content of the notification to be sent to users. Ex-

ample: Warning: Your administrator is about to power off your machine and you will be logged off from the session. Save your work and log back on to get a different machine.

## Cloud Health Check

August 19, 2021

Cloud Health Check allows you to run checks that gauge the health and availability of the site and its components. You can run health checks for Virtual Delivery Agents (VDAs), StoreFront servers, and Profile Management. VDA health checks identify possible causes for common VDA registration, session launch, and time zone redirection issues.

If issues are present during the checks, Cloud Health Check provides a detailed report and the actions to fix the issues. Each time Cloud Health Check starts, it checks for the latest version of scripts on the Content Delivery Network (CDN) and automatically downloads the scripts if they do not exist on the local machine. Cloud Health Check always chooses the latest local version of scripts to run health checks.

**Note:**

Cloud Health Check does not update every time it runs.

In a Citrix Cloud environment, run Cloud Health Check from a domain-joined machine to run checks on one or more VDAs or StoreFront servers.

**Note:**

You cannot install or run Cloud Health Check on a Cloud Connector.

The log for the Cloud Health Check application is stored in `C:\ProgramData\Citrix\TelemetryService\CloudHealthCheck.log`. You can use this file for troubleshooting.

View an introduction to Cloud Health Check.

[This is an embedded video. Click the link to watch the video](#)

View when to use Cloud Health Check.

[This is an embedded video. Click the link to watch the video](#)

## Installation

To prepare your environment for installation of Cloud Health Check, you must have a domain-joined Windows machine.


**Note:**

You cannot install or run Cloud Health Check on Cloud Connector.

1. On the domain-joined machine, download the [Cloud Health Check installer](#).
2. Double-click the CloudHealthCheckInstaller\_x64.msi file.
3. Click the box to accept the terms.
4. Click Install.



5. After installation has completed, click **Finish**.



## Permissions and requirements

Permissions:

- To run health checks:
  - You must be a member of the domain users group.
  - You must be a full administrator or have a custom role with read-only and **Run Environment Tests** permissions for the site.
  - Set the script execution policy to at least `RemoteSigned` to allow the scripts to run. For example: `Set-ExecutionPolicy RemoteSigned`. **Note:** other script execution privileges can work as well.
- Use **Run as administrator** when launching Cloud Health Check.

For each VDA or StoreFront machine that you run health checks on:

- The OS must be 64-bit.
- Cloud Health Check must be able to communicate with the machine.
- File and printer sharing must be turned on.
- PSRemoting and WinRM must be enabled. The machine must also be running PowerShell 3.0 or later.
- Windows Management Infrastructure (WMI) access must be enabled on the machine.

## About health checks

Health check data is stored in folders under `C:\ProgramData\Citrix\TelemetryService\`.

### VDA health checks

For registration on the VDA, Cloud Health Check checks:

- VDA software installation
- VDA machine domain membership
- VDA communication port availability
- VDA service status
- Windows firewall configuration
- Communication with Controller
- Time sync with Controller
- VDA registration status

For session launches on VDAs, Cloud Health Check checks:

- Session launch communication port availability
- Session launch services status
- Session launch Windows firewall configuration
- VDA Remote Desktop Services Client Access Licenses
- VDA application launch path

For time zone redirection on VDAs, Cloud Health Check checks:

- Windows hotfix installation
- Citrix hotfix installation
- Microsoft group policy settings
- Citrix group policy settings

For Profile Management on VDAs, Cloud Health Check checks:

- Hypervisor detection
- Provisioning detection
- Citrix Virtual Apps and Desktops
- Personal vDisk configuration
- User store
- Profile Management Service status detection
- Winlogon.exe hooking test

To run checks on Profile Management, you must install and enable Profile Management on the VDA. For more information on Profile Management configuration checks, see Knowledge Center article [CTX132805](#).



## StoreFront health checks

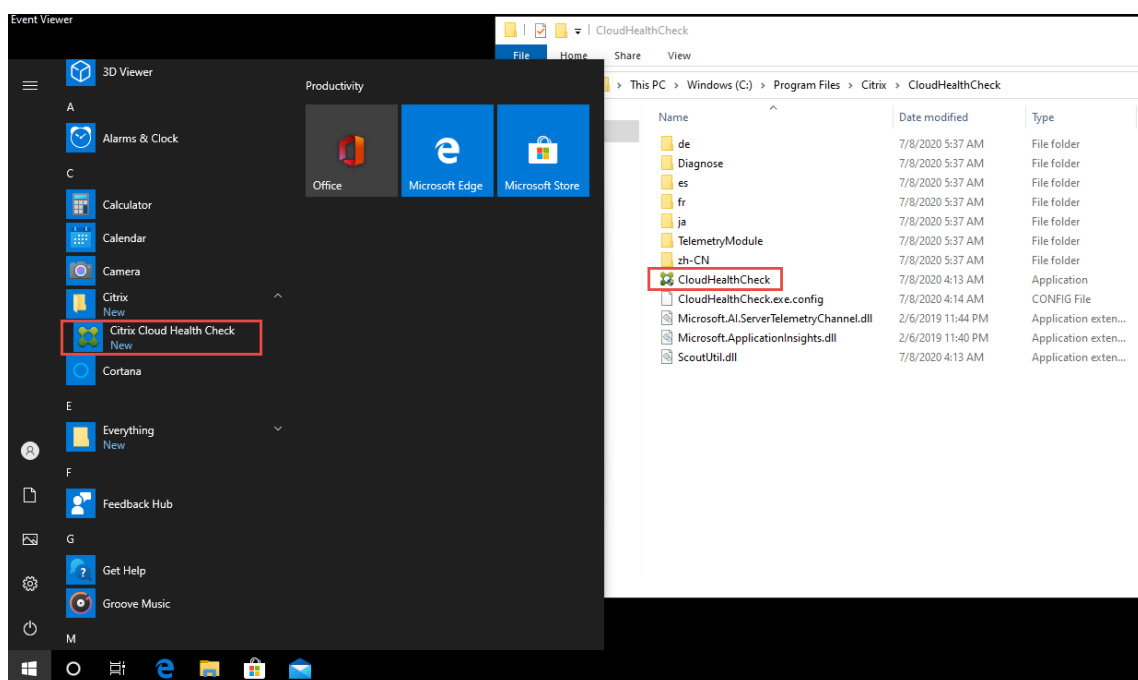
StoreFront checks verify whether:

- Citrix Default Domain service is running
- Citrix Credential Wallet service is running
- The connection from the StoreFront server to Active Directory is port 88
- The connection from the StoreFront server to Active Directory is port 389
- The connection from the StoreFront server to Active Directory is port 464
- The base URL has a valid FQDN
- The correct IP address from the base URL can be retrieved
- The IIS application pool is using .NET 4.0
- The certificate is bound to the SSL port for the host URL
- The certificate chain is complete
- The certificates have expired
- A certificate is expiring within 30 days

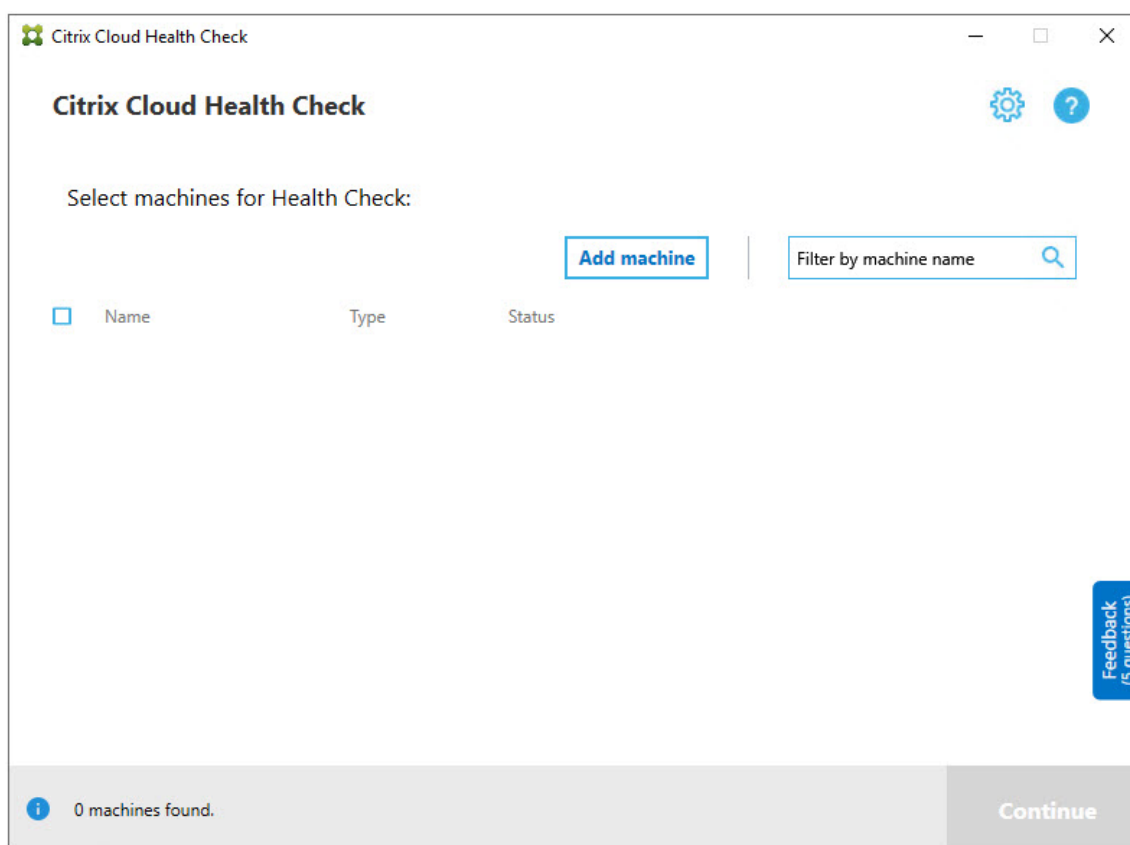
## Running Cloud Health Check

To run Citrix Cloud Health Check:


1. Select **Citrix > Citrix Cloud Health Check** from the machine's Start menu, or run `CloudHealthCheck.exe` in `C:\Program Files\Citrix\CloudHealthCheck`.



2. On the main Cloud Health Check screen, click **Add machine**.



3. Type the FQDN of the machine you want to add. **Note:** Although entering a DNS alias instead of an FQDN can appear valid, the health checks might fail.
4. Click **Continue**.
5. Repeat to add other machines as needed.



6. To remove a manually added machine, click the **X** on the right end of the row and confirm the deletion. Repeat to delete other manually added machines.

Cloud Health Check remembers manually added machines until you remove them. When you close and then reopen Cloud Health Check, the manually added machines are still listed at the top of the list.

## Import VDA machines

You can import VDA machines in the deployment when running health checks.

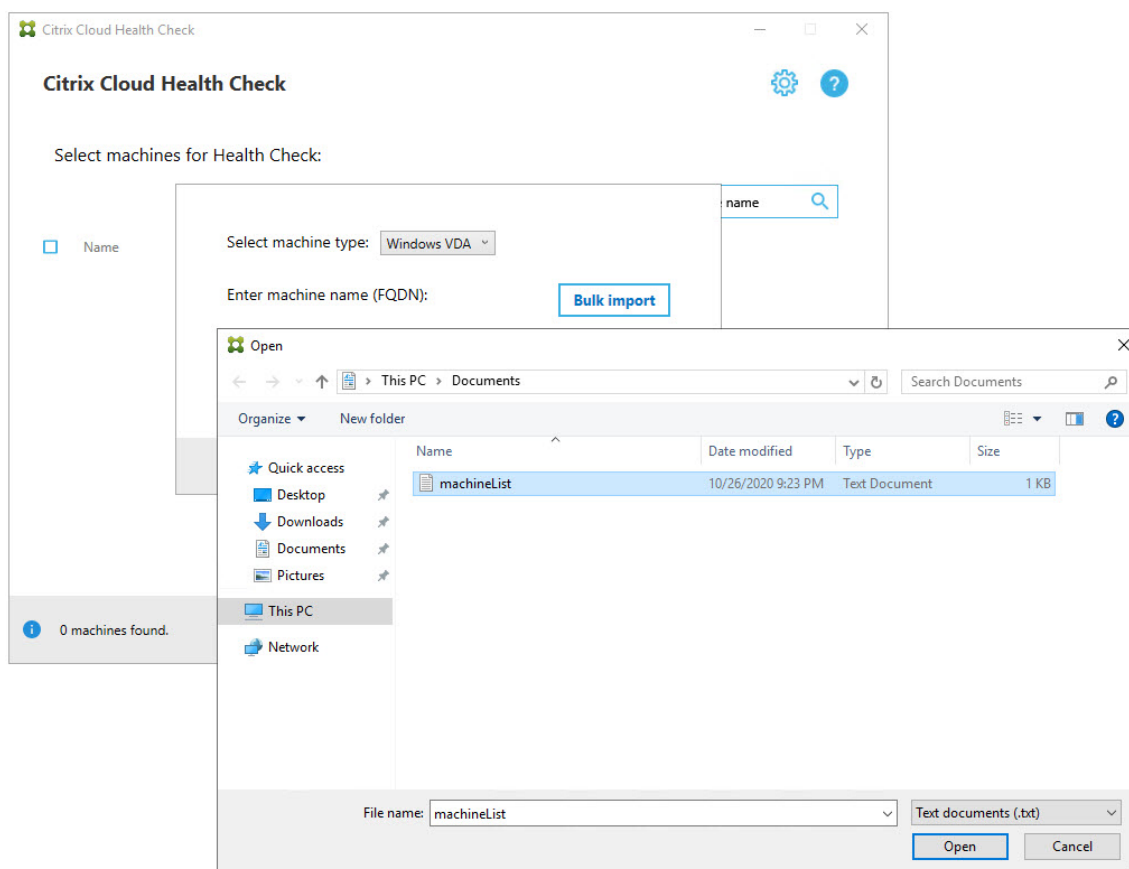
1. On Connector, generate the machine list file with the following PowerShell command. On Connector, you must input Citrix credentials and select the customer in the pop-up dialog.

```
Get-BrokerMachine | foreach { $_.DnsName } | out-file C:\machineList.txt
```

1. Copy the machineList.txt file to the domain-joined machine you want to run Cloud Health Check on.
2. On the Cloud Health Check page, click **Add Machine**.
3. Select the Windows VDA machine type.
4. Click **Import VDA machines**.

5. Select the machineList.txt file.

6. Click **Open**.



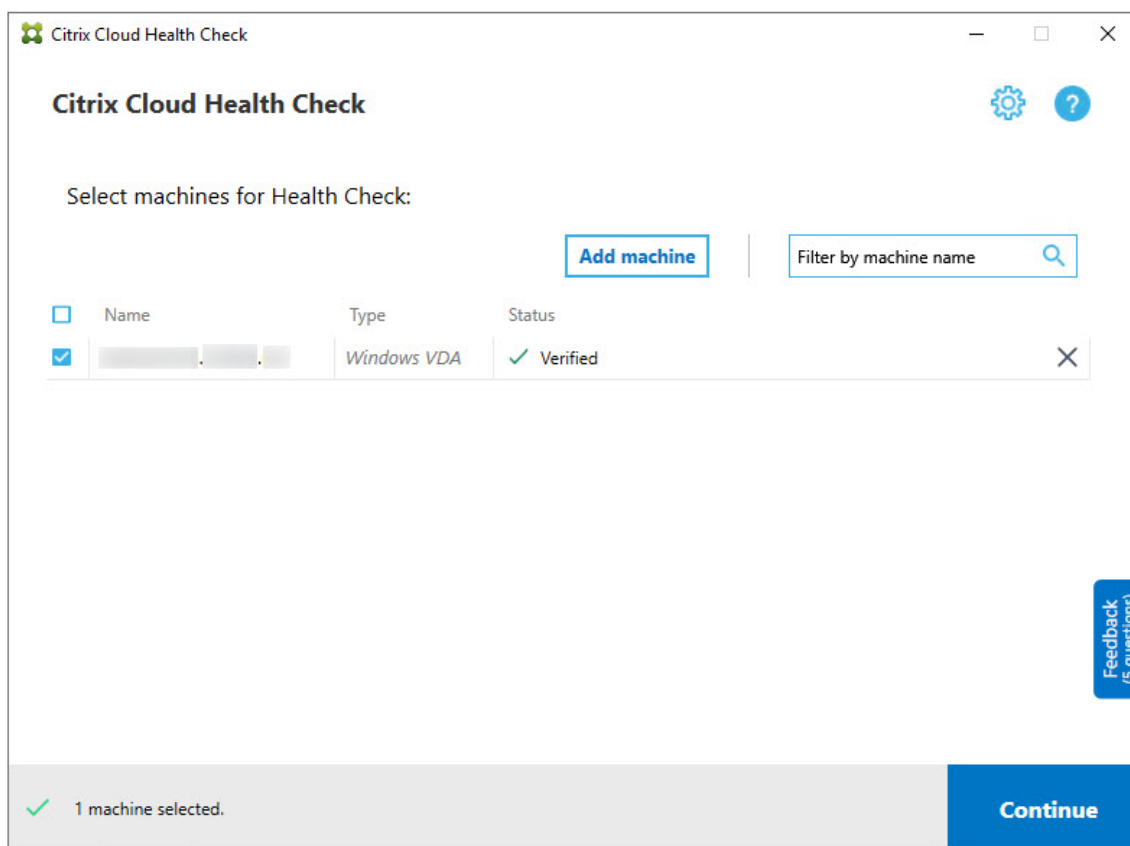
The imported VDA machines are listed on the Cloud Health Check page.

7. Select the check box next to each machine you want to run health checks on.

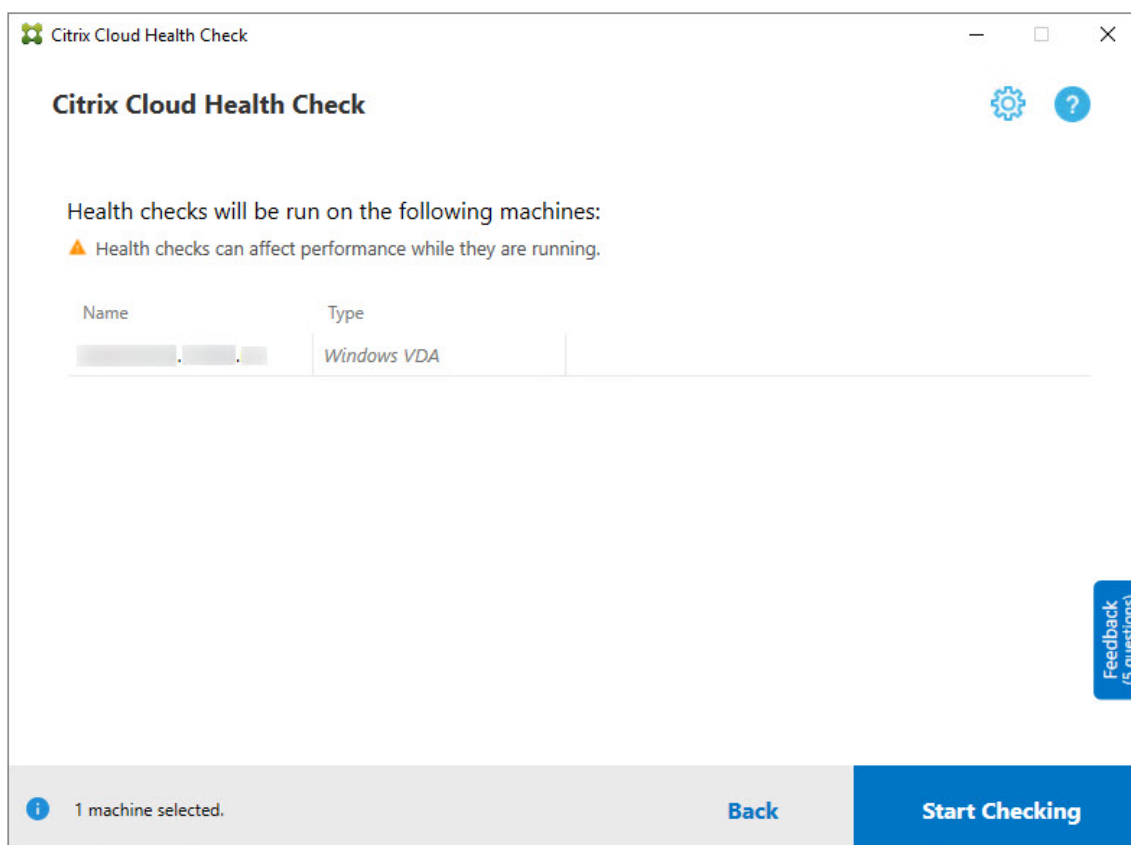
Cloud Health Check automatically launches verification tests on each selected machine, making sure it meets the criteria listed in verification tests. If verification fails, a message appears in the **Status** column, and that machine's check box is cleared. You can then:

- Resolve the issue and then select the machine's check box again. This triggers a retry of the verification tests.
- Skip that machine by leaving its check box unselected. Health checks are not run for that machine.

8. When the verification tests are complete, click **Continue**.

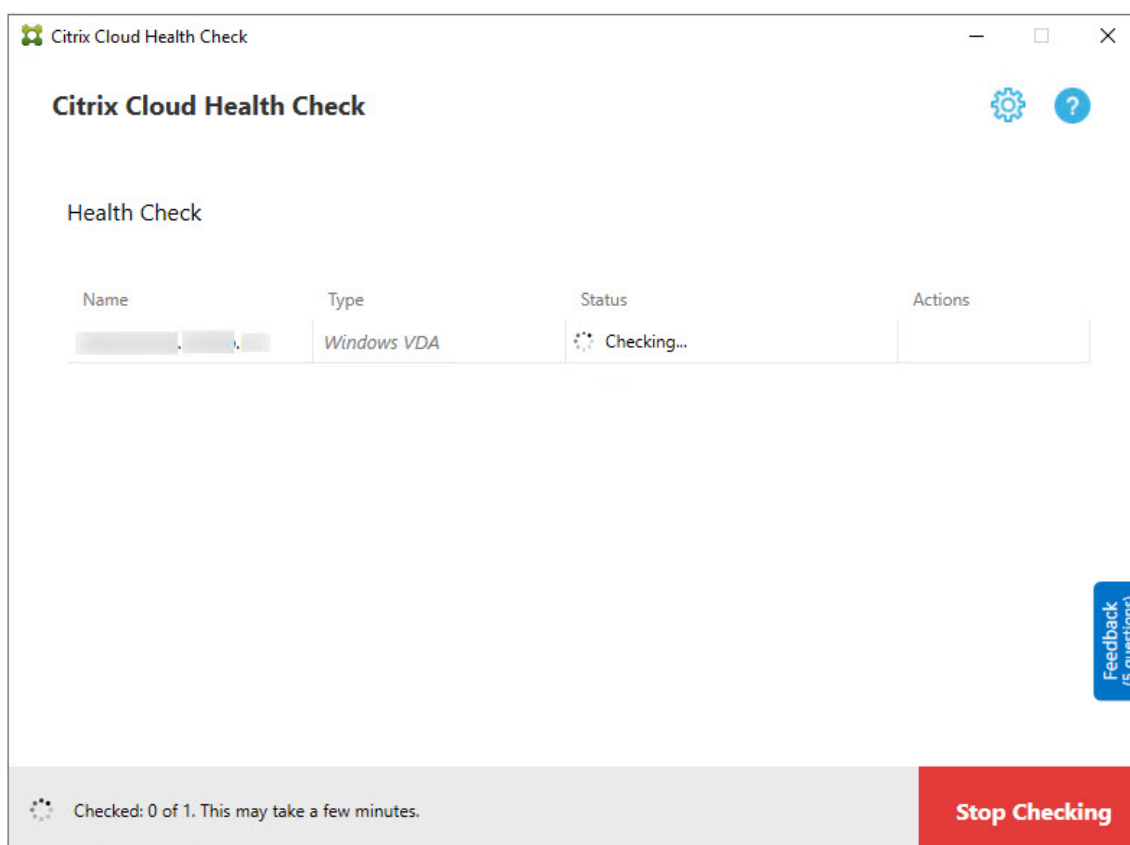


9. Run the health checks on the selected machines. The summary lists the machines where the tests run (the machines you selected that passed the verification tests).
10. Click **Start Checking**.

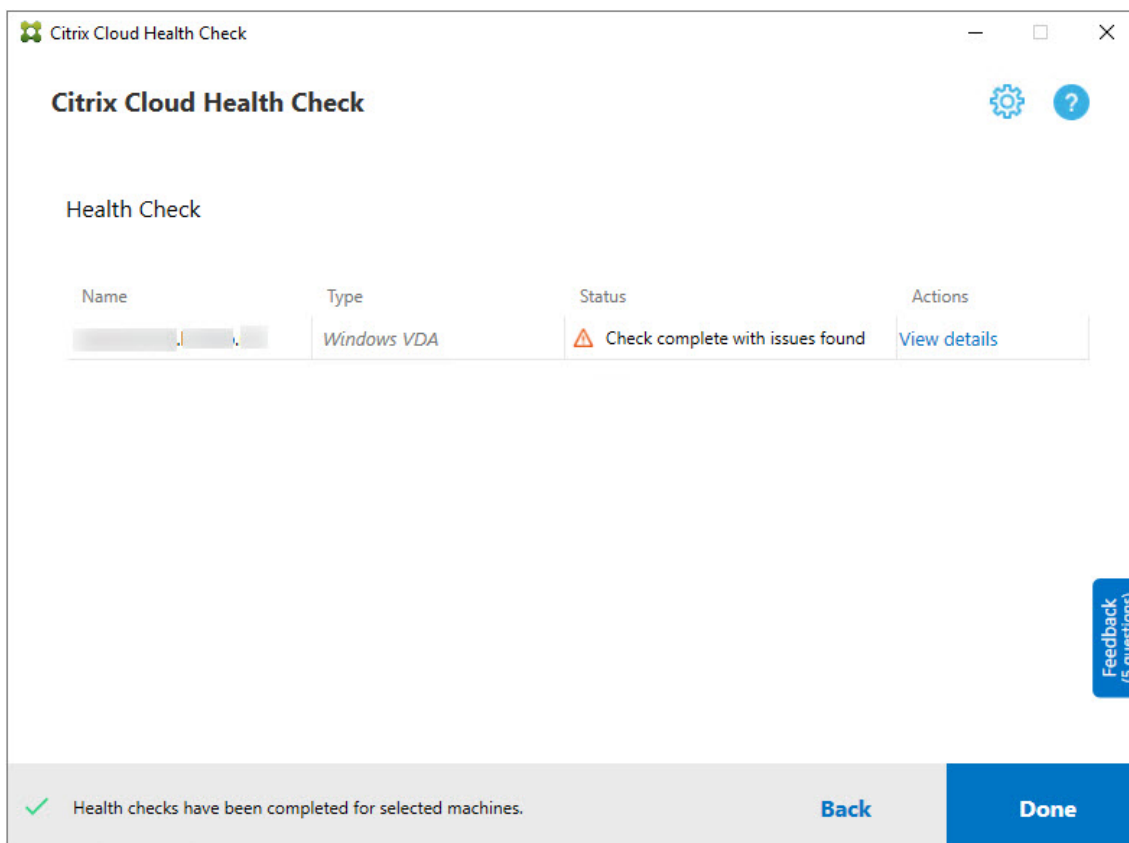


During and after checking, the **Status** column indicates the current checking state for a machine.

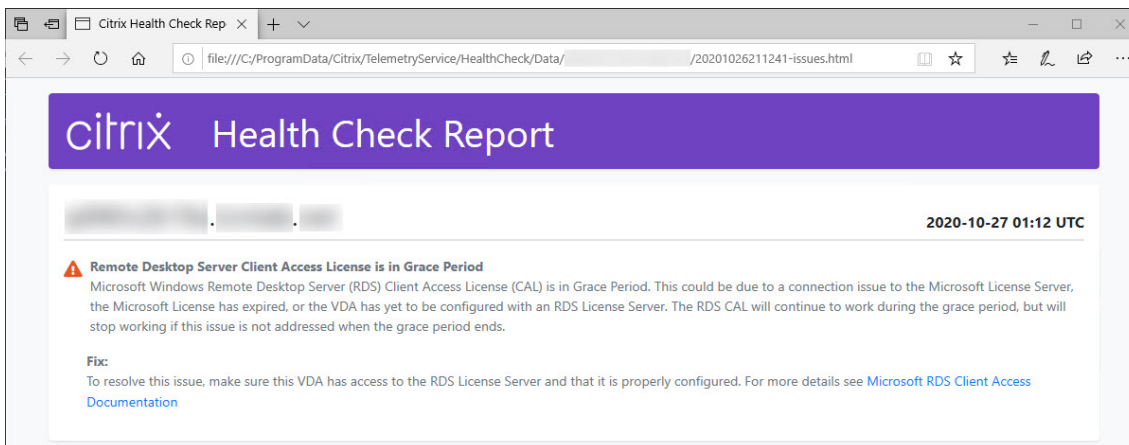
11. To stop all in-progress checks, click **Stop Checking** in the lower right corner of the page. You can't cancel a single machine's health check, you can only cancel the check for all selected machines.



12. When the checks are complete for all selected machines, the **Stop Checking** button in the lower right corner changes to **Done**.



- If a check fails, you can click **Retry** in the **Action** column.
- If a check completes with no issues found, the **Action** column is empty.
- If a check finds issues, click **View Details** to view the results.



If you use Internet Explorer to view the report, you must click **Allow blocked content** to display the hyperlink.



The screenshot shows the Citrix Health Check Report interface. At the top, there is a purple header with the Citrix logo and the text "Health Check Report". Below the header, there are three blurred status indicators. On the right side, the date and time "2020-10-27 01:29 UTC" are displayed. The main content area contains a warning message: "Remote Desktop Server Client Access License is in Grace Period". The text explains that the Microsoft Windows Remote Desktop Server (RDS) Client Access License (CAL) is in Grace Period, which could be due to a connection issue to the Microsoft License Server, an expired license, or a VDA not configured with an RDS License Server. A "Fix:" section provides instructions to ensure the VDA has access to the RDS License Server and is properly configured, with a link to Microsoft RDS Client Access Documentation. At the bottom of the screenshot, a yellow warning bar from Internet Explorer is visible, stating "Internet Explorer restricted this webpage from running scripts or ActiveX controls." and includes an "Allow blocked content" button.

After the check completes for all selected machines, clicking **Back** causes you to lose your check results.

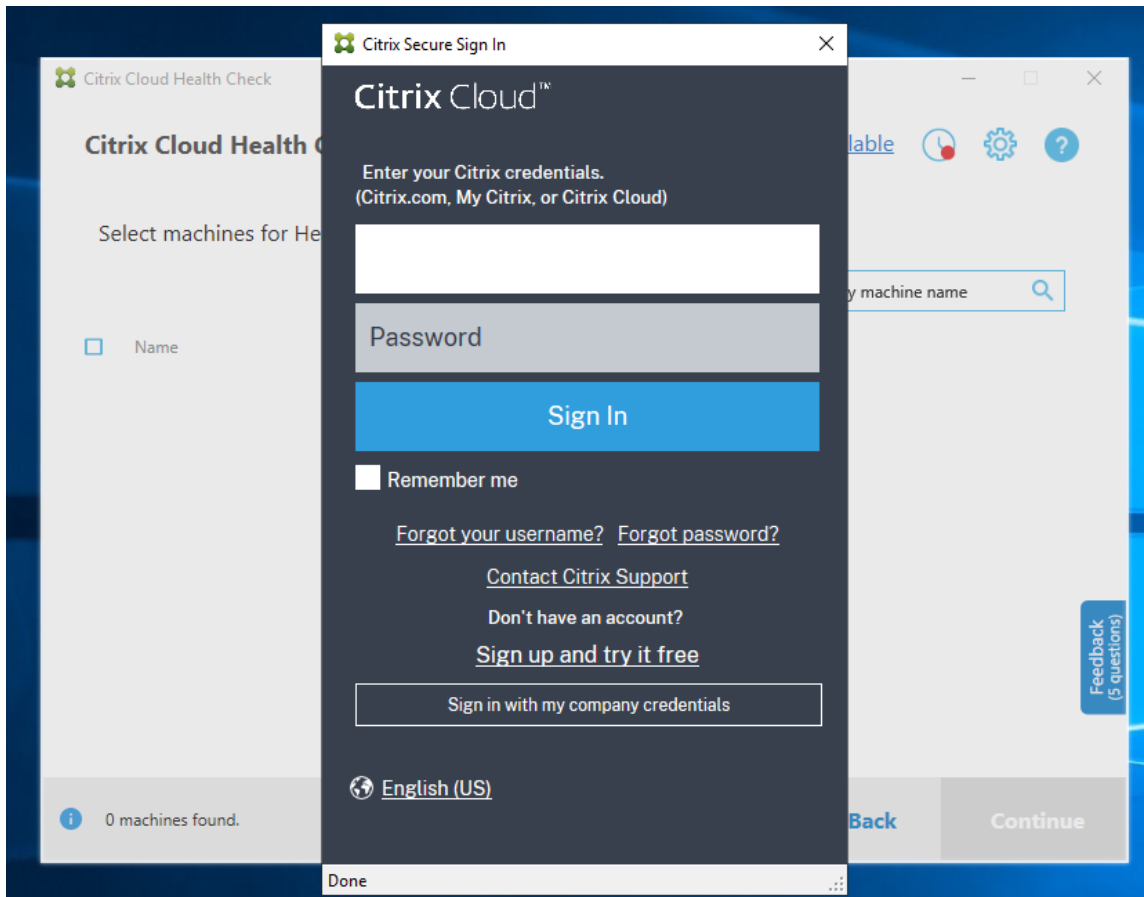
When the checks complete, click **Done** to return to the Cloud Health Check main screen.

### Retrieve VDA machines

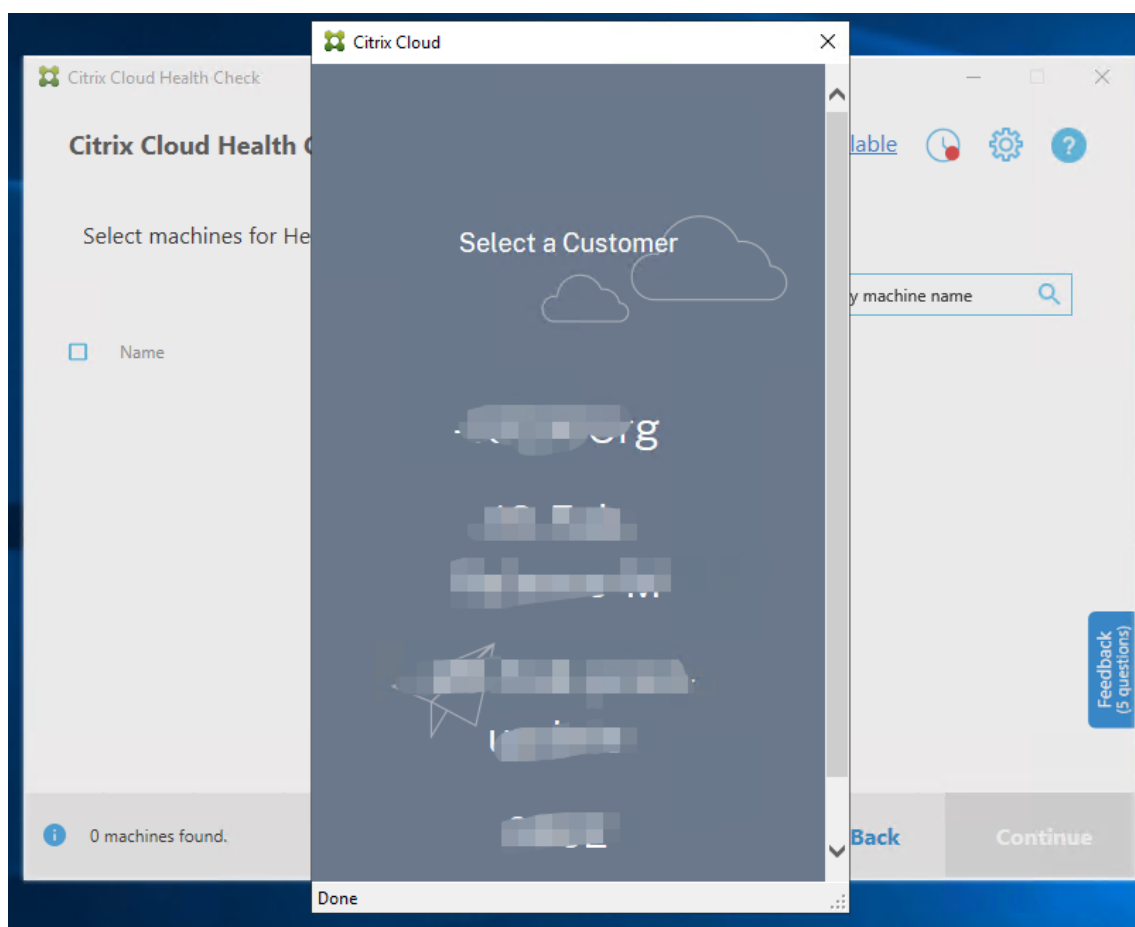
Cloud Health Check can automatically detect and retrieve VDAs from your Citrix Virtual Apps and Desktops service deployments.

To retrieve your VDAs:

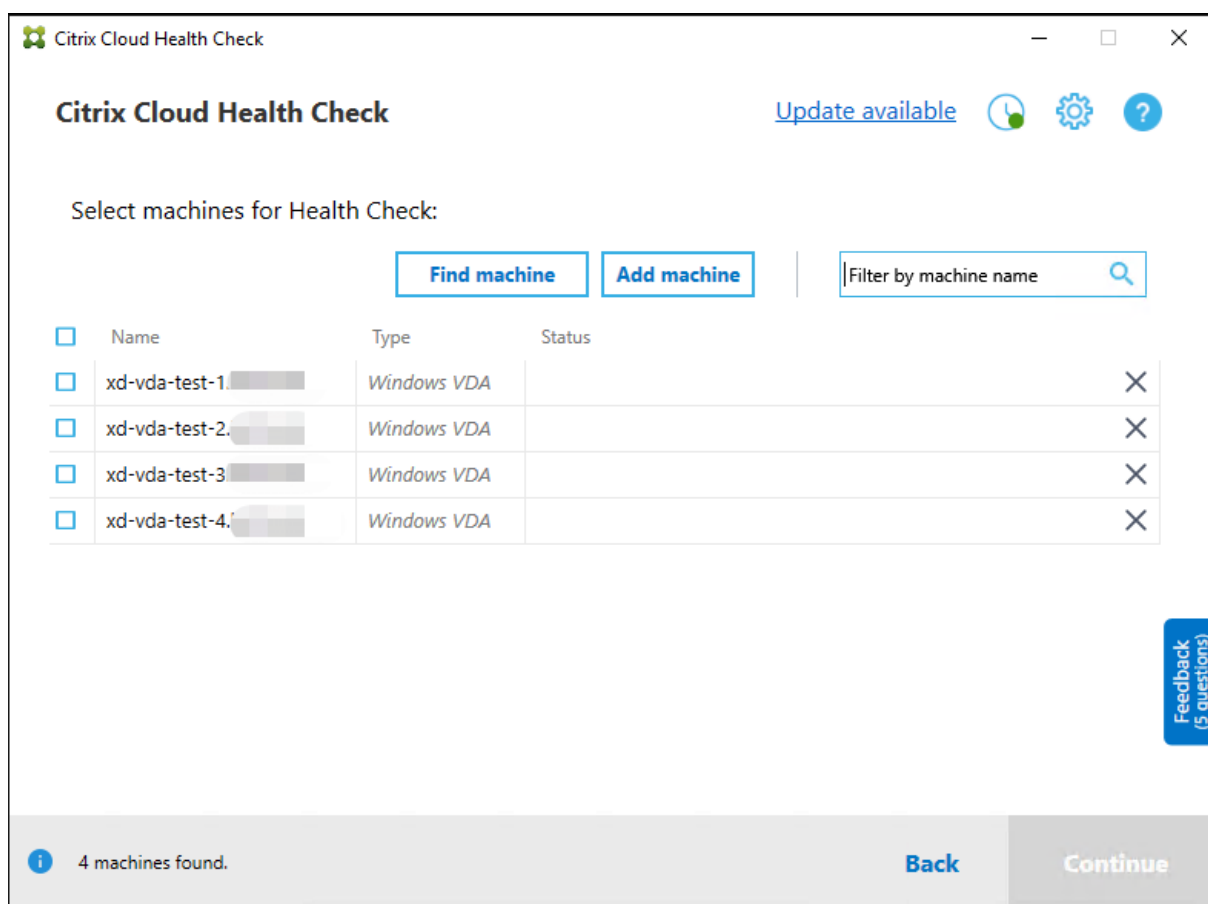
1. Prepare a new machine that is joined to the same domain forest as the machine Cloud Health Check runs on.
2. Open Cloud Health Check and click **Find machine** to sign in to Citrix Cloud.



3. Select the customer with the cloud site you want to retrieve.



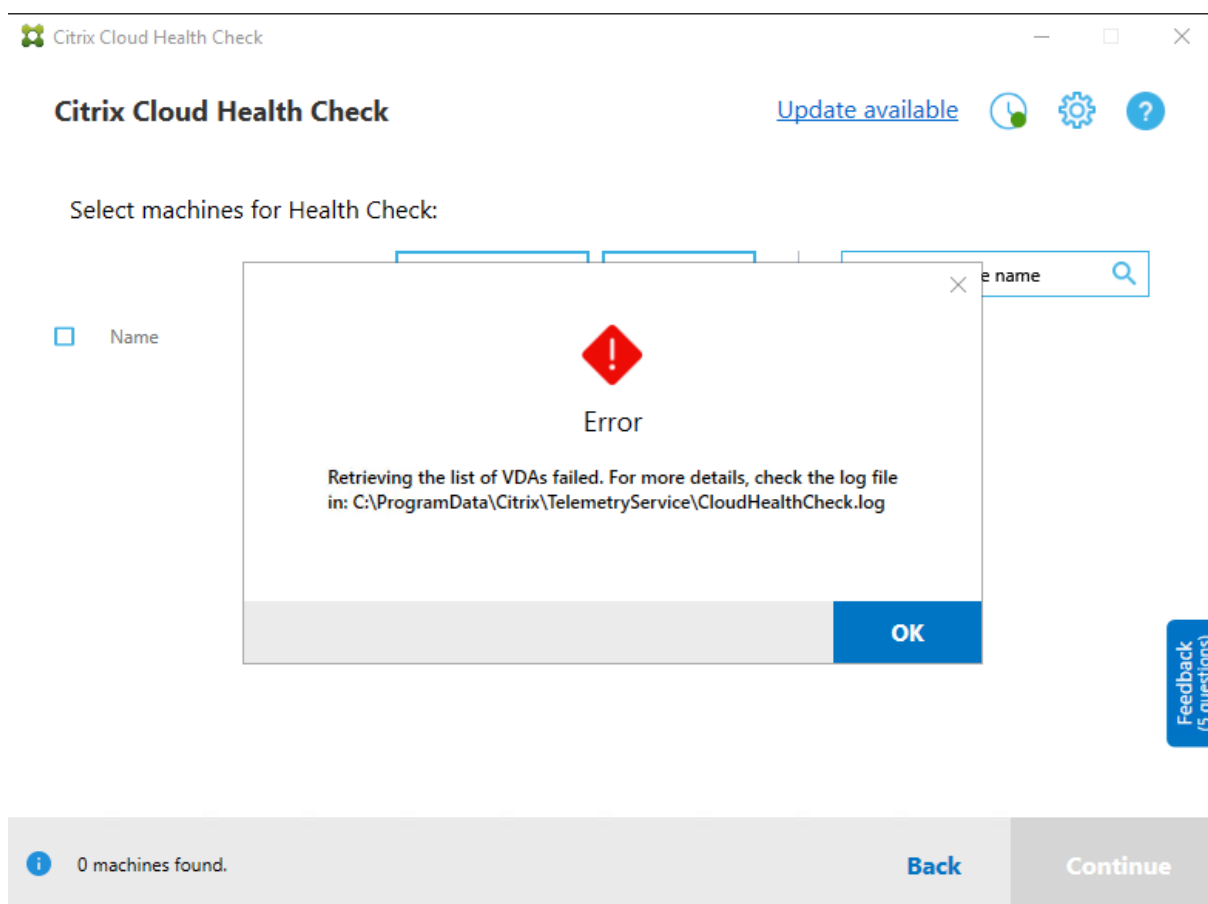
The VDA list displays in Cloud Health Check. The list is also saved in a local file located at `\ProgramData\Citrix\TelemetryService\ChcDiscovery\ChcDiscoveredMachineList.json`.



Your machine list loads the local cache when you open Cloud Health Check again. If you have made any updates in your deployment, you must click **Find machine** to refresh the machine list.

**Note:**

- Cloud Health Check finds machines only in the same domain forest as the machine Cloud Health Check runs on.
- Citrix Cloud sessions expire in one hour. After one hour, you must click **Find machine** again to get the latest VDA list.
- An error message pops up if retrieving the VDA list fails. You can check the details in `C:\ProgramData\Citrix\TelemetryService\CloudHealthCheck.log`.



## Health check results

Health checks that generate reports contain the following elements:

- Time and date when the results report was generated
- FQDNs of the machines that were checked
- Conditions checked on the target machines

## Running Cloud Health Check on the command line

Cloud Health Check can be run on the command line to help customers to perform health checks. To use Cloud Health Check on the command line, you must be an administrator on the machine Cloud Health Check is running on.

### Note:

When using Cloud Health Check on the command line, only one machine can be checked at a time. Only one instance of `CloudHealthCheck.exe` can be run at the same time on the target machine. If you want to check multiple machines, the machines must be checked one by one, by

wrapping the cmdlets in a loop in cmdlet/PowerShell scripts. Any opened UI instance of Cloud Health Check must also be closed.

## Cmdlets

The supported command line cmdlets are:

- **MachineFQDN** - This cmdlet is **mandatory**. This is the fully qualified domain name of the target machine.
- **MachineType** - This cmdlet is optional. The cmdlet value can be the Windows VDA (default value) or StoreFront.
- **ReportName** - This cmdlet is optional. The cmdlet value must be a valid file name on Windows. The default value is **HealthCheckReport**.

Examples:

```
CloudHealthCheck.exe -MachineFQDN machine.domain.local
```

```
CloudHealthCheck.exe -MachineFQDN machine.domain.local -ReportName checkreport
```

### Note:

The parameter names are not case sensitive.

By default, the console output is not shown in the command line console window. You can manually display the output by appending `|more` to the cmdlet.

Example: `CloudHealthCheck.exe -MachineFQDN machine.domain.local|more`

## Multiple machine checks

To run Cloud Health Check on the command line for multiple machines, use the following example:

```
1 @echo off
2 for %%n in (machine1.domain.local,machine2.domain.local,machine3.domain
   .local) do (
3 start /wait CloudHealthCheck.exe -FQDN %%n
4 echo %errorlevel%)
5 <!--NeedCopy-->
```

## Exit codes

Exit codes explain the result of Cloud Health Check checks within the command line. To get the exit code, you must add `start /wait` before the cmdlet.

Example: `start /wait CloudHealthCheck.exe -MachineFQDN machine.domain.local`

Exit codes are:

- 0 - Normal, check completed and passed.
- 1 - Failure, check completed with issues.
- 2 - Error, check not completed with errors.

You can also use the cmdlet `echo %errorlevel%` to get the exit code for the last ran command.

## Reports

Cloud Health Check creates folders with the name of the machine in `HealthCheckDataFolder` for the target machine. An `.html` file and a `.json` file are created on the machine Cloud Health Check is installed on. Health check reports are located in the `HealthCheckDataFolder` in `%ProgramData%\Citrix\TelemetryService\HealthCheck\Data`.

Reports are only created when issues exist on the target machine.

### Note:

The report files are overwritten if the specified report name exists.

Alerts and basic information are stored in the `.json` report.

```
JSON
├── version : 1
├── id : 9547e4ae-022c-4d36-b3a6-77ee61aa72cd
├── siteId : 00000000-0000-0000-0000-000000000000
├── generatedTime : 2020-09-08T06:53:25Z
└── machineReports
    ├── 0
    │   ├── startTime : 2020-09-08T02:53:13.000Z
    │   ├── endTime : 2020-09-08T02:53:23.000Z
    │   ├── fqdn : machine.domain.local
    │   ├── machineType : VDA
    │   └── alerts
    │       ├── 0
    │       │   ├── issueKey : citrix.vda.network.registration-port-unreachable
    │       │   ├── issueUuid : a3547960-fdad-4594-96bd-ebf9c0af7f4a
    │       │   ├── fixRecommendation : To resolve this issue, see [CTX227516](https://support.citrix.com/article/CTX227516)
    │       │   ├── severity : error
    │       │   ├── issueName : Invalid Windows Firewall configuration
    │       │   ├── issueDescription : The following Windows Firewall rules are not enabled on the VDA:* Inbound agent connections on TCP port 80* Outbound Broker connections on TCP port 80 (default)<br>
    │       │   ├── tags : null
    │       │   └── checkNames
    │       │       ├── 0 : VDA Health Check
    │       │       └── HtmlFix : Fix:
    │       └── 1
    │       └── 2
    │       └── 3
    │       └── 4
    └── HtmlReportName : Health Check Report
```

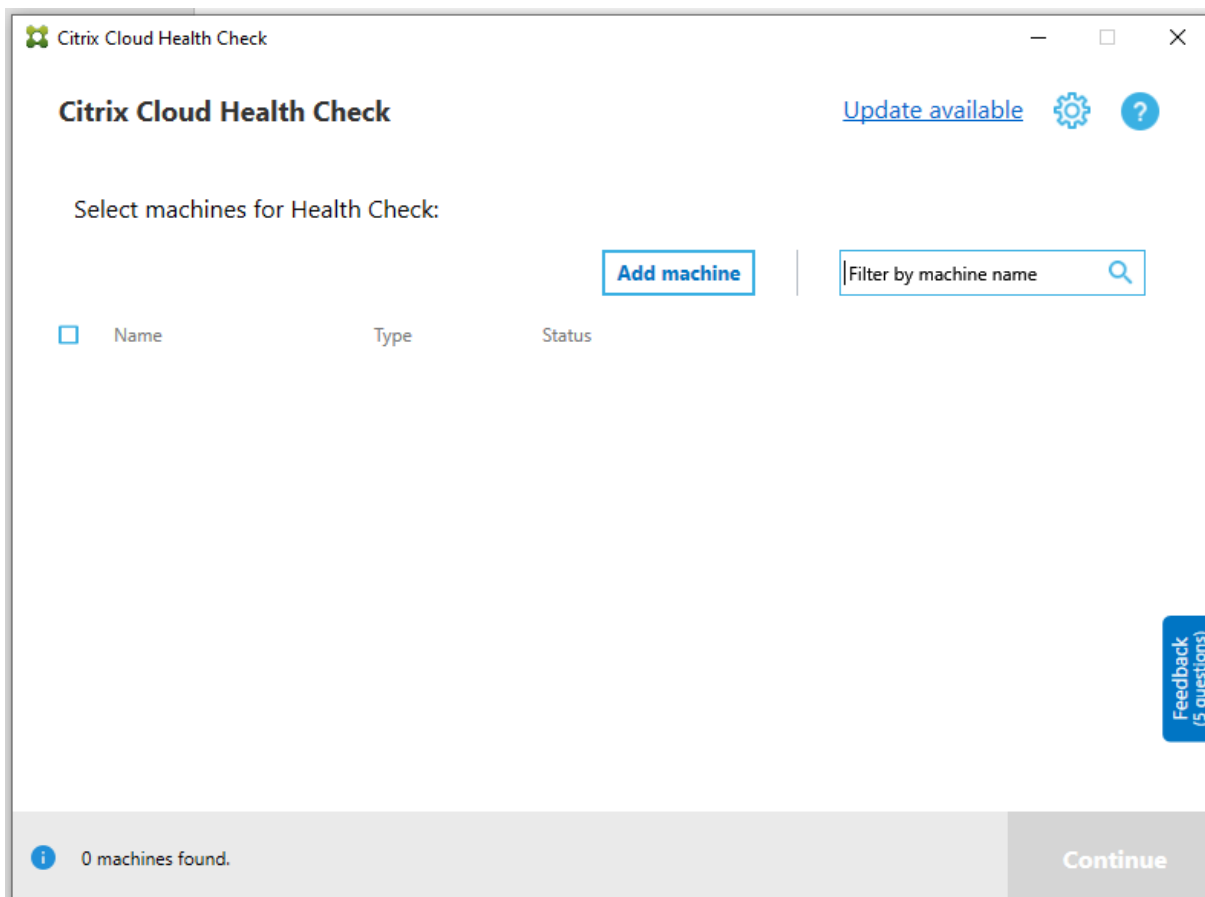
Report codes are:

- **issueKey**: a plain text description of the issue.
- **issueUuid**: a unique identifying string for the issue.
- **fixRecommendation**: the fix recommendation for the issue.
- **severity**: indicates if the issue must be fixed. An error can indicate that the component (VDA or StoreFront) malfunctioned, and a warning indicates that the component can work but might have some potential issues.

- **issueName:** the title of the issue.
- **issueDescription:** a detailed description of the issue.

## Updating Cloud Health Check

If there is a new version of Cloud Health Check available, an Update available link displays on the top right of the Cloud Health Check window. Click the link to go to Citrix Downloads to get the new version.



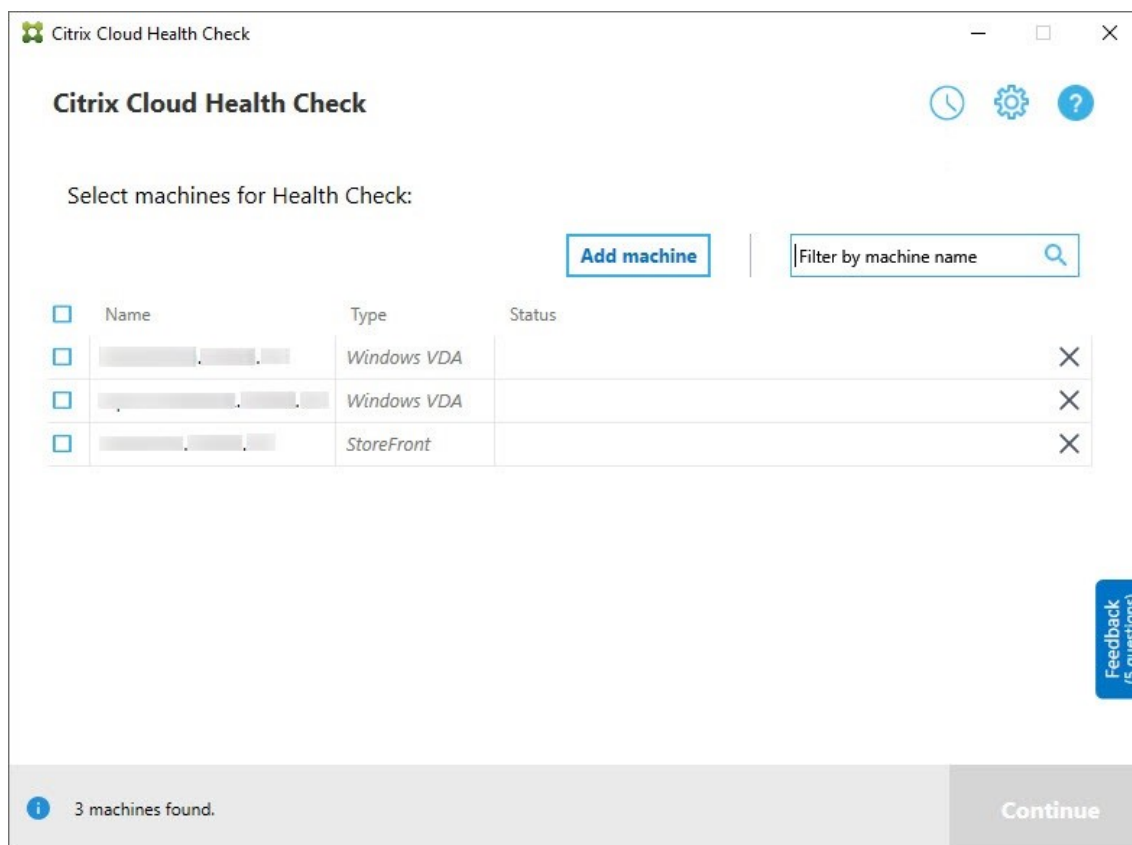
## Cloud Health Check scheduler

Use the Cloud Health Check scheduler to perform periodic health checks.


### Set up the schedule

1. Click **Add machine** in the Cloud Health Check main window to add machines that you want to run periodic checks on.





2. Click the clock icon, then click **Set schedule**.



3. Select a time for your schedule, then click **Next**. The task can be set to repeat by selecting the **Repeat task every** check box.
4. Choose to output results to Windows Event Log. The task can be set to write the results to Windows Event Log.
5. Choose to trigger a custom PowerShell script after the scheduled check finishes, and then click **Next**.
  - Click **Edit** to edit the script content in Windows PowerShell ISE if needed.
  - Click **Locate** to open the file location and to use a different editor to open the file to edit the script.
  - Click **Reset** to reset the script to its original setting.

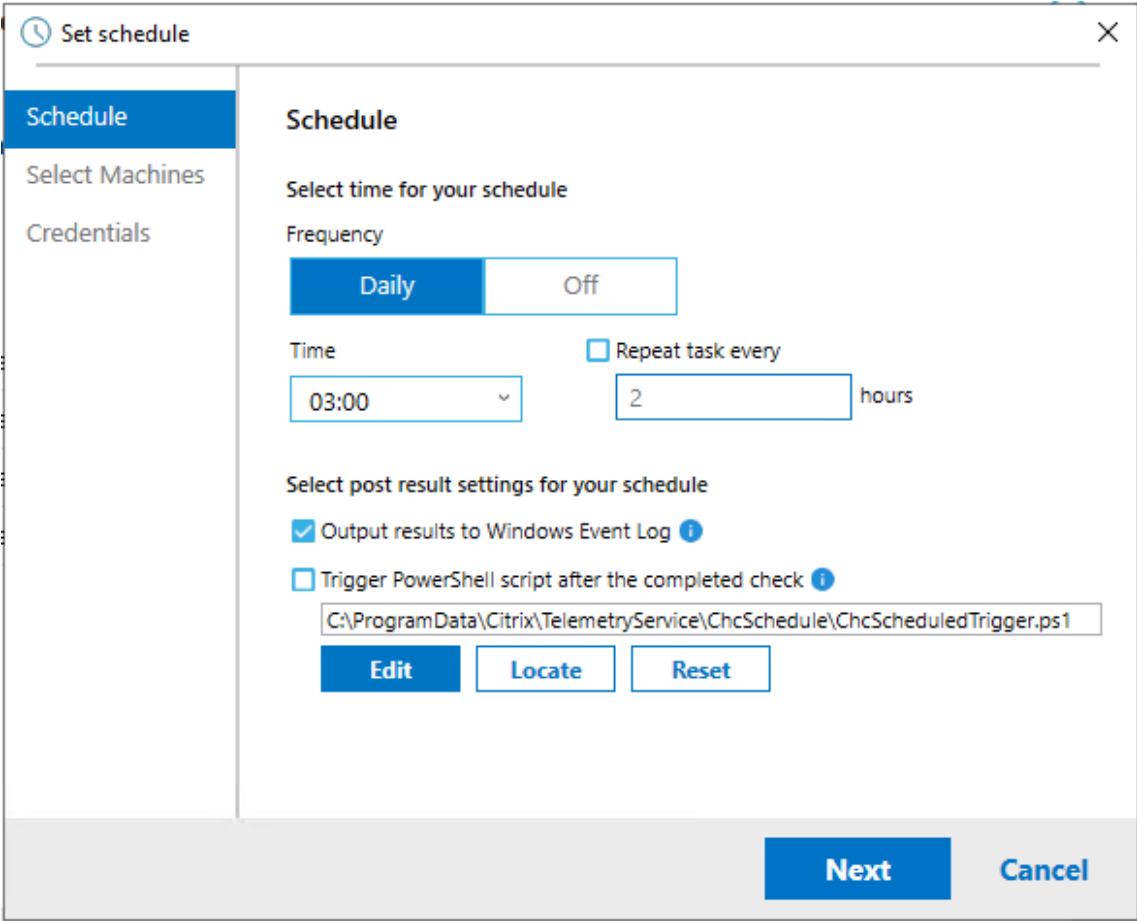
**Note:**

- You cannot change the script name and path for the script.
- You can implement custom actions using the ChcShceduledTrigger.ps1 script, such as sending an email after the scheduled check report is ready. Add the following code to the end of the script. Customize the code to add the correct email accounts and the SMTP server address. An email notification is sent using the


credentials of the account that the scheduled task runs.

```

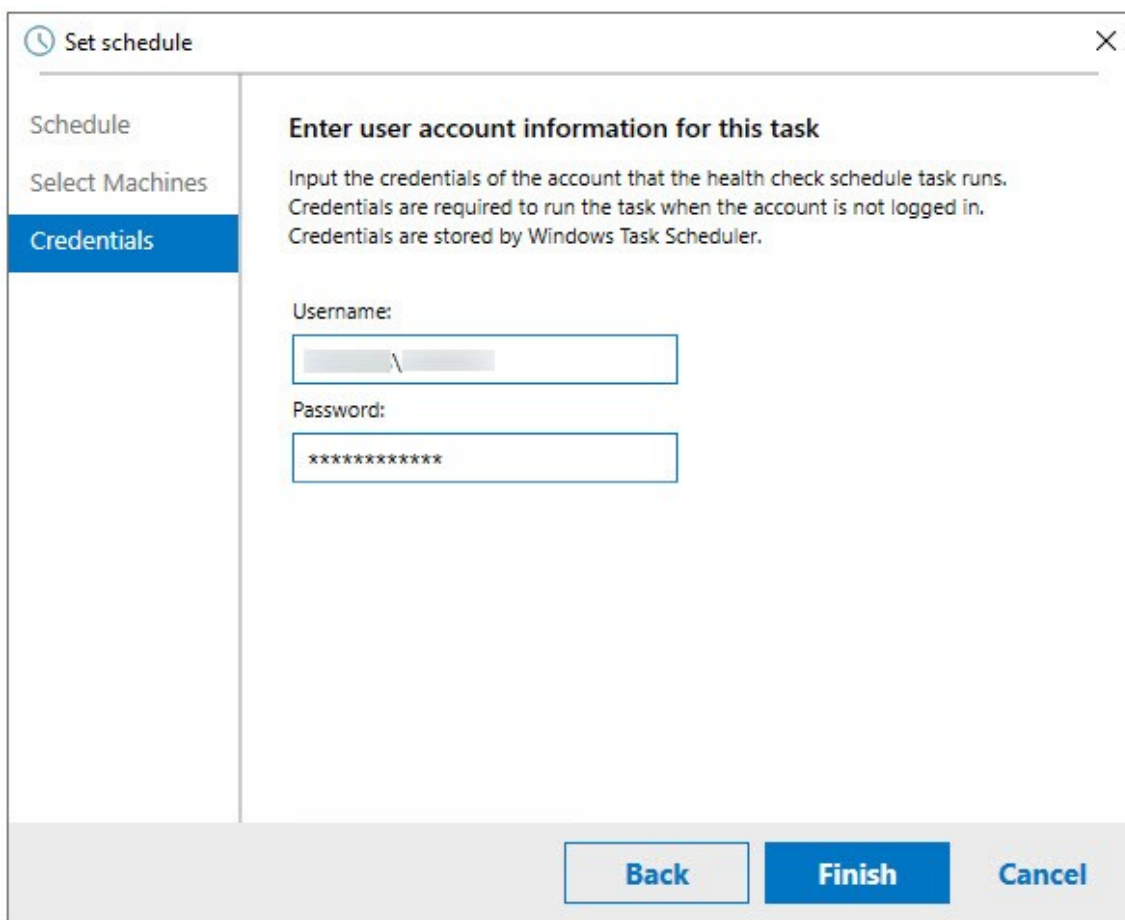
1 #Sending email example code:
2 $body = "CreatedTime: $($report.CreatedTime)"
3 $body = $body + "`nStatusCode: $($report.StatusCode)"
4 $body = $body + "`nMachineCount: $($report.MachineReports.Count)"
5 $from = "mock_email_accout"
6 $to = "mock_email_accout"
7 $smtpServer = "mock_smtp_server"
8
9 Send-MailMessage -Subject "Citrix Cloud Health Check Scheduler
   Report" -Body $body -From $from -To $to -SmtpServer $smtpServer
10 <!--NeedCopy-->
    
```



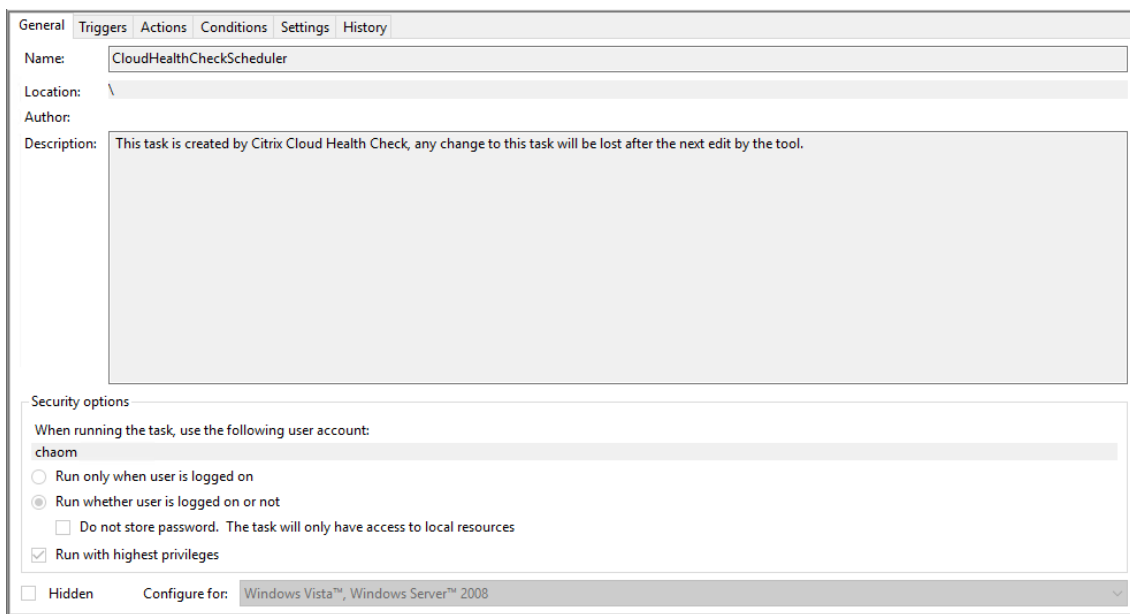
6. Select the machines for your schedule, then click **Next**.



7. Input the credentials of the account that the task runs on, then click **Finish**.

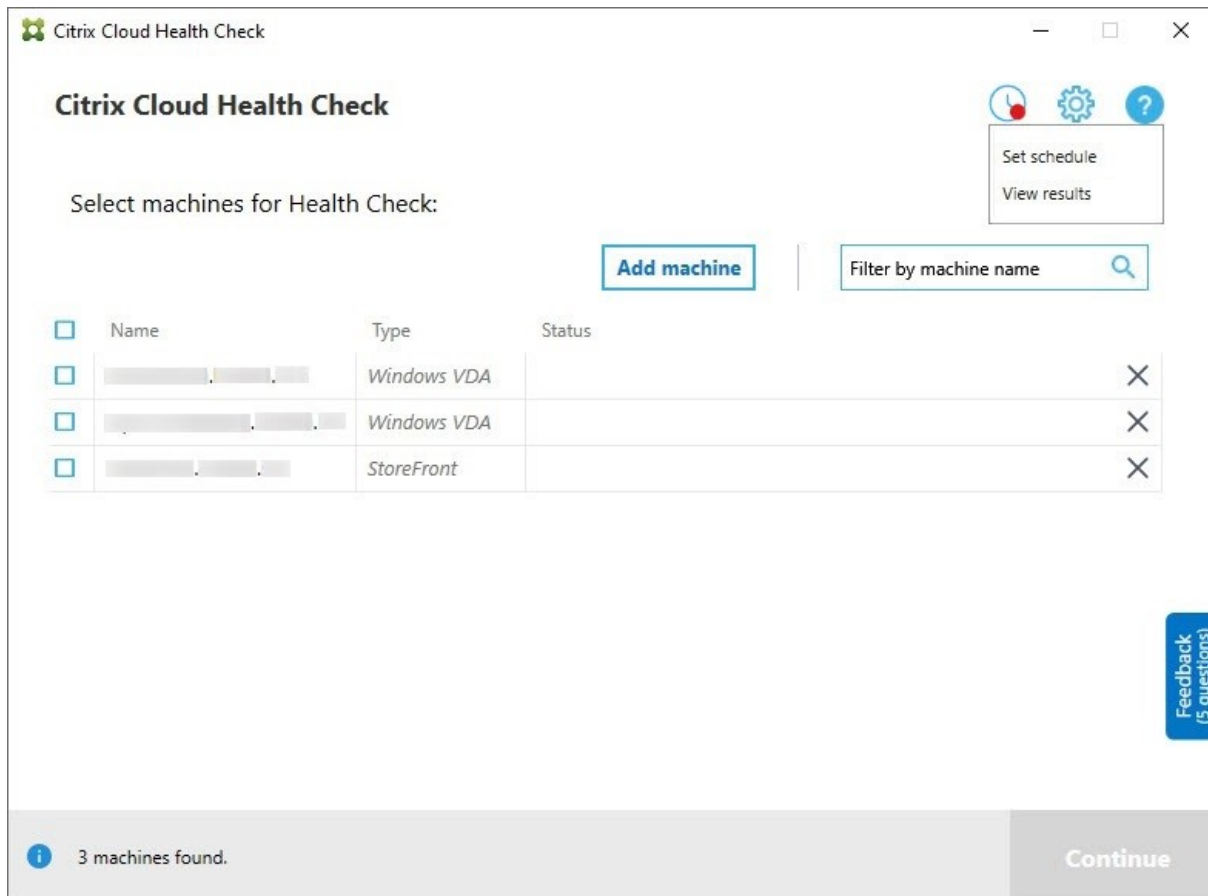


8. A CloudHealthCheckScheduler task is created in Windows Task Scheduler.



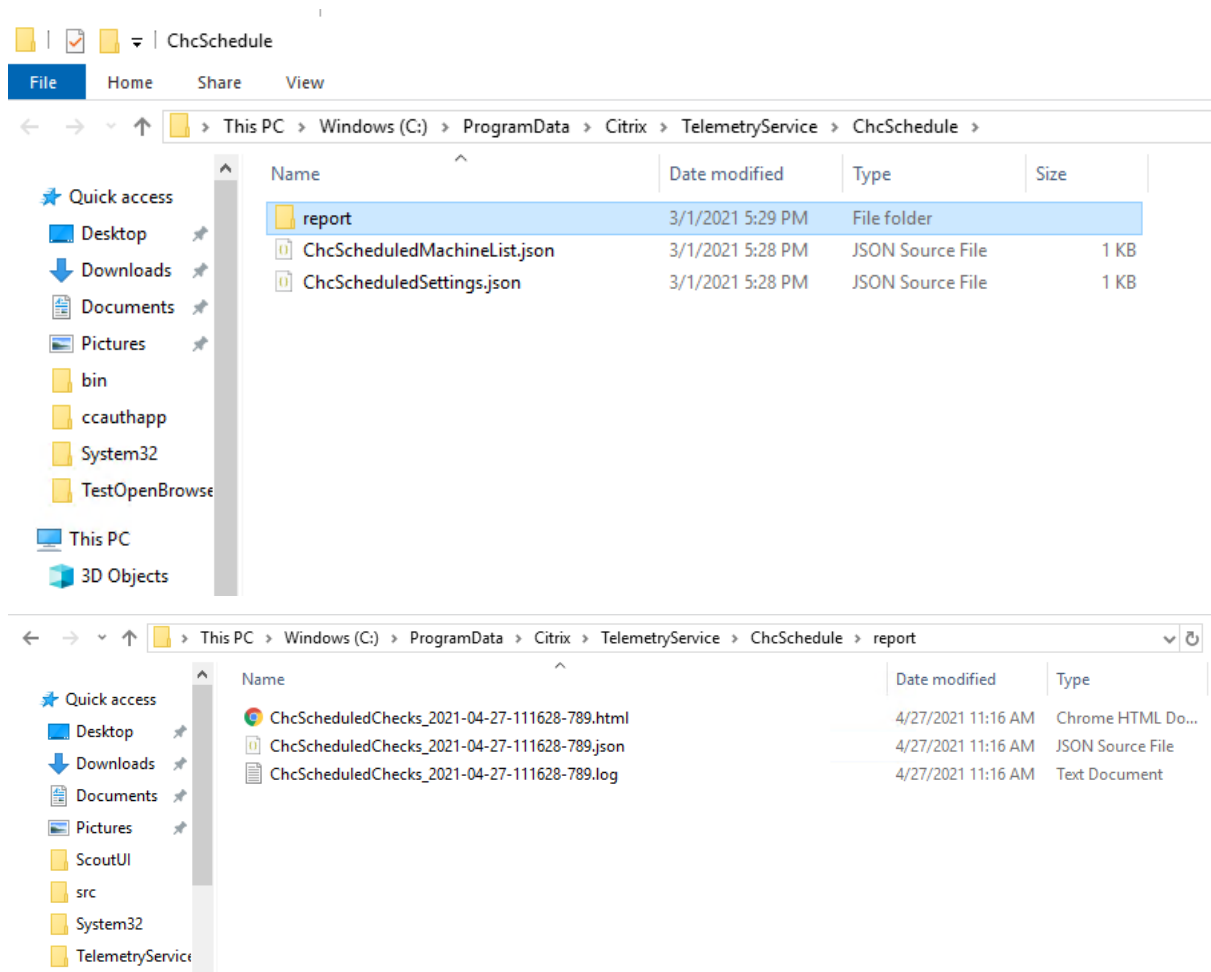
### View schedule results

The clock icon with a red dot indicates that issues were found in the last check. To view the results, click the clock icon, then click **View results**.

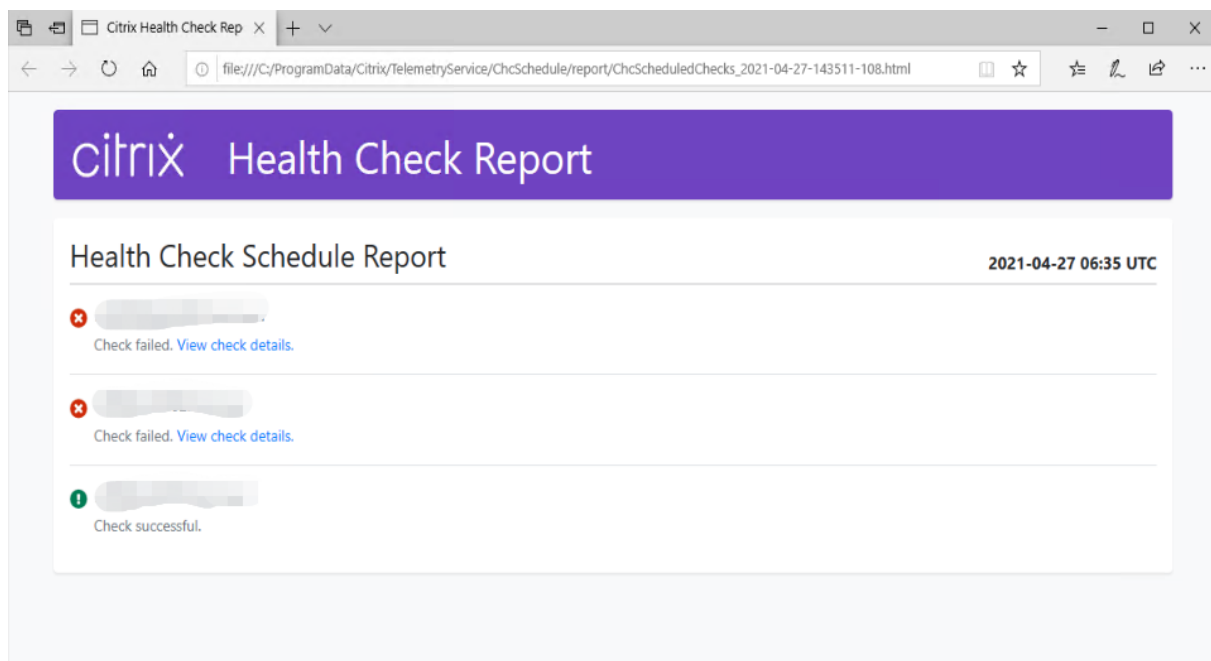


All health check results are stored in a folder called ChcSchedule. Cloud Health Check creates three files during each check run. Up to 500 iteration logs are kept.


## Citrix Virtual Apps and Desktops service



The html report lists the overall report for each schedules. An example of the report is below:

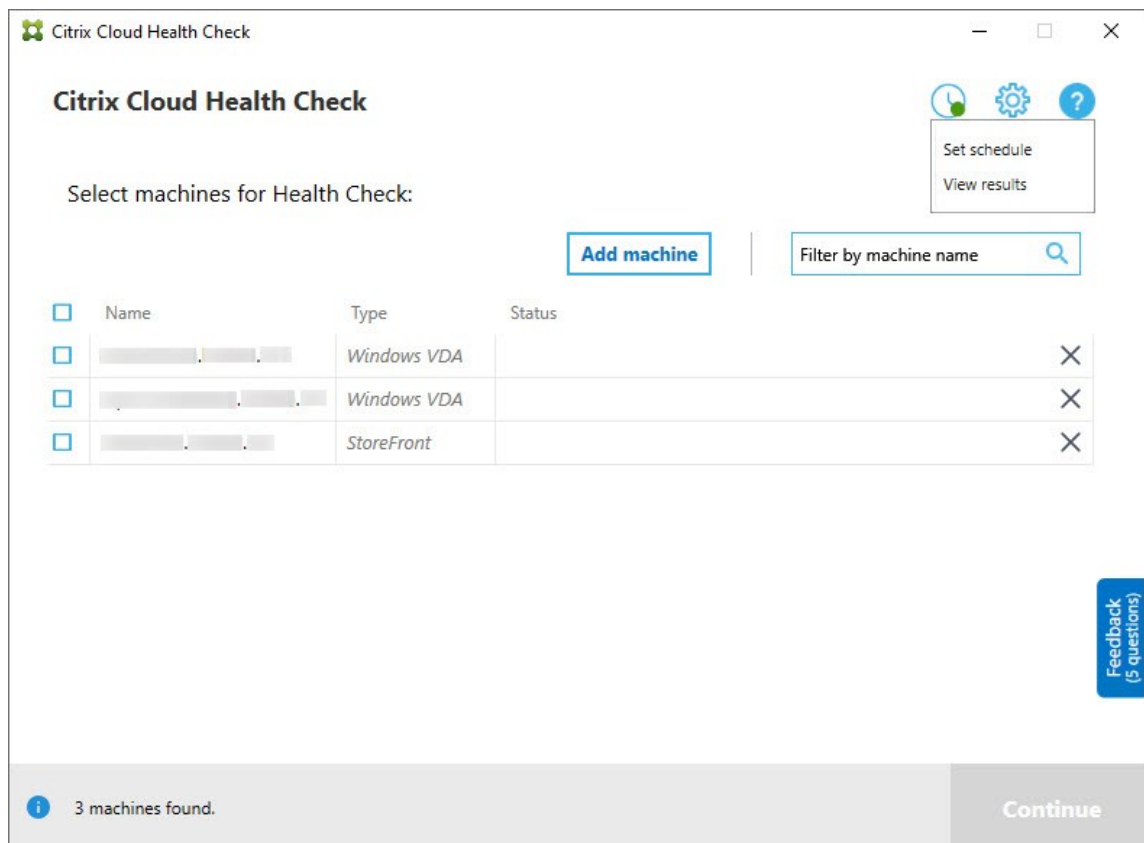


If the **Output results to Windows Event Log** checkbox is selected, the check result is also sent to Windows Event Log.




## Disable schedules

1. Click the clock icon, then click **Set schedule**.



2. Click **Off**, then click **Finish** to disable the scheduler.





### More information

- You must add or import VDAs to Cloud Health Check first. For more information, see [Import VDA Machines](#).
- The Cloud Health Check scheduler can only schedule one task at a time on a domain-joined machine. If you set the schedule multiple times, only the latest one takes effect.

### Verification tests

Before a health check starts, verification tests run automatically for each selected machine. These tests make sure that the requirements are met for a health check to run. If a test fails for a machine, Cloud Health Check displays a message with suggested corrective actions.

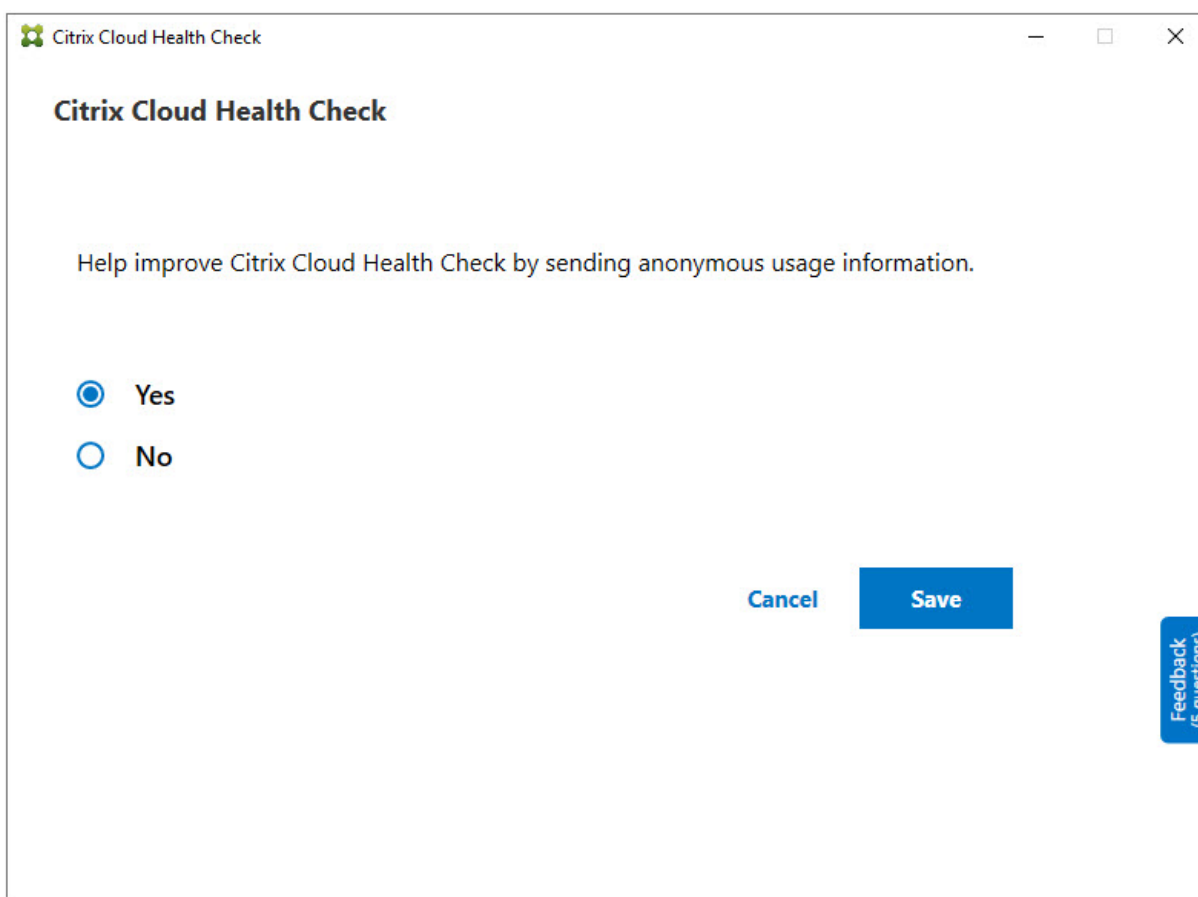
- **Cloud Health Check cannot reach this machine** - Ensure that:
  - The machine is powered on.
  - The network connection is working properly. (This can include verifying that your firewall is properly configured.)
  - File and printer sharing is turned on. See the Microsoft documentation for instructions.

- **Enable PSRemoting and WinRM** - You can enable PowerShell remoting and WinRM by running PowerShell as an administrator, then running the Enable-PSRemoting cmdlet. For details, see the Microsoft help for the cmdlet.
- **Cloud Health Check requires PowerShell 3.0 or later** - Install PowerShell 3.0 or later on the machine, and then enable PowerShell remoting.
- **WMI is not running on the machine** - Ensure that Windows Management Instrumentation (WMI) access is enabled.
- **WMI connections blocked** - Enable WMI in the Windows Firewall service.

## Usage data collection

When you use Cloud Health Check, Citrix uses Google Analytics to collect anonymous usage data to be used for future product features and improvements. Data collection is enabled by default.

To change usage data collection and upload, click the **Settings** gear in the Cloud Health Check UI. You can then choose whether to send the information by selecting **Yes** or **No** and then clicking **Save**.



The screenshot shows a window titled "Citrix Cloud Health Check" with a standard Windows title bar (minimize, maximize, close). The main content area has the heading "Citrix Cloud Health Check" and the text "Help improve Citrix Cloud Health Check by sending anonymous usage information." Below this text are two radio button options: "Yes" (which is selected) and "No". At the bottom right of the dialog are two buttons: "Cancel" and "Save". On the far right edge, there is a vertical blue button labeled "Feedback (5 questions)".

## Automatic fix (preview)

**Note:**

Automatic fix is available as a preview.

Automatic fix allows Cloud Health Check to automatically detect and fix certain issues by changing the settings or restarting the services. This feature works only on local machines where Cloud Health Check is running.

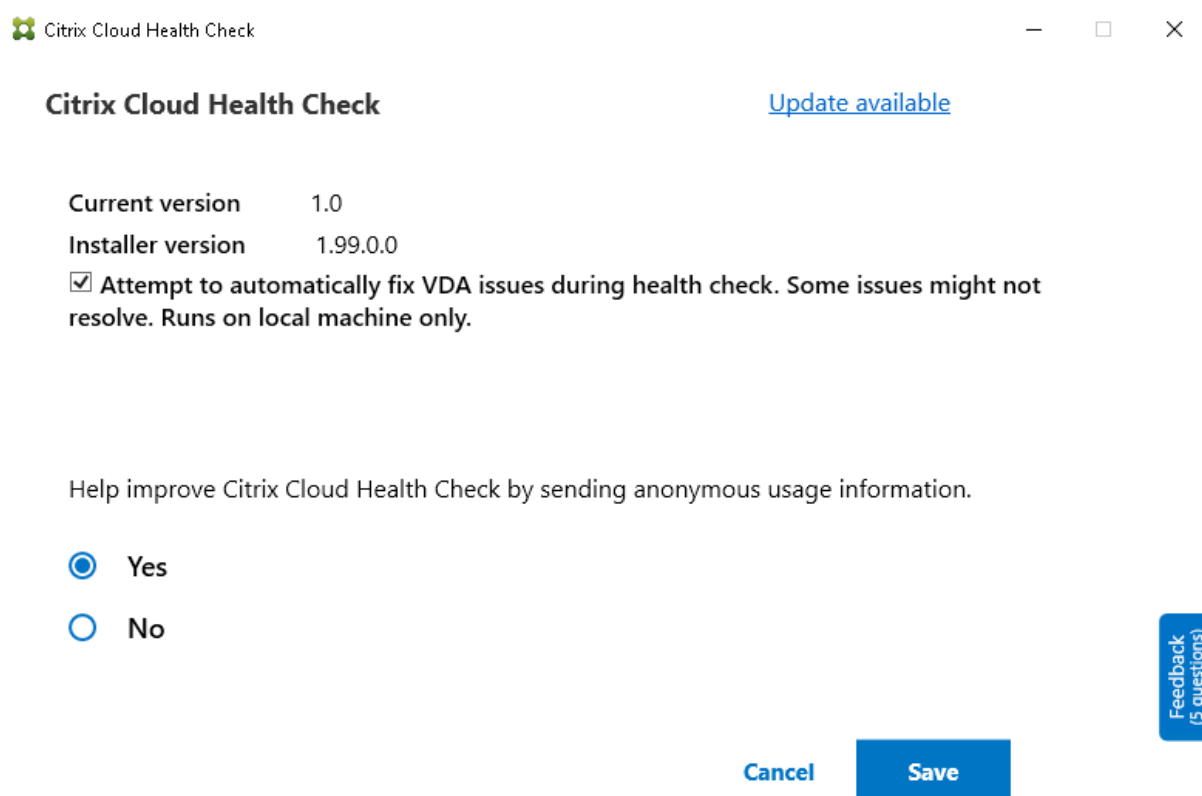
Automatic fix checks the following VDA registration items, with the recommended fixes:

- VDA machine domain membership
  - Fix: Test connection security channel with a “repair” model to fix
- VDA services status
  - Fix: Restart BrokerAgent service
- Communication with Controller
  - Fix: Restart BrokerAgent service
- Time sync with Controller
  - Fix: Run W32tm command

For session launches, automatic fix checks the following item, with the recommended fix:

- Session launch service status
  - Fix: Restart BrokerAgent service

This feature is enabled by default. To disable it, click the gear icon in the upper right corner of the Cloud Health Check main window and then clear **Attempt to automatically fix VDA issues during health check**.



## Troubleshooting

When Cloud Health Check fails to run or any exception occurs, check the Cloud Health Check log in `C:\ProgramData\Citrix\TelemetryService\CloudHealthCheck.log`.

The Cloud Health Check log for each target machine is in `C:\ProgramData\Citrix\TelemetryService\HealthCheck\Data\${TargetMachineFQDN}\log.txt`.

To enable the debug log:

Edit `C:\Program Files\Citrix\CloudHealthCheck\CloudHealthCheck.exe.config`, update `<add name="TraceLevelSwitch" value="3"/>` to `<add name="TraceLevelSwitch" value="4"/>`, save the file and reopen Cloud Health Check.

## Feedback

To leave feedback on Cloud Health Check, fill out the [Citrix survey](#).

## Configuration logging

July 26, 2021

### Note:

Configuration log content displays only in English, regardless of which language you select for your Citrix Cloud account.

Configuration logging is a feature that captures Citrix Virtual Apps and Desktops deployment configuration changes and administrative activities to a logging database in Citrix Cloud. You can use the logged content to:

- Diagnose and troubleshoot problems after configuration changes are made. The log provides a breadcrumb trail.
- Assist change management and track configurations.
- Report administrative activities.

In this Citrix Cloud service, configuration logging is always enabled. You cannot disable it.

From the Full Configuration management interface, you can view configuration log content, filtered by date ranges or by full text search. You can also generate a CSV report using PowerShell. From this console, you cannot edit or delete log content. You can use the Remote PowerShell SDK to schedule periodic data deletion from the log.

Configuration logs are localized when they are created. For example, a log created in English is read in English, regardless of the locale of the reader.

Permissions required (see [Delegated administration](#)):

- Full Administrators in Citrix Cloud, plus Citrix Virtual Apps and Desktops service Cloud Administrators and Read Only Administrators can view configuration logs in the **Manage** console.
- Full Administrators and Cloud Administrators can also download a CSV report of logging activity, using PowerShell.

### What is logged

The following operations are logged:

- Configuration changes and administrative activities initiated from the **Manage** and **Monitor** tabs
- PowerShell scripts
- REST API requests

**Note:**

You cannot see log entries for Citrix Cloud platform internal operations, such as database setup and management.

Examples of logged configuration changes include working with (creating, editing, deleting, assigning):

- Machine catalogs
- Delivery groups (including changing power management settings)
- Administrator roles and scopes
- Host resources and connections
- Citrix policies through the **Manage** console

Examples of logged administrative changes include:

- Power management of a virtual machine or a user desktop
- Manage or monitor functions sending a message to a user

The following operations are not logged. (Many of them are not available to customer administrators.)

- Automatic operations such as pool management power-on of virtual machines.
- Policy actions implemented through the Group Policy Management Console (GPMC). Use Microsoft tools to view logs of those actions.
- Changes made through the registry or from sources other than the Full Configuration management interface, Monitor, or PowerShell.

## View configuration log content

To view configuration log content, follow these steps:

1. Sign in to [Citrix Cloud](#). Select **My Services > Virtual Apps and Desktops** in the upper left menu.
2. From **Manage > Full Configuration**, select **Logging > Events** in the left pane.

By default, the display in the center pane lists the log content chronologically (newest entries first), separated by date. You can:

- Sort the display by column heading.
- Filter the display by specifying a day interval, or entering text in the Search box. To return to the standard display after using search, clear the text in the Search box.

Display characteristics:

- High-level operations created during management and monitoring are listed in the upper middle pane. A high-level operation results in one or more services and PowerShell SDK calls, which are low-level operations. When you select a high-level operation in the upper middle pane, the lower pane displays the low-level operations.

- If you create a low-level operation in PowerShell without specifying a parent high-level operation, configuration logging creates a surrogate high-level operation.
- If an operation fails before completion, the log operation might not be completed in the database. For example, a start record has no corresponding stop record. In such cases, the log indicates that there is missing information. When you display logs based on time ranges, incomplete logs are shown if the data in the logs matches the criteria. For example, if you request logs for the last five days, and a log with a start time in the last five days has no end time, it is included.
- Remember: You cannot see log entries for Citrix Cloud platform internal operations, such as database setup and management.

### View tasks related to machine catalog operations

To view tasks related to machine catalog operations, navigate to **Manage > Full Configuration > Logging > Tasks**. The **Tasks** tab displays only tasks related to catalogs created through Machine Creation Services (MCS) or Provisioning Services (PVS). Specifically, tasks associated with the following machine catalog operations appear:

- Create catalogs
- Clone catalogs
- Add machines
- Remove machines
- Update a catalog (update images or machines)
- Roll back machine updates

#### Tip:

The **Tasks** tab displays only tasks related to provisioning scheme changes (creating or modifying a provisioning scheme).

A task can be in the following state:

- Completed
- Not started
- Running
- Canceled
- Failed
- Unknown

To cancel a running task, select the task and then click **Cancel**. The cancelation takes some time to complete.

Examples of logged tasks include:

- Image update completed for a certain catalog

- Error while updating the image for a certain catalog
- Canceled image update for a certain catalog
- Provisioning VMs to a certain catalog
- Removing VMs from a certain catalog
- Created a certain catalog

By default, the display in the center pane lists logged tasks chronologically (newest entries first), separated by date. You can sort the display by column heading. To clear completed tasks, click **Clear Completed Tasks** under the **Tasks** tab.

## View API logs

To view REST API logs, navigate to **Manage > Full Configuration > Logging > APIs**. The **APIs** tab displays REST API requests made during a certain time period.

Be aware of the following considerations:

- REST API logs are cleared after you sign out of the console. (They are also cleared if you refresh your browser window.)
- Any operations in the console that result in API calls will have their corresponding API requests displayed on the **APIs** tab.
- The display lists the API requests chronologically (newest entries first), separated by date. The maximum number of API requests in the display is 1,000.

## Generate reports

To generate a CSV or HTML report containing configuration log data, use PowerShell cmdlets for the ConfigLogging Service in the Citrix Virtual Apps and Desktops Remote PowerShell SDK. For details, see:

- [Export-LogReportCsv](#)
- [Export-LogReportHtml](#)

## Schedule periodic data deletion

Use the Remote PowerShell SDK to specify how long data is retained in the configuration logging database. (This feature is not available in the Full Configuration management interface.) In the Citrix Virtual Apps and Desktops service, you must have Full access.

In the `Set-LogSite` cmdlet, the `-LoggingDBPurgeDurationDays` parameter specifies how many days data is retained in the configuration logging database before it's deleted automatically.

- By default, the value of this parameter is 0. A zero value means that data in the configuration logging database is never deleted automatically.



- When you set a non-zero value, the database is checked once every 120 minutes. Data older than the retention period is deleted.

Use `Get-LogSite` to view the current value of the parameter.

## Differences from on-premises Citrix Virtual Apps and Desktops

If you're familiar with configuration logging in the on-premises Virtual Apps and Desktops product, the Citrix Cloud version has several differences. In Citrix Cloud:

- Configuration logging is always enabled. You cannot disable it. Mandatory logging is not available.
- You cannot change the location of the configuration logging database, because the database is managed in the Citrix Cloud platform.
- Configuration log displays do not include operations and activities that are performed within the Citrix Cloud platform.
- You can create a CSV or HTML report of logged operations using PowerShell cmdlets only. In the on-premises product, reports can be generated from Studio or PowerShell.
- You cannot delete configuration log content.

## Delegated administration

July 7, 2021

### Overview

With delegated administration in Citrix Cloud, you can configure the access permissions that all of your administrators need, in accordance with their role in your organization.

By default, administrators have full access. This setting enables access to all available customer administration and management functions in Citrix Cloud, plus all subscribed services. To tailor an administrator's access:

- Configure custom access for an administrator's general management permissions in Citrix Cloud.
- Configure custom access for subscribed services. In the Citrix Virtual Apps and Desktops service, you can configure custom access when you invite a new administrator. You can change an administrator's access later.

For information about displaying the list of administrators and defining access permissions, see [Add administrators to a Citrix Cloud account](#).

This article describes how to configure custom access in the Citrix Virtual Apps and Desktops service.

## Administrators, roles, and scopes

Delegated administration uses three concepts for custom access: administrators, roles, and scopes.

- **Administrators:** An administrator represents a person identified by their Citrix Cloud sign-in, which is typically an email address. Each administrator is associated with one or more role and scope pairs.
- **Roles:** A role represents a job function, and has permissions associated with it. These permissions allow certain tasks that are unique to the service. For example, the Delivery Group Administrator role has permission to create a delivery group and remove a desktop from a delivery group, plus other associated permissions. An administrator can have multiple roles. An administrator might be a Delivery Group Administrator and a Machine Catalog Administrator.

The service offers several built-in custom access roles. You cannot change the permissions within these built-in roles, or delete those roles.

You can create your own custom access roles to meet your organization's requirements, and delegate permissions with more detail. Use custom roles to allocate permissions at the granularity of an action or task. You can delete a customized role only if it is not assigned to an administrator.

You can change which roles an administrator has.

A role is always paired with a scope.

- **Scopes:** A scope represents a collection of objects. Scopes are used to group objects in a way that is relevant to your organization. Objects can be in more than one scope.

There is one built-in scope: All, which contains all objects. Citrix Cloud and Help Desk administrators are always paired with the All scope. That scope cannot be changed for those administrators.

When you invite (add) an administrator for this service, a role is always paired with a scope (by default, the All scope).

You create and delete scopes in the service's **Manage > Full Configuration** interface. You assign role/scope pairs in the Citrix Cloud console.

A scope is not shown for Full access administrators. By definition, those administrators can access all customer-managed Citrix Cloud and subscribed services objects.

## Built-in roles and scopes

The service has the following built-in roles.

- **Cloud Administrator:** Can perform all tasks that can be initiated from the service.  
Can see the **Manage** and **Monitor** tabs in the console. This role is always combined with the All scope. You cannot change the scope.  
Do not be confused by this role's name. A custom access Cloud Administrator cannot perform Citrix Cloud-level tasks (Citrix Cloud tasks require Full access).
- **Read Only Administrator:** Can see all objects in the specified scopes (in addition to global information), but cannot change anything. For example, a Read Only Administrator with a scope of London can see all global objects and any objects in the London scope (for example, London Delivery Groups). However, that administrator cannot see objects in the New York scope (assuming that the London and New York scopes do not overlap).  
Can see the **Manage** tab in the console. Cannot see the **Monitor** tab. You can change the scope.
- **Help Desk Administrator:** Can view delivery groups, and manage the sessions and machines associated with those groups. Can see the machine catalog and host information for the delivery groups being monitored. Can also perform session management and machine power management operations for the machines in those delivery groups.  
Can see the **Monitor** tab in the console. Cannot see the **Manage** tab. This role is always combined with the All scope. You cannot change the scope.
- **Machine Catalog Administrator:** Can create and manage machine catalogs and provision the machines into them. Can manage base images and install software, but cannot assign applications or desktops to users.  
Can see the **Manage** tab in the console. Cannot see the **Monitor** tab. You can change the scope.
- **Delivery Group Administrator:** Can deliver applications, desktops, and machines. Can also manage the associated sessions. Can manage application and desktop configurations such as policies and power management settings.  
Can see the **Manage** tab in the console. Cannot see the **Monitor** tab. You can change the scope.
- **Host Administrator:** Can manage host connections and their associated resource settings. Cannot deliver machines, applications, or desktops to users.  
Can see the **Manage** tab in the console. Cannot see the **Monitor** tab. You can change the scope.
- **Session Administrator:** Can view delivery groups being monitored and manage their associated sessions and machines.  
Can see the **Monitor** tab in the console. Cannot see the **Manage** tab. You cannot change the scope.
- **Full Administrator:** Can perform all tasks and operations. A full administrator is always combined with **All scope**.

Can see the **Manage** and **Monitor** tabs in the console. This role is always combined with **All scope**. You cannot change the scope.

- **Full Monitor Administrator:** Has full access to all views and commands on the **Monitor** tab.

Can see the **Monitor** tab in the console. Cannot see the **Manage** tab. You cannot change the scope.

- **Probe Agent Administrator:** Has access to Probe Agent APIs.

Can see the **Monitor** tab in the console. Cannot see the **Manage** tab. Has read-only access to the **Applications** page but cannot access any other views.

The following table summarizes which console tabs are visible for each custom access role in the service, and whether the role can be used with custom scopes.

Custom access administrator role	Can see <b>Manage</b> tab in console?	Can see <b>Monitor</b> tab in console?	Can role be used with custom scopes?
Cloud Administrator	Yes	Yes	No
Read Only Administrator	Yes	No	Yes
Help Desk Administrator	No	Yes	No
Machine Catalog Administrator	Yes	No	Yes
Delivery Group Administrator	Yes	No	Yes
Host Administrator	Yes	No	Yes
Session Administrator	No	Yes	No
Full Administrator	Yes	Yes	No
Full Monitor Administrator	No	Yes	No
Probe Agent Administrator	No	Yes	No

**Note:**

Custom access administrator roles (except Cloud Administrator and Help Desk Administrator) are not available for Citrix Virtual Apps and Desktops Standard for Azures, Virtual Apps Essentials, and Virtual Desktops Essentials.

To view the permissions associated with a role:

1. Sign in to [Citrix Cloud](#). Select **My Services > Virtual Apps and Desktops** in the upper left menu.
2. From **Manage > Full Configuration**, select **Administrators** in the left pane.
3. Select the **Roles** tab.
4. Select a role in the upper middle pane. The **Role definition** tab in the lower pane lists the categories and permissions. Select a category to see the specific permissions. The **Administrators** tab lists the administrators who have been assigned the selected role.

Known issue: A Full Administrator entry does not display the correct set of permissions for a full access service administrator.

### How many administrators you need

The number of administrators and the granularity of their permissions generally depend on the size and complexity of the deployment.

- In small or proof of concept deployments, one or a few administrators do everything. There is no custom access delegation. In this case, each administrator has Full access, which always has the All scope.
- In larger deployments with more machines, applications, and desktops, more delegation is needed. Several administrators might have more specific functional responsibilities (roles). For example, two have Full access, and others are Help Desk Administrators. Also, an administrator might manage only certain groups of objects (scopes), such as machine catalogs in a particular department. In this case, create new scopes, plus administrators with the appropriate custom access role and scopes.

### Administrator management summary

Setting up administrators for the service follows this sequence:

1. If you want the administrator to have a role other than a Full administrator (which covers all subscribed services in Citrix Cloud) or a built-in role, create a custom role.
2. If you want the administrator to have a scope other than All (and a different scope is allowed for the intended role, and has not already been created), create scopes.
3. From Citrix Cloud, invite an administrator. If you want the new administrator to have anything other than the default Full access, specify a custom access role and scope pair.

Later, if you want to change an administrator's access (roles and scope), see [Configure custom access](#).

## Invite an administrator

To add administrators, follow the guidance in [Add administrators to a Citrix Cloud account](#). A subset of that information is repeated here.

### Important:

Do not confuse how “custom” and “custom access” are used.

- When creating administrators and assigning roles for the service in the Citrix Cloud console, the term “custom access” includes both the built-in roles and any additional custom roles that were created in the service’s **Manage > Full Configuration** interface.
- In the service’s **Manage > Full Configuration** interface, “custom” simply differentiates that role from a built-in role.

To add and invite an administrator:

1. After signing in to [Citrix Cloud](#), select **Identity and Access Management** in the upper left menu.
2. On the **Identity and Access Management** page, select **Administrators**. The display lists the current administrators in the account.
3. Select **Add administrators from**, and then select your authentication method. Enter the person’s email address. Optionally, select a role and scope pair.

If you do not select a custom access role and scope pair, the new administrator is assigned Full access by default. That setting includes access to all customer administrator functions in Citrix Cloud and in all subscribed services.

If you want that administrator to have more limited access, select a custom access role and scope pair. In that way, new administrators have the intended permissions when they sign in to Citrix Cloud for the first time.

4. Select **Invite**. Citrix Cloud sends an invitation to the email address you specified and adds the administrator to the list.

When the administrator receives the email, they select the **Join** link to accept the invitation.

## Create and manage roles

When administrators create or edit a role, they can enable only the permissions that they themselves have. This control prevents administrators from creating a role with more permissions than they currently have and then assigning it to themselves (or editing a role that they are already assigned).

Custom role names can contain up to 64 Unicode characters. Names cannot contain: backslash, forward slash, semicolon, colon, pound sign, comma, asterisk, question mark, equal sign, left arrow, right arrow, pipe, left or right bracket, left or right parenthesis, quotation marks, and apostrophe.

Role descriptions can contain up to 256 Unicode characters.

1. Sign in to [Citrix Cloud](#) if you haven't already. Select **My Services > Virtual Apps and Desktops** in the upper left menu.
2. From **Manage > Full Configuration**, select **Administrators** in the left pane.
3. Select the **Roles** tab.
4. Follow the instructions for the task you want to complete:
  - **View role details:** Select the role in the middle pane. The lower portion of the middle pane lists the object types and associated permissions for the role. Select the **Administrators** tab in the lower pane to display a list of administrators who currently have this role.
  - **Create a custom role:** Select **Create Role**. Enter a name and description. Select the object types and permissions. When you're done, select **Save**.

Create Role

Define a role for this administrator based on the permissions to manage different features.

Name:  
Example: My New Role

Description:  
Example: My New Role Description

Permissions: Select at least one permission for this Role

- Administrators
- AppDisks
- AppDNA
- Application Groups
- App-V
- Cloud
- Delivery Groups
- Director
- DirectorProbeAgent
- Hosts
- Logging
- Machine Catalogs
- Other permissions

Save Cancel

- **Copy a role:** Select the role in the middle pane and then select **Copy Role** in the action bar. Change the name, description, object types, and permissions, as needed. When you're done, select **Save**.
- **Edit a custom role:** Select the role in the middle pane and then select **Edit Role** in the action bar. Change the name, description, object types, and permissions, as needed. You cannot edit a built-in role. When you're done, select **Save**.
- **Delete a custom role:** Select the role in the middle pane and then select **Delete Role** in the action bar. When prompted, confirm the deletion. You cannot delete a built-in role. You cannot delete a custom role if it is assigned to an administrator.

## Create and manage scopes

By default, all roles have the All scope for their relevant objects. For example, a Delivery Group Administrator can manage all Delivery Groups. For some administrator roles, you can create a scope that allows that administrator role to access a subset of the relevant objects. For example, you might want to give a Machine Catalog Administrator access to only catalogs that contain a certain type of machines, rather than all catalogs.

- Full access administrators or custom access Cloud Administrators can create scopes for the Read Only Administrator, Machine Catalog Administrator, Delivery Group Administrator, and Host Administrator roles.
- Scopes cannot be created for Full access administrators, nor can they be created for Cloud Administrators or Help Desk Administrators. Those administrators always have the All scope.

Rules for creating and managing scopes:

- Scope names can contain up to 64 Unicode characters. Names cannot include: backslash, forward slash, semicolon, colon, pound sign, comma, asterisk, question mark, equal sign, left or right arrow, pipe, left or right bracket, left or right parenthesis, quotation marks, and apostrophe.
- Scope descriptions can contain up to 256 Unicode characters.
- When you copy or edit a scope, keep in mind that removing objects from the scope can make those objects inaccessible to an administrator. If the edited scope is paired with one or more roles, ensure that your scope updates do not make any role/scope pair unusable.

To create and manage scopes:

1. Sign in to [Citrix Cloud](#). Select **My Services > Virtual Apps and Desktops** in the upper left menu.
2. From **Manage > Full Configuration**, select **Administrators** in the left pane.
3. Select the **Scopes** tab.
4. Follow the instructions for the task you want to complete:
  - **View scope details:** Select the scope. The lower portion of the pane lists the objects and administrators that have that scope.
  - **Create a scope:** Select **Create Scope** in the action bar. Enter a name and description. The objects are listed by type, such as delivery group and machine catalog.
    - To include all objects of a particular type (for example, all delivery groups), select the check box for the object type.
    - To include individual objects within a type, expand the type and then select the check boxes for the objects (for example, specific delivery groups).

When you're done, select **Save**.



**Create Scope**

Define a scope based on the objects in your deployment.

**Name:**

Example: Sales

**Description: (Optional)**

Example: Sales team members

**Objects:**

- Delivery Groups
- Hosting
- Machine Catalogs

Select all objects of a particular type or specific objects within a type.

Save Cancel

- **Copy a scope:** Select the scope in the middle pane and then select **Copy Scope** in the action bar. Change the name, description. Change the object types and objects, as needed. When you're done, select **Save**.
- **Edit a scope:** Select the scope in the middle pane and then select **Edit Scope** in the action bar. Change the name, description, object types, and objects, as needed. When you're done, select **Save**.
- **Delete a scope:** Select the scope in the middle pane and then select **Delete Scope** in the action bar. When prompted, confirm the deletion.

You cannot delete a scope if it is assigned to a role. If you attempt to do this, an error message indicates that you do not have permission. In fact, the error occurs because the role/scope pair that uses this scope is assigned to an administrator. First, remove the role/scope pair assignment for all administrators who use it. Then delete the scope in the **Manage** console.

After you create a scope, it appears in the **Custom access** list in the Citrix Cloud console, paired with its appropriate role. You can then assign it to an administrator.

For example, let's say you create a scope named CAD, and select the catalogs that contain machines suitable for CAD applications. When you return to the Citrix Cloud console, the list of service-level

custom access role/scope pairs now has new entries (shown in bold):

- Cloud Administrator,All
- Delivery Group Administrator,All
- **Delivery Group Administrator,CAD**
- Help Desk Administrator,All
- Host Administrator,All
- **Host Administrator,CAD**
- Machine Catalog Administrator,All
- **Machine Catalog Administrator,CAD**
- Read Only,All
- **Read Only,CAD**

The Cloud Administrator and Help Desk Administrator always have the All scope, so the CAD scope does not apply to them.

### Configure custom access for an administrator

This feature lets you define access permissions of existing administrators or administrators you invite in a way that aligns with their role in your organization.

Changes you made to access permissions take 5 minutes to take effect. Logging out of the Full Configuration management interface and logging back on makes the changes take effect immediately. In scenarios where administrators still use the management interface after the changes take effect without reconnecting to it, a warning appears when they attempt to access items to which they no longer have permissions.

By default, when you invite administrators, they have Full access.

Remember: Full access allows the administrator to manage all subscribed services plus customer administrator Citrix Cloud operations (such as inviting more administrators). A Citrix Cloud deployment needs at least one administrator with Full access.

To configure custom access for an administrator:

1. Sign in to [Citrix Cloud](#). Select **Identity and Access Management > Administrators** in the upper left menu.
2. Locate the administrator you want to manage, select the ellipsis menu, and select **Edit access**.
3. Select **Custom access**. To configure service-specific custom access, under **Virtual Apps and Desktops**, select or clear the check marks next to one or more role and scope pairs in the **Custom access** list.

If you have not created any scopes and assigned them to a role, every role in the **Custom access** list has the All scope. For example, the role/scope entry **Delivery Group Administrator,All** indicates that role has the All scope.

When you create a role or scope, it appears in the custom access list for the service and can be selected. For example, if you created a scope named Catalog1, the **Custom access** list includes a **Machine Catalog Administrator,Catalog1** entry, in addition to the default **Machine Catalog Administrator,All** entry.

4. If the administrator you're editing already has custom access and you want to give that administrator full access, select **Full access**.
5. When you're done, select **Save**.

The following screenshot shows the full access and the custom access built-in administrator roles.

⚠ Custom Access requires at least one role to be selected.

---

Full access  
Full access allows administrators management control of Citrix Cloud and its services, as we

Custom access  
ⓘ Switching to custom access will remove management access to certain services.  
Custom access allows you to determine exactly which part of Citrix Cloud your administrato

[Select all](#)

**General Management**

---

Domains

Library

Notifications

Resource Location

**Virtual Apps and Desktops**

---

Cloud Administrator, All

Delivery Group Administrator, All

Help Desk Administrator, All - Access to 'Monitor' tab only

Host Administrator, All

Machine Catalog Administrator, All

Read Only Administrator, All

## Differences from on-premises Citrix Virtual Apps and Desktops

If you're familiar with delegated administration in the on-premises Citrix Virtual Apps and Desktops product, the service version has several differences.

In Citrix Cloud:

- Administrators are identified by their Citrix Cloud login, rather than their Active Directory account. You can create role/scope pairs for Active Directory individuals, but not groups.
- Administrators are created, configured, and deleted in the Citrix Cloud console, rather than the service.
- Role/scope pairs are assigned to administrators in the Citrix Cloud console, rather than the service.
- Reports are not available. You can view administrator, role, and scope information in the service's **Manage > Full Configuration** interface.
- The custom access Cloud Administrator is similar to a Full Administrator in the on-premises version. Both have full management and monitoring permissions for the Citrix Virtual Apps and Desktops version being used.

However, in the service, there is no named Full Administrator role. Do not equate "Full access" in Citrix Cloud with the "Full administrator" in on-premises Citrix Virtual Apps and Desktops. Full access in Citrix Cloud spans the platform-level domains, library, notifications, and resource locations, plus all subscribed services.

## Differences from earlier service releases

Before the release of the expanded custom access feature (September 2018), there were two custom access administrator roles: Full Administrator and Help Desk Administrator. When your deployment has delegated administration enabled (which is a platform setting), those roles are mapped automatically.

- An administrator who was formerly configured as a custom access **Virtual Apps and Desktops (or XenApp and XenDesktop) Service: Full Administrator** is now a custom access **Cloud Administrator**.
- An administrator who was formerly configured as a custom access **Virtual Apps and Desktops (or XenApp and XenDesktop) Service: Help Desk Administrator** is now a custom access **Help Desk Administrator**.

## More information

See [Delegated administration and monitoring](#) for information about administrators, roles, and scopes used in the service's **Monitor** console.

## Load balance machines

July 6, 2021

### Note:

- This feature applies to all your catalogs and only to multi-session OS machines.

This feature lets you control how to load balance machines. You have two options: vertical and horizontal. By default, horizontal load balancing is enabled.

- **Vertical load balancing.** Assigns an incoming user session to the most loaded machine that has not yet reached the maximum load. This saturates existing machines before moving on to new machines. Users disconnecting from existing machines frees up capacity on those machines. Incoming loads are then assigned to those machines. Vertical load balancing degrades the user experience but reduces costs (sessions maximize powered-on machine capacity).

Example: You have two machines configured for 10 sessions each. The first machine handles the first 10 concurrent sessions. The second machine handles the eleventh session.

Alternatively, you can use PowerShell to enable or disable vertical load balancing site-wide. Use the `UseVerticalScalingForRdsLaunches` setting in the `Set-BrokerSite` cmdlet. Use `Get-BrokerSite` to display the value of the `UseVerticalScalingForRdsLaunches` setting. See the cmdlet help for details.

- **Horizontal load balancing.** Assigns an incoming user session to the least-loaded, powered-on machine available. Horizontal load balancing improves the user experience but increases costs (because more machines are kept powered on). By default, horizontal load balancing is enabled.

Example: You have two machines configured for 10 sessions each. The first machine handles five concurrent sessions. The second machine also handles five.

To configure this feature, from **Manage > Full Configuration**, select **Settings** in the left pane. Select an option under **Load balance multi-session catalogs**.

## Local Host Cache

September 8, 2021

Local Host Cache enables connection brokering operations in a Citrix Virtual Apps and Desktops service deployment to continue when a Cloud Connector cannot communicate with Citrix Cloud. Local Host Cache engages when the network connection is lost for 60 seconds.

With Local Host Cache, users who are connected when an outage occurs can continue working uninterrupted. Reconnections and new connections experience minimal connection delays.

**Important:**

Local Host Cache requires a customer-deployed on-premises StoreFront as part of the deployment. You must add all Cloud Connectors that have (or can have) VDAs registered with them to the StoreFront as Delivery Controllers. A Cloud Connector that is not added to the StoreFront cannot transition to outage mode, which might result in user launch failures.

For deployments with no on-premises StoreFront, use the service continuity Citrix Workspace feature to allow users to connect to resources during outages. For more information, see [Service continuity](#).

## Data content

Local Host Cache includes the following information, which is a subset of the information in the main database:

- Identities of users and groups who are assigned rights to resources published from the site.
- Identities of users who are currently using, or who have recently used, published resources from the site.
- Identities of VDA machines (including Remote PC Access machines) configured in the site.
- Identities (names and IP addresses) of client Citrix Workspace app machines being actively used to connect to published resources.

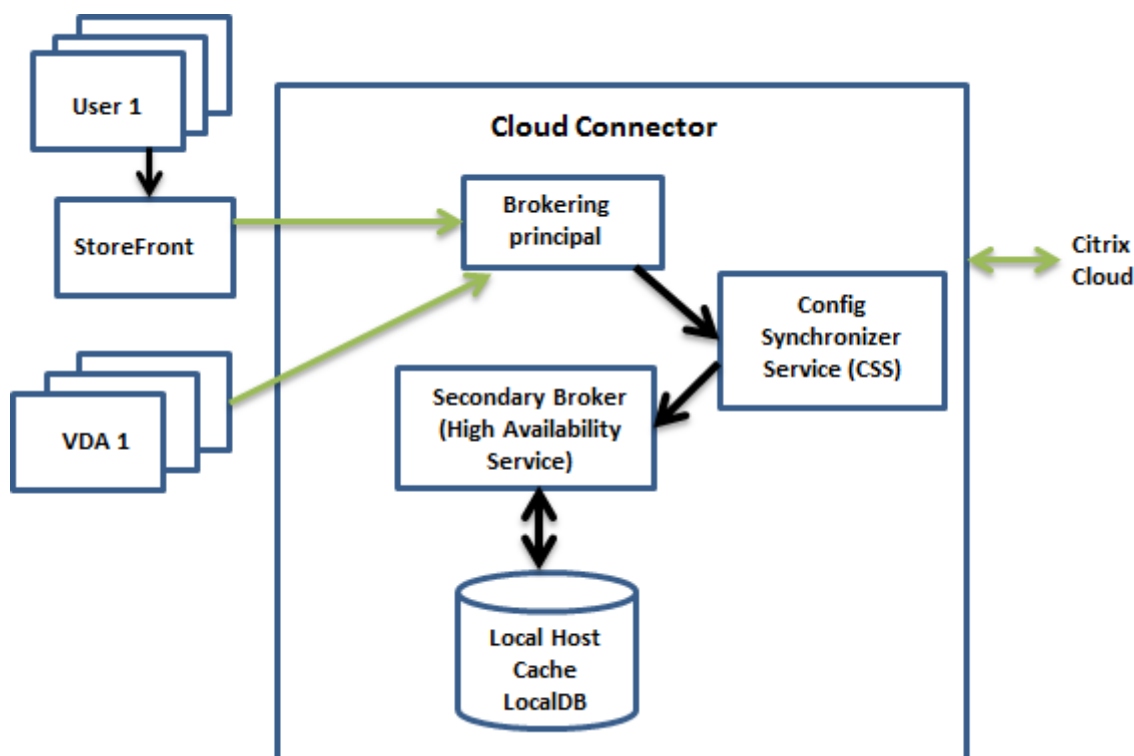
It also contains information for currently active connections that were established while the main database was unavailable:

- Results of any client machine endpoint analysis performed by Citrix Workspace app.
- Identities of infrastructure machines (such as Citrix Gateway and StoreFront servers) involved with the site.
- Dates, times, and types of recent activity by users.

## How it works

View how Local Host Cache interacts with Citrix Cloud.

[This is an embedded video. Click the link to watch the video](#)

**During normal operations**

- The Brokering Principal (also known as the Citrix Remote Broker Provider Service) on a Cloud Connector accepts connection requests from StoreFront. The Brokering Principal communicates with Citrix Cloud to connect users with VDAs that are registered with the Cloud Connector.
- The Citrix Config Synchronizer Service (CSS) checks with the broker in Citrix Cloud approximately every minute to see if any configuration changes were made. Those changes can be administrator-initiated (such as changing a delivery group property) or system actions (such as machine assignments).
- If a configuration change occurred since the previous check, the CSS synchronizes (copies) information to a secondary broker on the Cloud Connector. (The secondary broker is also known as the High Availability Service, or HA broker, as shown in the preceding figure.)

All configuration data is copied, not just items that changed since the previous check. The CSS imports the configuration data into a Microsoft SQL Server Express LocalDB database on the Cloud Connector. This database is referred to as the Local Host Cache database. The CSS ensures that the information in the Local Host Cache database matches the information in the site database in Citrix Cloud. The Local Host Cache database is re-created each time synchronization occurs.

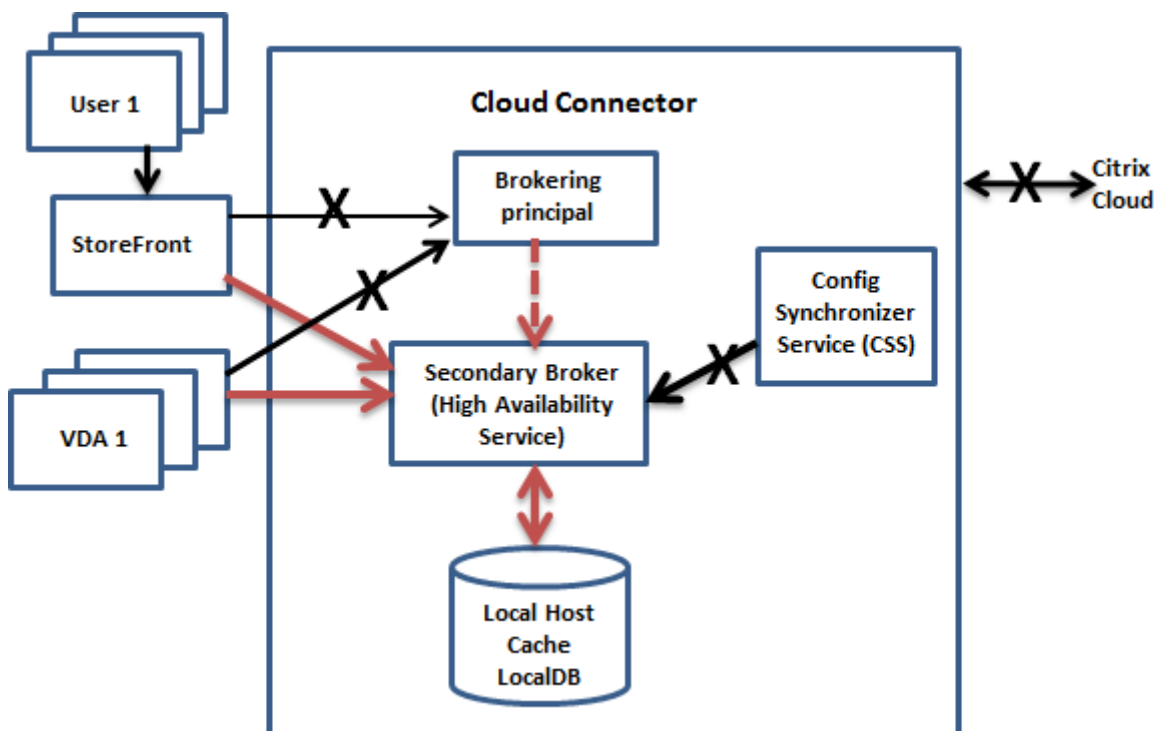
Microsoft SQL Server Express LocalDB (used by the Local Host Cache database) is installed automatically when you install a Cloud Connector. The Local Host Cache database cannot be shared across Cloud Connectors. You do not need to back up the Local Host Cache database. It is recre-



ated every time a configuration change is detected.

- If no changes occurred since the last check, the configuration data is not copied.

### During an outage



When an outage begins:

- The secondary broker starts listening for and processing connection requests.
- When the outage begins, the secondary broker does not have current VDA registration data, but when a VDA communicates with it, a registration process is triggered. During that process, the secondary broker also gets current session information about that VDA.
- While the secondary broker is handling connections, the Brokering Principal continues to monitor the connection to Citrix Cloud. When the connection is restored, the Brokering Principal instructs the secondary broker to stop listening for connection information, and the Brokering Principal resumes brokering operations. The next time a VDA communicates with the Brokering Principal, a registration process is triggered. The secondary broker removes any remaining VDA registrations from the previous outage. The CSS resumes synchronizing information when it learns that configuration changes have occurred in Citrix Cloud.

In the unlikely event that an outage begins during a synchronization, the current import is discarded and the last known configuration is used.

The event log indicates when synchronizations and outages occur.

There is no time limit imposed for operating in outage mode.

You can also intentionally trigger an outage. See Force an outage for details about why and how to do this.

### Resource locations with multiple Cloud Connectors

Among its other tasks, the CSS routinely provides the secondary broker with information about all Cloud Connectors in the resource location. Having that information, each secondary broker knows about all peer secondary brokers running on other Cloud Connectors in the resource location.

The secondary brokers communicate with each other on a separate channel. Those brokers use an alphabetical list of FQDN names of the machines they're running on to determine (elect) which secondary broker will broker operations in the zone if an outage occurs. During the outage, all VDAs re-register with the elected secondary broker. The non-elected secondary brokers in the zone actively reject incoming connection and VDA registration requests.

If an elected secondary broker fails during an outage, another secondary broker is elected to take over, and VDAs register with the newly elected secondary broker.

During an outage, if a Cloud Connector is restarted:

- If that Cloud Connector is not the elected broker, the restart has no impact.
- If that Cloud Connector is the elected broker, a different Cloud Connector is elected, causing VDAs to register. After the restarted Cloud Connector powers on, it automatically takes over brokering, which causes VDAs to register again. In this scenario, performance can be affected during the registrations.

The event log provides information about elections.

### What is unavailable during an outage, and other differences

There is no time limit imposed for operating in outage mode. However, if the outage is due to loss of Citrix Cloud connectivity from their resource location, Citrix recommends restoring connectivity from the resource location as quickly as possible.

During an outage:

- You cannot use the **Manage** interfaces.
- You have limited access to the Remote PowerShell SDK.
  - You must first:
    - \* Add a registry key `EnableCssTestMode` with a value of 1: `New-ItemProperty -Path HKLM:\SOFTWARE\Citrix\DesktopServer\LHC -Name EnableCssTestMode -PropertyType DWORD -Value 1`
    - \* Set the SDK auth to `OnPrem` so that the SDK proxy does not try to redirect the cmdlet calls: `$XDSDKAuth="OnPrem"`

\* Use port 89: `Get-BrokerMachine -AdminAddress localhost:89 | Select MachineName, ContollerDNSName, DesktopGroupName, RegistrationState`

- After running those commands, you can access:

\* All `Get-Broker*` cmdlets.

\* The power management cmdlets `New-BrokerHostingPowerAction`, `Set-BrokerHostingPowerAction`, and `Remove-BrokerHostingPowerAction`.

- Monitoring data is not sent to Citrix Cloud during an outage. So, the **Monitor** functions do not show activity from an outage interval.
- Hypervisor credentials cannot be obtained from the Host Service. All machines are in the unknown power state, and no power operations can be issued. However, VMs on the host that are powered-on can be used for connection requests.

(Preview. For instructions, see [this blog post](#).) Exceptions for supported Citrix Hypervisor and VMware hypervisors only:

- Hypervisor credentials are synchronized before the outage, and stored securely on the Cloud Connector.
  - VMs from local hypervisors have the latest power state in the Local Host Cache database.
  - Power actions supported by the hypervisor are available during outage mode. VMs that are powered on can be used for connection requests.
  - VMs on the hypervisor can be powered-on on-demand automatically during session launch.
- An assigned machine can be used only if the assignment occurred during normal operations. New assignments cannot be made during an outage.
  - Automatic enrollment and configuration of Remote PC Access machines is not possible. However, machines that were enrolled and configured during normal operation are usable.
  - Server-hosted applications and desktop users might use more sessions than their configured session limits, if the resources are in different zones.
  - Users can launch applications and desktops only from registered VDAs in the zone containing the currently active/elected broker. Launches across zones (from a broker in one zone to a VDA in a different zone) are not supported during an outage.
  - If a site database outage occurs before a scheduled restart begins for VDAs in a delivery group, the restarts begin when the outage ends. This scenario can have unintended results. For more information, see [Scheduled restarts delayed due to database outage](#).
  - [Zone preference](#) cannot be configured. If configured, preferences are not considered for session launch.

- [Tag restrictions](#) are not considered for session launch. When tag restrictions are configured, and a StoreFront store's [advanced health check](#) option is enabled, sessions might intermittently fail to launch.

## StoreFront requirement

Local Host Cache requires a customer-deployed on-premises StoreFront as part of the deployment. You must add all Cloud Connectors that have (or can have) VDAs registered with them to the StoreFront as Delivery Controllers. A Cloud Connector that is not added to the StoreFront cannot transition to outage mode, which might result in user launch failures.

## Resource availability

You can ensure the availability of resources (apps and desktops) during an outage in two ways:

- Publish the resources in every resource location in your deployment.
- Publish the resources to at least one resource location. Then use the following procedure to enable the advanced health check feature in each StoreFront store.
  1. Upgrade the StoreFront installation in each resource location to minimum version 1912 CU1. For guidance, see the StoreFront documentation.
  2. For each StoreFront store, enable the advanced health check option. In the store's [web.config](#) file, under [farmsets](#), add `advancedHealthCheck="on"`.

Option example:

```
{farmsets}>
  <farmset name="Default" enableFileTypeAssociation="on" pooledSockets="off"
    serverCommunicationAttempts="1" communicationTimeout="30" connectionTimeout="6"
    multiFarmAuthenticationMode="ANY" backgroundHealthCheckPollingPeriod="00:01:00"
    advancedHealthCheck="on">
  <farm name="Controller12345" xmlPort="80" transport="HTTP" sslRelayPort="443"
    bypassDuration="60" allFailedBypassDuration="0" loadBalance="on"
    ticketTimeToLive="200" farmType="XenDesktop" maxServersPerRequest="0"
    zones="">
```

3. After you update the file, manually restart IIS. Repeat the web.config file update and IIS restart for other stores.

## Application and desktop support

Local Host Cache works only with customer-deployed StoreFront. It does not support Workspace.

Local Host Cache supports server-hosted applications and desktops, and static (assigned) desktops.

Local Host Cache supports desktop (single-session) VDAs in pooled delivery groups, as follows.

- By default, power-managed desktop VDAs in pooled delivery groups (created by MCS or Citrix Provisioning) that have the `ShutdownDesktopsAfterUse` property enabled are placed into maintenance mode when an outage occurs. You can change this default, to allow those desktops to be used during an outage.

However, you cannot necessarily rely on the power management during the outage. (Power management resumes after normal operations resume.) Also, those desktops might contain data from the previous user, because they have not been restarted.

- To override the default behavior, it must be enabled site-wide and for each affected delivery group, using PowerShell commands.

For site-wide, run the following command:

```
Set-BrokerSite -ReuseMachinesWithoutShutdownInOutageAllowed $true
```

By default, all delivery groups are not enabled for this feature. There are two options to enable it at the delivery group level:

- **Enable for selected delivery groups:** For each affected delivery group, run the following command.

```
Set-BrokerDesktopGroup -Name "name"-ReuseMachinesWithoutShutdownInOutage $true
```

- **Enable for all delivery groups:** To enable the delivery group level setting by default, run the following command. This setting applies to all newly created delivery groups (that is, all delivery groups you create after enabling the setting).

```
Set-BrokerSite -DefaultReuseMachinesWithoutShutdownInOutage $true
```

To enable this for existing delivery groups, run the command noted previously (`Set-BrokerDesktopGroup -Name "name"-ReuseMachinesWithoutShutdownInOutage $true`).

Enabling this feature in the site and the delivery groups does not affect how the configured `ShutdownDesktopsAfterUse` property works during normal operations.

## Differences from XenApp 6.x releases

Although this Local Host Cache implementation shares the name of the Local Host Cache feature in XenApp 6.x and earlier XenApp releases, this Local Host Cache is an entirely different implementation technically. This implementation has significant improvements, is more robust, is more immune to corruption, and requires less maintenance.

## Manage Local Host Cache

View how to configure Local Host Cache.

[This is an embedded video. Click the link to watch the video](#)

In a Citrix Virtual Apps and Desktops service deployment, Local Host Cache is always enabled. You don't have to do anything else to configure or manage it.

As noted previously, the Microsoft SQL Server Express LocalDB database is installed automatically when you install a Cloud Connector in a resource location. Do not attempt to disable or remove it. Citrix updates the Cloud Connector regularly. If you disable or remove the SQL Server Express LocalDB software manually, the next Cloud Connector update replaces it.

## Verify that Local Host Cache is working

View how to verify that Local Host Cache is configured correctly.

[This is an embedded video. Click the link to watch the video](#)

To verify that Local Host Cache is set up and working correctly:

- Verify that the resource location contains a local StoreFront that points to all the Cloud Connectors in that resource location.
- Ensure that synchronization imports complete successfully. Check the event logs.
- Ensure that the Local Host Cache database was created on each Cloud Connector. This confirms that the High Availability Service can take over, if needed.
  - On the Cloud Connector server, browse to `c:\Windows\ServiceProfiles\NetworkService`.
  - Verify that `HaDatabaseName.mdf` and `HaDatabaseName_log.ldf` are created.
- Force an outage on all Cloud Connectors in the resource location. After you've verified that Local Host Cache works, remember to place all the Cloud Connectors back into normal mode. This can take approximately 15 minutes.

## Event logs

Event logs indicate when synchronizations and outages occur. In event viewer logs, outage mode is referred to as *HA mode*.

## Config Synchronizer Service

During normal operations, the following events can occur when the CSS imports the configuration data into the Local Host Cache database using the Local Host Cache broker.

- 503: The Citrix Config Sync Service received an updated configuration. This event occurs each time an updated configuration is received from Citrix Cloud. It indicates the start of the synchronization process.
- 504: The Citrix Config Sync Service imported an updated configuration. The configuration import completed successfully.
- 505: The Citrix Config Sync Service failed an import. The configuration import did not complete successfully. If a previous successful configuration is available, it is used if an outage occurs. However, it will be out-of-date from the current configuration. If there is no previous configuration available, the service cannot participate in session brokering during an outage. In this case, see the Troubleshoot section, and contact Citrix Support.
- 507: The Citrix Config Sync Service abandoned an import because the system is in outage mode and the Local Host Cache broker is being used for brokering. The service received a new configuration, but the import was abandoned because an outage occurred. This is expected behavior.
- 510: No Configuration Service configuration data received from primary Configuration Service.
- 517: There was a problem communicating with the primary Broker.
- 518: Config Sync script aborted because the secondary Broker (High Availability Service) is not running.

### High Availability Service

This service is also known as the Local Host Cache broker.

- 3502: An outage occurred and the Local Host Cache broker is performing broker operations.
- 3503: An outage was resolved and normal operations have resumed.
- 3504: Indicates which Local Host Cache broker is elected, plus other Local Host Cache brokers involved in the election.

### Force an outage

You might want to deliberately force an outage.

- If your network is going up and down repeatedly. Forcing an outage until the network issues are resolved prevents continuous transition between normal and outage modes (and the resulting frequent VDA registration storms).
- To test a disaster recovery plan.
- To help ensure that Local Host Cache is working correctly.

Although a Cloud Connector can be updated during a forced outage, unforeseen issues can occur. We recommend you [set a schedule for Cloud Connector updates](#) that avoids forced outage mode intervals.

To force an outage, edit the registry of each Cloud Connector server. In `HKLM\Software\Citrix\DesktopServer\LHC`, create and set `OutageModeForced` as `REG_DWORD` to 1. This setting instructs

the Local Host Cache broker to enter outage mode, regardless of the state of the connection to Citrix Cloud. Setting the value to 0 takes the Local Host Cache broker out of outage mode.

To verify events, monitor the `Current_HighAvailabilityService` log file in `C:\ProgramData\Citrix\WorkspaceCloud\Logs\Plugins\HighAvailabilityService`.

## Troubleshoot

Several troubleshooting tools are available when a synchronization import to the Local Host Cache database fails and a 505 event is posted.

**CDF tracing:** Contains options for the ConfigSyncServer and BrokerLHC modules. Those options, along with other broker modules, can identify the problem.

**Report:** If a synchronization import fails, you can generate a report. This report stops at the object causing the error. This report feature affects synchronization speed, so Citrix recommends disabling it when not in use.

To enable and produce a CSS trace report, enter the following command:

```
New-ItemProperty -Path HKLM:\SOFTWARE\Citrix\DesktopServer\LHC -Name EnableCssTraceMode -PropertyType DWORD -Value 1
```

The HTML report is posted at: `C:\Windows\ServiceProfiles\NetworkService\AppData\Local\Temp\CitrixBrokerConfigSyncReport.html`

After the report is generated, enter the following command to disable the reporting feature:

```
Set-ItemProperty -Path HKLM:\SOFTWARE\Citrix\DesktopServer\LHC -Name EnableCssTraceMode -Value 0
```

## More information

See [Scale and size considerations for Local Host Cache](#) for information about:

- Testing methodologies and results
- RAM size considerations
- CPU core and socket configuration considerations
- Storage considerations

## Manage security keys

August 17, 2021



**Note:**

- You must use this feature in combination with StoreFront 1912 LTSR CU2 or later.
- The Secure XML feature is only supported on Citrix ADC and Citrix Gateway release 12.1 and later.

This feature lets you allow only approved StoreFront and Citrix Gateway machines to communicate with Citrix Cloud. After you enable this feature, any requests that do not contain the key are blocked. Use this feature to add an extra layer of security to protect against attacks originating from the internal network.

A general workflow to use this feature is as follows:

1. Enable the feature in Studio by using the PowerShell SDK.
2. Configure settings in Studio. (Use the Studio console or PowerShell).
3. Configure settings in StoreFront. (Use PowerShell).
4. Configure settings in Citrix ADC.

## Enable the security key feature

By default, the feature is disabled. To enable it, use the Remote PowerShell SDK. For more information about the Remote PowerShell SDK, see [SDKs and APIs](#).

To enable the feature, perform these steps:

1. Run the Citrix Virtual Apps and Desktops Remote PowerShell SDK.
2. In a command window, run the following commands:
  - `Add-PSSnapIn Citrix*`. This command adds the Citrix snap-ins.
  - `Set-ConfigSiteMetadata -Name "Citrix_DesktopStudio_SecurityKeyManagementEnabled" -Value "True"`

## Configure settings in Studio

You can configure settings in Studio by using the Studio console or PowerShell.


### Use the Studio console


After enabling the feature, navigate to **Studio > Settings > Manage security key** and click **Edit**. The **Manage Security Key** window appears. Click **Save** to apply your changes and to exit the window.


### Manage Security Key


This feature lets you manage the security key used to authenticate Citrix Gateway and StoreFront when they communicate with the Delivery Controller.


[Learn more](#)


Key1: 



Key2: 



Require key for communications over XML port (StoreFront only) 

Require key for communications over STA port 

**Important:**

- There are two keys available for use. You can use the same key or different keys for communications over the XML and STA ports. We recommend that you use only one key at a time. The unused key is used only for key rotation.
- Do not click the refresh icon to update the key already in use. If you do, service interruption will occur.

Click the refresh icon to generate new keys.

**Require key for communications over XML port (StoreFront only).** If selected, require a key to authenticate communications over the XML port. StoreFront communicates with Citrix Cloud over this port. For information about changing the XML port, see Knowledge Center article [CTX127945](#).

**Require key for communications over STA port.** If selected, require a key to authenticate communications over the STA port. Citrix Gateway and StoreFront communicate with Citrix Cloud over this port. For information about changing the STA port, see Knowledge Center article [CTX101988](#).

After applying your changes, click **Close** to exit the **Manage Security Key** window.

**Use PowerShell**

The following are PowerShell steps equivalent to Studio operations.

1. Run the Citrix Virtual Apps and Desktops Remote PowerShell SDK.

2. In a command window, run the following command:
  - `Add-PSSnapIn Citrix*`
3. Run the following commands to generate a key and set up Key1:
  - `New-BrokerXmlServiceKey`
  - `Set-BrokerSite -XmlServiceKey1 <the key you generated>`
4. Run the following commands to generate a key and set up Key2:
  - `New-BrokerXmlServiceKey`
  - `Set-BrokerSite -XmlServiceKey2 <the key you generated>`
5. Run one or both of the following commands to enable the use of a key in authenticating communications:
  - To authenticate communications over the XML port:
    - `Set-BrokerSite -RequireXmlServiceKeyForNFuse $true`
  - To authenticate communications over the STA port:
    - `Set-BrokerSite -RequireXmlServiceKeyForSta $true`

See the PowerShell command help for guidance and syntax.

## Configure settings in StoreFront

After completing the configuration in Studio, you need to configure relevant settings in StoreFront by using PowerShell.

On the StoreFront server, run the following PowerShell commands:

- To configure the key for communications over the XML port, use the `Get-STFStoreService` and `Set-STFStoreService` commands. For example:
  - `PS C:\> Set-STFStoreFarm $farm -Farmtype XenDesktop -Port 80 -TransportType HTTP -Servers <domain name1, domain name2> -XMLValidationEnabled $true -XMLValidationSecret <the key you generated in Studio>`
- To configure the key for communications over the STA port, use the `New-STFSecureTicketAuthority` command. For example:
  - `PS C:\> $sta = New-STFSecureTicketAuthority -StaUrl <STA URL> -StaValidationEnabled $true -StavalidationSecret <the key you generated in Studio>`

See the PowerShell command help for guidance and syntax.

## Configure settings in Citrix ADC

**Note:**

Configuring this feature in Citrix ADC is not required unless you use Citrix ADC as your gateway. If you use Citrix ADC, follow the steps below.

1. Ensure that the following prerequisite configuration is already in place:

- The following Citrix ADC related IP addresses are configured.
  - Citrix ADC Management IP (NSIP) address for accessing the Citrix ADC console. For details, see [Configuring the NSIP address](#).

Dashboard	Configuration	Reporting	Documentation	Downloads
-----------	---------------	-----------	---------------	-----------



### Citrix ADC IP Address

If you change the Citrix ADC IP address and subnet mask, click **Reboot** for the changes to become effective. Citrix recommends that you change the default administrator (nsroot) password.

Citrix ADC IP Address\*

Netmask\*

Change Administrator Password

**Done**   **Back**

- Subnet IP (SNIP) address for enabling communication between the Citrix ADC appliance and the back-end servers. For details, see [Configuring Subnet IP Addresses](#).
- Citrix Gateway virtual IP address and load balancer virtual IP address to log in to the ADC appliance for session launch. For details, see [Create a virtual server](#).



### Subnet IP Address

A subnet IP address is used by the Citrix ADC to communicate with the backend servers. Citrix ADC uses this subnet IP address as a source IP address to proxy the client connections as well as to send monitor probes to check the health of the backend servers.

The infographic shows the usage of SNIP in client server communication.

Depending on your network topology, you might have to configure additional subnet IP addresses.

For more information about subnet IP addresses, [click here](#).

The screenshot shows a configuration form with two input fields. The first field is labeled 'Subnet IP Address\*' and is empty, with a red error message 'Please enter value' to its right. The second field is labeled 'Netmask\*' and contains the value '255 . 255 . 255 . 0'. At the bottom of the form are two buttons: 'Done' and 'Back'.

- The required modes and features in the Citrix ADC appliance are enabled.
  - To enable the modes, in the Citrix ADC GUI navigate to **System > Settings > Configure Mode**.
  - To enable the features, in the Citrix ADC GUI navigate to **System > Settings > Configure Basic Features**.
- Certificates related configurations are complete.
  - The Certificate Signing Request (CSR) is created. For details, see [Create a certificate](#).

Dashboard Configuration Reporting Documentation Dow

## ← Create RSA Key

Key Filename\*

Choose File ▾ SSLTest ⓘ

Key Size(bits)\*

2048 ▾

Public Exponent Value\*

F4 ▾

Key Format\*

PEM ▾

PEM Encoding Algorithm

▾

PEM Passphrase

▾

Confirm PEM Passphrase

▾

PKCS8

Create Close

- The server and CA certificates and root certificates are installed. For details, see [Install, link, and updates](#).

## ← Install Server Certificate

Certificate-Key Pair Name\*  
 ⓘ

Certificate File Name\*  
 CSR\_DER ⓘ

Key File Name  
 ns-server.key ⓘ

Notify When Expires

---

2 SNMP Trap destination found.

Notification Period

## ← Install CA Certificate

Certificate-Key Pair Name\*  
 ⓘ

Certificate File Name\*  
 ns-server.cert ⓘ

Notify When Expires

---

2 SNMP Trap destination found.

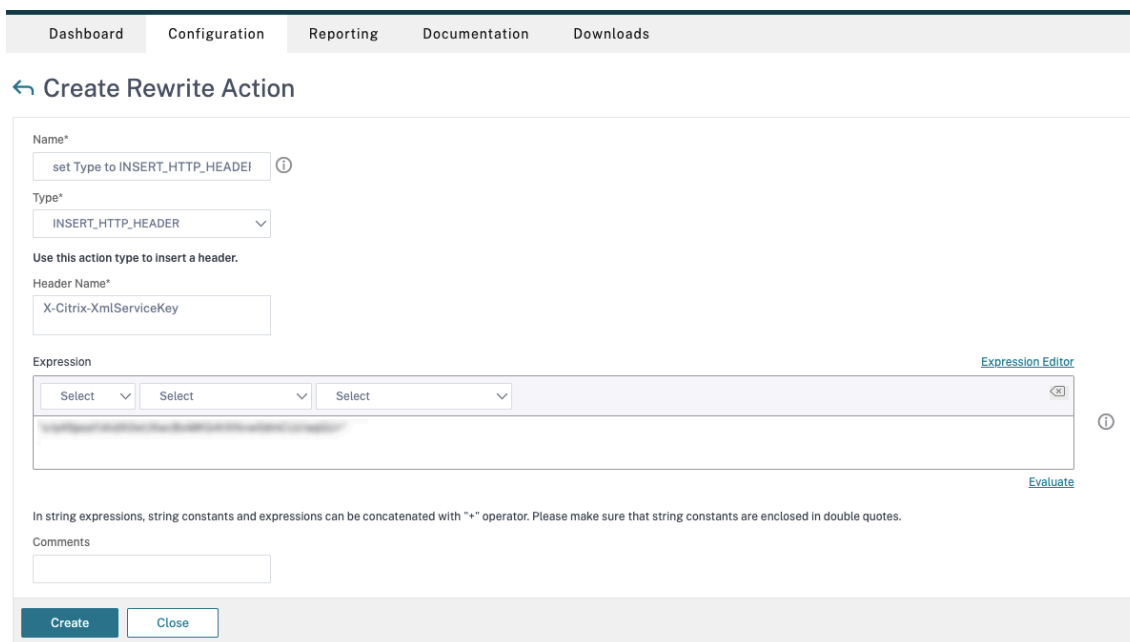
Notification Period

- A Citrix Gateway has been created for Citrix Virtual Desktops. Test the connectivity by clicking the **Test STA Connectivity** button to confirm that the virtual servers are online. For details, see [Setting up Citrix ADC for Citrix Virtual Apps and Desktops](#).



2. Add a rewrite action. For details, see [Configuring a Rewrite Action](#).

- a) Navigate to **AppExpert > Rewrite > Actions**.
- b) Click **Add** to add a new rewrite action. You can name the action as “set Type to INSERT\_HTTP\_HEADER”.



- a) In **Type**, select **INSERT\_HTTP\_HEADER**.
- b) In **Header Name**, enter X-Citrix-XmlServiceKey.
- c) In **Expression**, add `<XmlServiceKey1 value>` with the quotes. You can copy the XmlServiceKey1 value from your Desktop Delivery Controller configuration.



```

PS C:\Users\tyadmin> Get-BrokerSite

BaseOU :
BrokerServiceGroupUid :
ColorDepth :
ConfigLastChangeTime :
ConfigurationServiceGroupUid :
ConnectionLeasingEnabled :
DefaultMinimumFunctionalLevel :
DesktopGroupIconUid :
DnsResolutionEnabled :
IsSecondaryBroker :
LicenseEdition :
LicenseGraceSessionsRemaining :
LicenseModel :
LicenseServerName :
LicenseServerPort :
LicensedSessionsActive :
LicensingBurnIn :
LicensingBurnInDate :
LicensingGraceHoursLeft :
LicensingGracePeriodActive :
LicensingOutOfBoxGracePeriodActive :
LocalHostCacheEnabled :
MetadataMap :
Name :
PeakConcurrentLicenseUsers :
RequireXmlServiceKeyForNFuse :
RequireXmlServiceKeyForSta :
ReuseMachinesWithoutShutdownInOutageAllowed :
SecureIcaRequired :
TotalUniqueLicenseUsers :
TrustManagedAnonymousXmlServiceRequests :
TrustRequestsSentToTheXmlServicePort :
UseVerticalScalingForRdsLaunches :
XmlServiceKey1 :
XmlServiceKey2 :
    
```

3. Add a rewrite policy. For details, see [Configuring a Rewrite Policy](#).

- a) Navigate to **AppExpert > Rewrite > Policies**.
- b) Click **Add** to add a new policy.

- a) In **Action**, select the action created in the earlier step.
  - b) In **Expression**, add HTTP.REQ.IS\_VALID.
  - c) Click **OK**.
4. Set up load balancing. You must configure one load balancing virtual server per STA server. If not the sessions fail to launch.

For details, see [Set up basic load balancing](#).

- a) Create a load balancing virtual server.
  - Navigate to **Traffic Management > Load Balancing > Servers**.
  - In **Virtual Servers** page, click **Add**.

Dashboard Configuration Reporting Documentation Downloads

← Load Balancing Virtual Server

**Basic Settings**

Create a virtual server by specifying a name, an IP address, a port, and a protocol type. If an application is accessible from the Internet, the virtual server IP (VIP) address is a public IP address. If the application is accessible only from the local area network (LAN) or wide area network (WAN), the VIP is usually a private (RFC1918) IP address. You can configure multiple virtual servers to receive client requests, thereby increasing the availability of resources to process client requests.

Name\*  
LBserver1 ⓘ

Protocol\*  
HTTP ▾

IP Address Type\*  
IP Address ⓘ

IP Address\*  
. . . . . ⓘ

Port\*  
80

▶ More

OK Cancel

- In **Protocol**, select **HTTP**.
  - Add the load balancing virtual IP address and in **Port** select **80**.
  - Click **OK**.
- b) Create a load balancing service.
    - Navigate to **Traffic Management > Load Balancing > Services**.

Dashboard Configuration Reporting Documentation Downloads

## ← Load Balancing Service

### Basic Settings

Service Name\*  
DDCService1 ⓘ

New Server  Existing Server

Server\*  
[Blurred] ▾


Protocol\*  
HTTP ▾

Port\*  
80

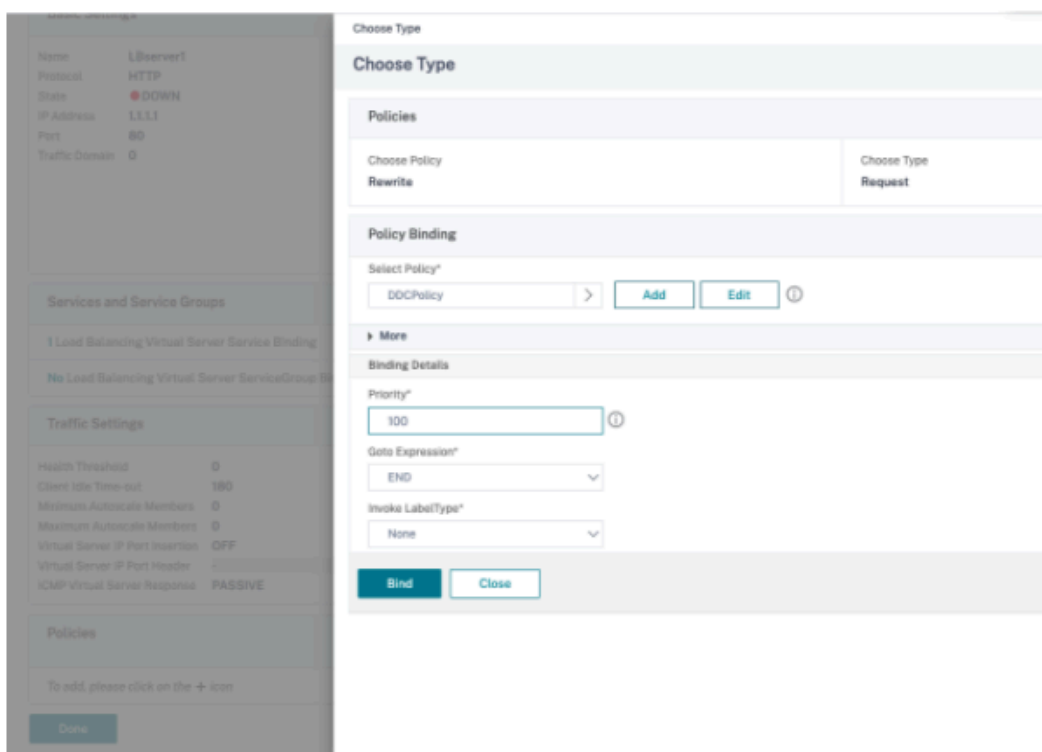
▶ More

OK Cancel

- In **Existing Server**, select the virtual server created in the previous step.
  - In **Protocol**, select **HTTP** and in **Port** select **80**.
  - Click **OK**, and then click **Done**.
- c) Bind the service to the virtual server.
- Select the virtual server created earlier and click **Edit**.
  - In **Services and Service Groups**, click **No Load Balancing Virtual Server Service Binding**.



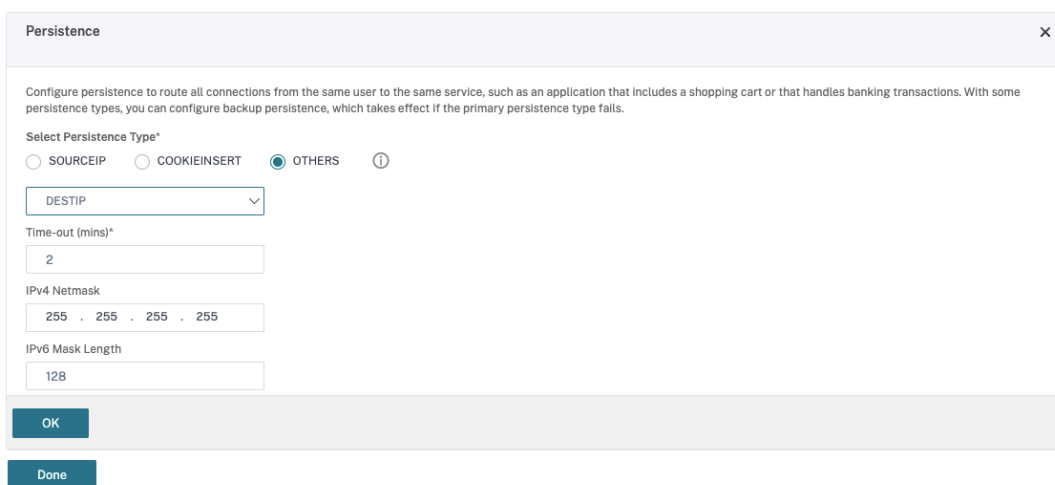
- In **Service Binding**, select the service created earlier.
  - Click **Bind**.
- d) Bind the rewrite policy created earlier to the virtual server.
- Select the virtual server created earlier and click **Edit**.
  - In **Advanced Settings**, click **Policies**, and then in **Policies** section click **+**.



- In **Choose Policy**, select **Rewrite** and in **Choose Type**, select **Request**.
- Click **Continue**.
- In **Select Policy**, select the rewrite policy created earlier.
- Click **Bind**.
- Click **Done**.

e) Set up persistence for the virtual server, if necessary.

- Select the virtual server created earlier and click **Edit**.
- In **Advanced Settings**, click **Persistence**.



- Select persistence type as **Others**.

- Select **DESTIP** to create persistence sessions based on the IP address of the service selected by the virtual server (the destination IP address)
- In **IPv4 Netmask**, add network mask same as that of the DDC.
- Click **OK**.


f) Repeat these steps for the other virtual server as well.

### Configuration changes if the Citrix ADC appliance is already configured with Citrix Virtual Desktops


If you have already configured the Citrix ADC appliance with Citrix Virtual Desktops, then to use the Secure XML feature, you must make the following configuration changes.

- Before the session launch, change the **Security Ticket Authority URL** of the gateway to use the FQDNs of the load balancing virtual servers.
- Ensure that the `TrustRequestsSentToTheXmlServicePort` parameter is set to False. By default, `TrustRequestsSentToTheXmlServicePort` parameter is set to False. However, if the customer has already configured the Citrix ADC for Citrix Virtual Desktops, then the `TrustRequestsSentToTheXmlServicePort` is set to True.

1. In the Citrix ADC GUI, navigate to **Configuration > Integrate with Citrix Products** and click **XenApp and XenDesktop**.
2. Select the gateway instance and click the edit icon.



3. In the StoreFront pane, click the edit icon.

StoreFront 	
StoreFront URL	https://yj-en2016-1.ddc.com
Storefront Status	
Receiver for Web Path	/Citrix/StoreWeb
Default Active Directory Domain	ddc.com
List of Secure Ticket Authority URL(s) with status	
http://[redacted].com	● DOWN
http://[redacted].com	● DOWN
http://[redacted].com	● DOWN
http://[redacted].com	● DOWN

4. Add the **Secure Ticket Authority URL**.

- If the Secure XML feature is enabled, then the STA URL must be the URL of the load balancing service.
- If the Secure XML feature is disabled, then the STA URL must be the URL of STA (DDC's address) and the TrustRequestsSentToTheXmlServicePort parameter on the DDC must be set to True.

**StoreFront**

StoreFront URL\*  
 ⓘ

**Retrieve Stores**

Receiver for Web Path\*

Default Active Directory Domain\*

Secure Ticket Authority URL\*

<input type="text" value="http://[redacted].com"/>	×
<input type="text" value="http://[redacted].com"/>	×
<input type="text" value="http://[redacted].com"/>	×
<input type="text" value="http://[redacted].com"/>	×

**Test STA Connectivity**

Use this StoreFront for Authentication

## Scale and size considerations for Local Host Cache

November 12, 2020

This article contains detailed information about Local Host Cache testing, and considerations when configuring your deployment. For general information about Local Host Cache and how it works, see



## Local Host Cache.

### Overview

The Local Host Cache feature in the Citrix Virtual Apps and Desktops service allows connection brokering in a site to continue if there is an outage. An outage happens if the WAN link between the site and the management console fails in a Citrix Cloud environment. In December 2017, we tested the Citrix Cloud Connector machine configuration using the Citrix Virtual Apps and Desktops service Local Host Cache feature. The test results provided in this document detail the tested maximums in December 2017. Best practice recommendations are based on those tested maximums.

This article assumes that the reader can set up and configure a Citrix Cloud environment according to recommended standards, with a minimum of three Cloud Connectors.

Local Host Cache supports only on-premises StoreFront in each resource location or zone.

While outage mode is active, if the elected Cloud Connector that brokers the sessions has an outage, the second Cloud Connector becomes the elected High Availability Service. After the election, the second Cloud Connector takes over to broker the sessions. The Local Host Cache feature uses only one socket for multi-core CPUs for the Cloud Connector VM configuration. In this scenario, we recommend a 4-core, 1-socket configuration.

### Summary

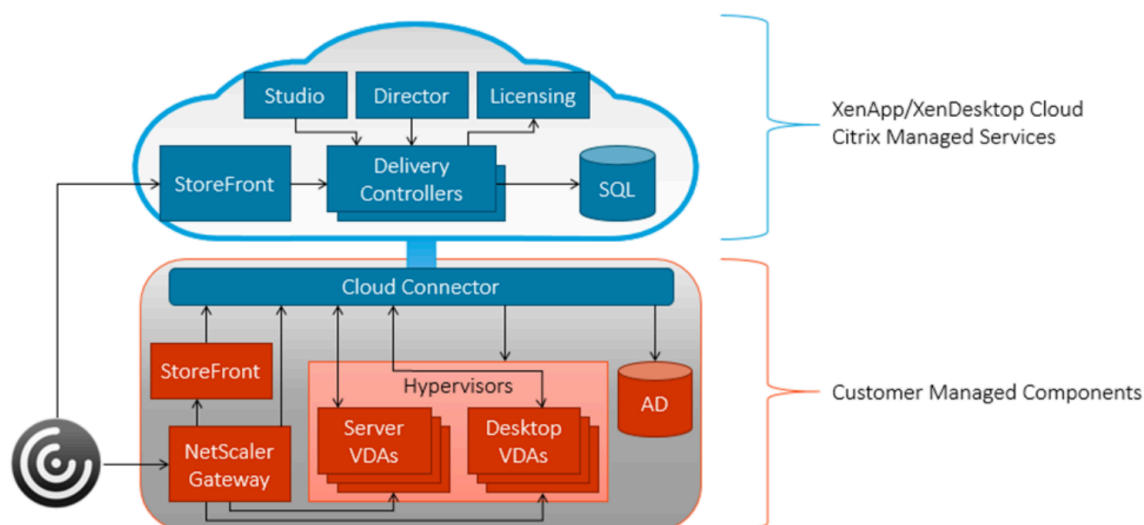
All results in this summary are based on the findings from test environments which we configured as detailed in the following sections. Different system configurations yield different results.

Key recommendations based on test results

- We recommend, for high availability sites that host no more than 5,000 workstations or 500 server VDAs, that you configure 3 VMs dedicated to the Cloud Connector. Each Cloud Connector VM requires 4 vCPU with 4 GB RAM. This configuration is an N+1 high availability configuration. Cloud Connectors are deployed in high availability sets. Cloud Connectors are not load-balanced. Because each CPU can process a limited number of connections, the CPU is the greatest limiting factor related to the number of workstations or server VDAs supported.
- Although this document focuses on testing with two Cloud Connectors, an N+1 set of three Cloud Connectors is recommended.
- We conducted session launch tests to compare Local Host Cache outage mode active and inactive after a new configuration was synchronized and imported. The launch tests covered scenarios with 5,000, 20,000, and 1,000 session launches against the respective number of available workstations.
  - 5,000 sessions launched against 5,000 workstation VDAs

- \* Tests used 2 Cloud Connector VMs, each had 4 vCPU and 4 GB RAM. Based on the recommendation for an N + 1 configuration, production environments should include 3 Cloud Connector VMs that meet these specifications.
- \* Local Host Cache service peak consumed 91% of CPU resources and there was an average of 563 MB available memory.
- \* It took approximately 10 minutes from when the High Availability Service detected an outage for all VDAs to reregister with the High Availability Service, which is now the broker. We measured from the time the High Availability Service entered outage mode until the High Availability Service was ready to broker sessions again.
- 20,000 sessions launched against 500 server VDAs
  - \* Tests used 2 Cloud Connector VMs, each had 4 vCPU and 4 GB RAM. Based on the recommendation for an N + 1 configuration, production environments should include 3 Cloud Connector VMs that meet these specifications.
  - \* Local Host Cache service peak consumed 90% of CPU resources and there was an average of 471 MB available memory.
  - \* It took approximately 8 minutes from when the High Availability Service detected an outage for all VDAs to reregister with the High Availability Service. We measured from the time the High Availability Service entered outage mode until the High Availability Service was ready to broker sessions again.
- 1,000 sessions launched against 1,000 workstation VDAs
  - \* Tests used 2 Cloud Connector VMs, each had 2 vCPU and 4 GB RAM. Based on the recommendation for an N + 1 configuration, production environments should include 3 Cloud Connector VMs that meet these specifications.
  - \* Local Host Cache service peak consumed 95% of CPU resources and there was an average of 589 MB available memory
  - \* It took approximately 7 minutes from when the High Availability Service detected an outage for all VDAs to reregister with the High Availability Service, which is now the broker. We measured from the time the High Availability Service entered outage mode until the High Availability Service was ready to broker sessions again.

## Environment Overview



Citrix Cloud manages Cloud Connector services, and the customer manages the machines.

## Test methodology

We conducted tests by adding load, and then measuring the performance of the environment components:

- CPU
- memory
- database load
- Citrix Remote Broker Provider service
- Citrix High Availability Service

We collected performance data, logon time, or both. In certain cases, proprietary Citrix simulation tools were used to simulate VDAs and sessions. The simulation tools are designed to exercise Citrix components the same way that traditional VDAs and sessions do, without the same resource requirements to host real sessions and VDAs.

Local Host Cache supports one elected High Availability Service per zone, not per site. For example, if you have five zones, one Cloud Connector is elected as the broker in each zone. The Citrix Config Synchronizer service is responsible for importing the Citrix-managed site database. Every configuration sync creates a database, so initial configurations are needed, such as compiling stored procedures the first time the database is used. We executed all tests after a configuration sync.

## Session launch tests

On customer-managed StoreFront servers, we started 5,000 and 20,000 session tests. The monitoring tools collect StoreFront logon time, resource enumeration, and ICA file retrieval.

Citrix uses simulation tools to facilitate high-volume user testing. The simulation tools, which are proprietary to Citrix, allow us to run the tests on less hardware than is required to run tests using real sessions at these levels (5,000 and 20,000 sessions). These simulated sessions go through the normal StoreFront logon, resource enumeration, and ICA file retrieval, but do not start active desktops. Instead, the simulation tool reports to the ICA stack that the session has launched and all communication between the broker agent and the broker service is consistent with that of an actual session. Performance metrics are gathered from Citrix Cloud Connectors. To determine how the environment responded to session launches, a sustained concurrency of 25 session launches was maintained at any given time throughout the duration of the test. The measurements therefore show results of a system under load throughout the test.

## Test results

### Session launch

The following tables compare session launch tests between Local Host Cache outage mode active and Local Host Cache outage mode inactive after a new configuration synchronization import. Each table shows the results for the number of sessions launched in the test.

#### 5,000 workstation VDA sessions

	Local Host Cache outage mode Inactive (Normal Operations) / Average Timing	Local Host Cache outage mode Active / Average Timing
Authenticate	193 ms	95 ms
Enumerate	697 ms	75 ms
Total logon time	890 ms	170 ms
Retrieve ICA File	4,191 ms	156 ms

#### 20,000 server VDA Sessions

	Local Host Cache outage mode Inactive (Normal Operations) / Average Timing	Local Host Cache outage mode Active / Average Timing
Authenticate	135 ms	112 ms
Enumerate	317 ms	91 ms
Total logon time	452 ms	203 ms

	Local Host Cache outage mode Inactive (Normal Operations) / Average Timing	Local Host Cache outage mode Active / Average Timing
Retrieve ICA File	762 ms	174 ms

- 5,000 workstation VDA session launch test
  - There were approximately 30 ms of latency between the Citrix Cloud Connectors and Citrix Delivery Controller while Local Host Cache outage mode was inactive.
  - There is a 720 ms difference in the logon process with Local Host Cache outage mode active versus inactive, while the StoreFront is under load.
  - The largest time difference is in the retrieval of the ICA file, which is 4 seconds. This is largely because the Cloud Connector is performing the brokering, whereas normally the StoreFront traffic traverses through the Cloud Connectors to the Citrix Delivery Controller in Azure and back.
- 20,000 server VDA session launch test
  - There is a 249 ms difference in the logon process with Local Host Cache outage mode active versus inactive, while the StoreFront is under load.
  - The difference in the retrieval of the ICA file is about 1 second.
- Compared to the 5,000-workstation VDA session launch, the 20,000-session launch test contains only 500 server VDAs, resulting in fewer calls from the Citrix Delivery Controller to the VDAs, which leads to lower response times.

### Average CPU usage comparison

Session launch test		Average CPU %	Peak CPU %
5,000 workstation VDA sessions	Connector 1	8.3	38.2
	Connector 2	8.4	33.3
5,000 workstation VDA sessions - Local Host Cache outage mode active	Connector 1 (elected High Availability Service)	42	<b>91</b>
	Connector 2	0.8	5
20,000 server VDA sessions	Connector 1	23	62
	Connector 2	23	55

Session launch test		Average CPU %	Peak CPU %
20,000 server VDA sessions - Local Host Cache outage mode active	Connector 1 (elected High Availability Service)	57	<b>90</b>
	Connector 2	0.8	6.6

- The table compares Citrix Cloud Connector CPU usage with Local Host Cache outage mode active and Local Host Cache mode inactive during 5,000 workstation VDA and 20,000 server VDA session launch tests.
- All Cloud Connectors are 4 vCPU and 4 GB RAM
- The elected High Availability Service machines peaked at 91% and 90% overall CPU respectively. It is worth noting that, while the non-elected High Availability Service does not have much usage, it may become the active if the elected High Availability Service has a failure. It is therefore critical for the Cloud Connectors to have identical Cloud Connector specifications.

### Available memory usage

Session launch test		Average Available Memory (working set MB)	Peak Available Memory (working set MB)
5,000 workstation VDA sessions	Connector 1	636	657
	Connector 2	786	801
5,000 workstation VDA sessions - Local Host Cache outage mode active	Connector 1 (elected High Availability Service)	563	618
	Connector 2	912	918
20,000 server VDA sessions	Connector 1	1030	1195
	Connector 2	1178	1329
20,000 server VDA sessions - Local Host Cache outage mode active	Connector 1 (elected High Availability Service)	471	687

Session launch test	Average Available Memory (working set MB)	Peak Available Memory (working set MB)
Connector 2	1210	1227

- The table compares available memory usage with Local Host Cache outage mode active and Local Host Cache mode inactive during 5,000 workstation VDA and 20,000 server VDA session launch tests.
- The number of sessions decreases the amount of available memory.
- There is a 54.35% (559 MB) increase in memory usage with 20,000 server VDA sessions when Local Host Cache outage mode is active, mainly due to SQL server memory consumption.

### Cloud Connector CPU usage by component

Session launch test	Component	Average CPU %	Peak CPU %
5,000 workstation VDA sessions	Connector 1 LSASS	2.4	10.7
	Connector 1 XaXdCloudProxy	3.5	18.5
	Connector 2 LSASS	2.5	12.9
	Connector 2 XaXdCloudProxy	3.5	21.2
5,000 workstation VDA sessions Local Host Cache outage mode active	Connector 1 (elected High Availability Service) LSASS	12.9	29.5
	Connector 1 (elected High Availability Service) HighAvailabilityService	14.7	49.7
20,000 server VDA sessions	Connector 1 LSASS	7	12.2
	Connector 1 XaXdCloudProxy	8.7	15.5
	Connector 2 LSASS	7	12.5

Session launch test	Component	Average CPU %	Peak CPU %
	Connector 2 XaXdCloudProxy	9	15.7
20,000 sessions Local Host Cache outage mode active	Connector 1 (elected High Availability Service) LSASS	4.3	17.2
	Connector 1 (elected High Availability Service) High Availability Service	4.5	18.2

- The preceding table shows the processes that consume the most overall CPU resources when Local Host Cache outage mode is active, compared to when Local Host Cache outage mode is inactive during 5,000 workstation VDA and 20,000 server VDA session launch tests.
- The Citrix Remote Broker Provider service (XaXdCloudProxy) is the top CPU consumer when Local Host Cache outage mode is inactive.
- LSASS (Local Security Authority Subsystem Service) uses CPU during session logons. All authentications from Citrix-managed services must traverse the Citrix Cloud Connectors to communicate with the customer-managed Active Directory.
- The Citrix High Availability Service is used to broker the sessions, resulting in higher CPU usage when Local Host Cache outage mode is active. Also, CPU usage peaked to 49.7% during the 5,000 workstation VDA session launch, while the usage was only 18.25% during the 20,000 server VDA session launch (500 VDAs). The difference is due to the number of VDAs.
- Cloud Connector 2 did not have any meaningful metrics, as it was not the elected High Availability Service.

### VDA reregistration time while switching to Local Host Cache

During a Delivery Controller outage, the 5,000 workstation VDAs must reregister with the elected Local Host Cache broker. This reregistration time was ~10 minutes. The reregistration time for 500 server VDAs was ~8 minutes.

Number of VDAs	reregistration time
5,000 workstation VDAs	~10 minutes
500 server VDAs	~8 minutes



### Outage timings

Outage event	Number of VDAs	Time
Enter outage mode		10 minutes
Reregistration time to elected High Availability Service	500	~8 minutes
	5000	~10 minutes
Exit outage mode		10 minutes
Reregistration time to Citrix Delivery Controller	500	~1.5 minutes
	5000	~5.5 minutes

- There is a total of 20 minutes to enter (10 minutes) and exit (10 minutes) outage mode, due to the number of Citrix Delivery Controller health checks required. The time required to reregister the VDAs adds to the overall outage time.
- If the network is going up and down repeatedly, forcing an outage until the network issues resolve prevents continuous transition between normal and outage modes.

### Database and High Availability Service metrics with Local Host Cache

Session launch test	Average High Availability Service Database Transactions/sec	Peak High Availability Service Database Transactions/sec
5,000 workstation VDA sessions	436	1344
20,000 server VDA sessions	590	2061

The preceding table shows the number of database transactions per second on the elected High Availability Service.

### StoreFront CPU usage comparison

Session launch test	Average CPU %	Peak CPU %
5,000 workstation VDA sessions	4.5	32.4
5,000 server VDA sessions Local Host Cache outage mode	13.8	32.6
20,000 server VDA sessions	11.4	22.1
20,000 server VDA sessions - Local Host Cache outage mode	18.6	33.2

- The preceding table compares StoreFront CPU usage when Local Host Cache outage mode is active to when Local Host Cache mode is inactive during 5,000 workstation VDA and 20,000 server VDA session launch tests.
- The StoreFront machine has the following specifications: Windows 2012 R2, 8 vCPU (2 sockets, 4 cores each), 8 GB RAM
- When Local Host Cache outage mode is active, there is approximately a 9% increase in average CPU usage with the 5,000 workstation VDA and about a 7% increase with the 20,000 server VDA session launch tests. The increase is mostly because the IIS worker processes more requests when Local Host Cache outage mode is active. There is more CPU usage because StoreFront is processing session launches at a faster rate than when outage mode is inactive.

### StoreFront available memory usage comparison

Session launch test	Average Available Memory (working set MB)	Peak Available Memory (working set MB)
5,000 workstation VDA sessions	5731	6821
5,000 workstation VDA sessions Local Host Cache outage mode	5345	5420
20,000 server VDA sessions	4671	4924
20,000 server VDA sessions - Local Host Cache outage mode	4730	5027

- The preceding table compares the StoreFront available memory usage when Local Host Cache outage mode is active and when Local Host Cache mode is inactive during 5,000 workstation VDA and 20,000 server VDA session launch tests.
- When Local Host Cache mode is active, there is a 6.73% increase in memory usage during the 5,000 workstation VDA session launch test.


The following table compares outage mode active vs inactive after a new configuration synchronization import, launching 1,000 sessions to 1,000 workstation VDAs with Local Host Cache, and using Citrix Cloud Connectors configured with 2 vCPU VMs.

### Session launch comparison

	Local Host Cache outage mode inactive (normal operations)	Local Host Cache outage mode active
Authenticate	359 ms	89 ms
Enumerate	436 ms	180 ms
Total logon time	795 ms	269 ms
Retrieve ICA File	804 ms	549 ms


- While the StoreFront is under load, there is a 526 ms difference in the logon process when Local Host Cache outage mode is active compared to when Local Host Cache mode is inactive.
- There is a 255 ms difference in the retrieval of the ICA file when Local Host Cache outage mode is active compared to when Local Host Cache mode is inactive. The difference increases with the number of sessions.

### Average CPU usage comparison



The elected High Availability Service peaked to 95% overall CPU, which indicates that 1,000 workstation VDA is an optimal configuration for a 2 vCPU Cloud Connector VM.


**Average memory usage comparison**




The preceding graph displays a comparison of Citrix Cloud Connector available usage when Local Host

Cache outage mode is active versus inactive, during a 1,000 workstation VDA session launch. There is not a significant difference in memory based on the Local Host Cache outage mode.

### Cloud Connector CPU usage by component comparison



The preceding graph displays the processes that consume the most CPU resources while Local Host Cache outage mode is inactive.




- The preceding graph displays the processes that consume the most CPU resources when Local Host Cache outage mode is active.
- Connector 2 did not have any meaningful metrics.

### **VDA reregistration time while switching to Local Host Cache**

During a Delivery Controller outage, the 1000 workstation VDAs must reregister with the elected Local Host Cache broker. The reregistration time was ~7 minutes.

### **Database and High Availability Service metrics with Local Host Cache**



The preceding graph displays the number of database transactions per second on the elected High Availability Service.

### Impact with increasing number of zones on database import times

An extra zone (with a pair of its own Cloud Connectors) was added to the test site to understand the impact. The first zone consists of 5,500 unique objects (2 catalogs). The secondary zone is a mirror of the first zone, and has its own unique objects, totaling 11,000 objects. It is important to note that Local Host Cache is recommended only for zones with no more than 10,000 objects. Before we added the secondary zone, database import time on the Cloud Connectors was about 4 minutes, 20 seconds. After we added the secondary zone and populated it with 11,000 objects, the import time increased to by ~30 seconds to ~4 minutes, 50 seconds. Adding more catalogs has marginal impact on import times. The largest contributing factors to performance degradation and increased import times are based on the number of assigned machines, users, and remote PCs. Additionally, 5,500 objects were split between 2 zones and the import time remained the same

Number of zones	Total Number of Objects	Import time
1	5,500	4 minutes 20 seconds
2	11,000	4 minutes 50 seconds
2	5,500	4 minutes 20 seconds

## Connector Sizing Guidance

For optimal performance, the following are the recommended configurations for Citrix Cloud Connector when Local Host Cache mode is enabled.

Recommendation 1: to support 1,000 workstation VDAs using Local Host Cache mode with Citrix Cloud Connector

- 2 Windows 2012 R2 VMs, each allocated with 2 vCPU (1 socket, 2 cores), 4 GB RAM
- This recommended sizing is based on the peak Citrix Cloud Connector overall 95% CPU usage and 589 MB average available memory while Local Host Cache mode is active

Recommendation 2: to support 5,000 workstation VDAs OR 500 server VDAs using Local Host Cache with Citrix Cloud Connector

- 2 Windows 2012 R2 VMs, each allocated with 4 vCPU (1 socket, 4 cores), 4 GB RAM
- This recommended sizing is based on
  - 5,000 workstation VDA sessions launched with Local Host Cache mode active
    - \* Overall 91% peak CPU usage
    - \* 563 MB average available memory
  - 20,000 server VDA sessions launched with Local Host Cache mode active
    - \* Overall 90% peak CPU usage
    - \* 471 MB average available memory

See the white paper Citrix Cloud Virtual Apps and Desktops service sizing and scalability considerations for more information about general scalability sizing.

## Test environment

The test environment employed internally developed, proprietary testing tools, and VMs configured to the specifications in the following sections.

### Tools used

We used an internal testing tool to collect performance data and metrics from the machines under test and to drive the session launches. The in-house testing tool orchestrates user session launches to the Citrix Virtual Apps and Desktops environment. The testing tool also provides a central location where we gather response time data and performance metrics. In essence, the test tool administers the tests and collects the results.

### Test configuration – Citrix Virtual Apps and Desktops service

The following is a list of the machine and OS specifications used with the Citrix Virtual Apps and Desktops service testing.



- **Cloud Connectors:**
  - 2 Windows 2012 R2 VMs, each allocated 4 vCPU (1 socket, 4 cores), 4 GB RAM
  - 2 Windows 2012 R2 VMs, each allocated 2 vCPU (1 socket, 2 cores), 4 GB RAM
- **StoreFront (Customer-managed):** Windows 2012 R2, 8 vCPU (2 sockets, 4 cores each), 8 GB RAM
- **Hypervisor:** Citrix XenServer 7.0 + updates, 5x HP Blade BL 460C Gen 9, 2x Intel E5-2620 CPU, 256 GB RAM
- **Hypervisor Storage:** 2 TB NFS share on NetApp 3250
- **VDA:** Windows 2012 R2

### Data Collection

We collect the following metrics from each test: average overall CPU, memory, component (cloud processes) usage increase.

- VDA reregistration time when switching to the elected Local Host Cache High Availability Service
- Database and High Availability Service metrics when Local Host Cache outage mode is active
- Session launch comparison, average timings for
  - Authentication
  - Enumeration
  - ICA file retrieval
- Impact to database synchronization times while increasing the number of zones
  - Time required to synchronize after a configuration change

### RAM size considerations

SQL Server Express LocalDB can use up to 1.2 GB of RAM (up to 1 GB for the database cache, plus 200 MB for running SQL Server Express LocalDB). The High Availability Service (the Local Host Cache broker) can use up to 1 GB of RAM if an outage lasts for an extended interval with many logons occurring (for example, 12 hours with 10K users). These memory requirements are in addition to the normal RAM requirements for the Cloud Connector. Consider increasing the total amount of RAM capacity.

### CPU core and socket configuration considerations

A Cloud Connector's CPU configuration, particularly the number of cores available to the SQL Server Express LocalDB, directly affects Local Host Cache performance, even more than memory allocation. This CPU overhead is observed only during the outage period when the database is unreachable and the Local Host Cache broker is active.

While SQL Server Express LocalDB can use multiple cores (up to 4), it's limited to only a single socket. Adding more sockets does not improve the performance (for example, having 4 sockets with 1 core each).

## Storage considerations

As users access resources during an outage, the Local Host Cache database grows. For example, during a logon/logoff test running at 10 logons per second, the database grew by 1 MB every 2 to 3 minutes. When normal operation resumes, the Local Host Cache database is recreated when a configuration change is detected. The Local Host Cache broker must have sufficient space on the drive where the Local Host Cache database is installed to allow for the database growth during an outage. Local Host Cache also incurs more I/O during an outage: approximately 3 MB of writes per second, with several hundred thousand reads.

## Use Search in the Full Configuration management interface

August 30, 2021


### Introduction

Use the search feature to view information about specific machines, sessions, machine catalogs, applications, or delivery groups. After selecting **Search** in the **Manage > Full Configuration** left pane, you have several options:

- Use tabs to list machines by type (single-session or multi-session OS), or list all sessions.
- Enter the name in the search box.
- Select the filter icon to perform an advanced search. Select the down arrow to display a list of search properties. Select the plus sign to build an expression from the properties in the list.

To save your search, select the ellipsis (...) icon and then select **Save As**. The search appears in the **Saved searches** list. (To access the list, select the search box.) To delete saved searches, select the search box and select **Clear**.

When you use filters to perform an advanced search, the **Add filters** window appears in the foreground, leaving the background view unchanged. After you select **Search**, the matched search results appear, with the filter criteria appearing next to **Filter**. When you close the **Add filters** window, the results remain there. To clear the filters, select the X icon next to the filter criteria.



## Search for machine catalogs or delivery groups

You cannot perform searches from the **Machine Catalogs** or **Delivery Groups** node because the Search box is not available. Use the **Search** node instead to search for machine catalogs or delivery groups. On the **Search** node, select the filter icon, add filters as follows, and then select **Search**.

The image contains two screenshots of the 'Add filters' dialog box. The first screenshot shows the 'Machine Catalog' dropdown menu selected and highlighted with a red box. The second screenshot shows the 'Delivery Group' dropdown menu selected and highlighted with a red box. Both screenshots show the 'Contains' operator, an empty search input field, and buttons for 'Search', 'Save', 'Clear', 'Delete', and 'Cancel'.

To show more search criteria in the display, select the plus sign. Remove search criteria by selecting the trash can icon.

## Tips to enhance a search

Consider the following tips when using the Search feature:

- On the **Search** node, select any column to sort items.
- To show more characteristics to include in the display where you can search and sort, select **Columns to Display** or right-click any column and select **Columns to Display**. In the **Columns to Display** window, select the check box next to the items you want to display and select **Save** to exit.

### Note:

Items that degrade performance are marked with the **Degrades performance** label.

- To locate a user device connected to a machine, use **Client (IP)** and **Is**, and enter the device IP address.
- To locate active sessions, use **Session State**, **Is**, and **Connected**.
- To list all machines in a delivery group, select **Delivery Groups** in the left pane. Select the group, and then select **View Machines** from the action bar or from the context menu.

Keep the following considerations in mind when performing sort operations:

- As long as the number of items does not exceed 5,000, you can click any column to sort the items in it. When the number exceeds 5,000, you can sort only by name or by current user (depending on which tab you are on). To enable sorting, use filters to reduce the number of items to 5,000 or fewer.
- When the number of items is greater than 500 but no more than 5,000:
  - We cache all data locally to improve sort performance. On the **Single-session OS Machines** and **Multi-session OS Machines** tabs, we cache the data the first time you click a column (any column except the **Name** column) to sort. On the **Sessions** tab, we cache the data the first time you click a column (any column except the **Current User** column) to sort. As a result, the sort takes longer to complete. For faster performance, sort by name or current user, or use filters to reduce the number of items.
  - The following message under the table indicates that the data is cached: Last refreshed: < the time when you refreshed the table>. In that case, sort operations are based on items that were loaded previously. Those items might not be up to date. To bring them up to date, click the refresh icon.

## Virtual IP and virtual loopback

February 9, 2021

### Important:

Windows 10 Enterprise multi-session doesn't support Remote Desktop IP Virtualization (Virtual IP) and we don't support Virtual IP nor virtual loopback on Windows 10 Enterprise multi-session.

Virtual IP and virtual loopback features are supported on Windows Server 2016 machines. These features do not apply to Windows desktop OS machines.

The Microsoft virtual IP address feature provides a published application with a unique dynamically assigned IP address for each session. The Citrix virtual loopback feature allows you to configure applications that depend on communications with localhost (127.0.0.1 by default) to use a unique virtual loopback address in the localhost range (127.\*).

Certain applications, such as CRM and Computer Telephony Integration (CTI), use an IP address for addressing, licensing, identification, or other purposes and thus require a unique IP address or a loopback address in sessions. Other applications may bind to a static port, so attempts to launch additional instances of an application in a multiuser environment will fail because the port is already in use. For such applications to function correctly in a Citrix Virtual Apps environment, a unique IP address is required for each device.

Virtual IP and virtual loopback are independent features. You can use either or both.

Administrator action synopsis:

- To use Microsoft virtual IP, enable and configure it on the Windows server. (Citrix policy settings are not needed.)
- To use Citrix virtual loopback, configure two settings in a Citrix policy.

## Virtual IP

When virtual IP is enabled and configured on the Windows server, each configured application running in a session appears to have a unique address. Users access these applications on a Citrix Virtual Apps server in the same way they access any other published application. A process requires virtual IP in either of the following cases:

- The process uses a hard-coded TCP port number
- The process uses Windows sockets and requires a unique IP address or a specified TCP port number

To determine if an application needs to use virtual IP addresses:

1. Obtain the TCPView tool from Microsoft. This tool lists all applications that bind specific IP addresses and ports.
2. Disable the Resolve IP Addresses feature so that you see the addresses instead of host names.
3. Launch the application and use TCPView to see which IP addresses and ports are opened by the application and which process names are opening these ports.
4. Configure any processes that open the IP address of the server, 0.0.0.0, or 127.0.0.1.
5. To ensure that an application does not open the same IP address on a different port, launch an additional instance of the application.

## How Microsoft Remote Desktop (RD) IP virtualization works

- Virtual IP addressing must be enabled on the Microsoft server.

For example, in a Windows Server 2016 environment, from Server Manager, expand **Remote Desktop Services > RD Session Host Connections** to enable the RD IP Virtualization feature and configure the settings to dynamically assign IP addresses using the Dynamic Host Configuration Protocol (DHCP) server on a per-session or per-program basis. See the Microsoft documentation for instructions.

- After the feature is enabled, at session start-up, the server requests dynamically assigned IP addresses from the DHCP server.
- The RD IP Virtualization feature assigns IP addresses to remote desktop connections per-session or per-program. If you assign IP addresses for multiple programs, they share a per-session IP address.

- After an address is assigned to a session, the session uses the virtual address rather than the primary IP address for the system whenever the following calls are made: `bind`, `closesocket`, `connect`, `WSAConnect`, `WSAAccept`, `getpeername`, `getsockname`, `sendto`, `WSASendTo`, `WSASocketW`, `gethostbyaddr`, `getnameinfo`, `getaddrinfo`.

When using the Microsoft IP virtualization feature within the Remote Desktop session hosting configuration, applications are bound to specific IP addresses by inserting a “filter” component between the application and Winsock function calls. The application then sees only the IP address it should use. Any attempt by the application to listen for TCP or UDP communications is bound to its allocated virtual IP address (or loopback address) automatically, and any originating connections opened by the application originate from the IP address bound to the application.

In functions that return an address (such as `GetAddrInfo()`, which is controlled by a Windows policy), if the local host IP address is requested, virtual IP looks at the returned IP address and changes it to the virtual IP address of the session. Applications that attempt to get the IP address of the local server through such name functions see only the unique virtual IP address assigned to that session. This IP address is often used in subsequent socket calls, such as `bind` or `connect`. For more information about Windows policies, see [RDS IP Virtualization in Windows Server](#).

Often, an application requests to bind to a port for listening on the address 0.0.0.0. When an application does this and uses a static port, you cannot launch more than one instance of the application. The virtual IP address feature also looks for 0.0.0.0 in these call types and changes the call to listen on the specific virtual IP address, which enables more than one application to listen on the same port on the same computer because they are all listening on different addresses. The call is changed only if it is in an ICA session and the virtual IP address feature is enabled. For example, if two instances of an application running in different sessions both try to bind to all interfaces (0.0.0.0) and a specific port (such as 9000), they are bound to `VIPAddress1:9000` and `VIPAddress2:9000` and there is no conflict.

## Virtual loopback

Enabling the Citrix virtual IP loopback policy settings allows each session to have its own loopback address for communication. When an application uses the localhost address (default = 127.0.0.1) in a Winsock call, the virtual loopback feature simply replaces 127.0.0.1 with 127.X.X.X, where X.X.X is a representation of the session ID + 1. For example, a session ID of 7 is 127.0.0.8. In the unlikely event that the session ID exceeds the fourth octet (more than 255), the address rolls over to the next octet (127.0.1.0), to the maximum of 127.255.255.255.

A process requires virtual loopback in either of the following cases:

- The process uses the Windows socket loopback (localhost) address (127.0.0.1)
- The process uses a hard-coded TCP port number

Use the [virtual loopback policy settings](#) for applications that use a loopback address for interprocess communication. No additional configuration is required. Virtual loopback has no dependency on

Virtual IP, so you do not have to configure the Microsoft server.

- Virtual IP loopback support. When enabled, this policy setting allows each session to have its own virtual loopback address. This setting is disabled by default. The feature applies only to applications specified with the Virtual IP virtual loopback programs list policy setting.
- Virtual IP virtual loopback programs list. This policy setting specifies the applications that use the virtual IP loopback feature. This setting applies only when the Virtual IP loopback support policy setting is enabled.

### Related feature

You can use the following registry settings to ensure that virtual loopback is given preference over virtual IP; this is called preferred loopback. However, proceed with caution:

- Use preferred loopback only if both Virtual IP and virtual loopback are enabled; otherwise, you may have unintended results.
- Editing the registry incorrectly can cause serious problems that may require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

Run regedit on the servers where the applications reside.

- HKEY\_LOCAL\_MACHINE\SOFTWARE\Wow6432Node\Citrix\VIP
- Name: PreferLoopback, Type: REG\_DWORD, Data: 1
- Name: PreferLoopbackProcesses, Type: REG\_MULTI\_SZ, Data: <list of processes>

## Sessions

June 30, 2020

Maintaining session activity is critical to providing the best user experience. Losing connectivity due to unreliable networks, highly variable network latency, and range limitations of wireless devices can lead to user frustration. Being able to move quickly between workstations and access the same set of applications each time they log on is a priority for many mobile workers such as health-care workers in a hospital.

The features described in this article optimize the reliability of sessions, reduce inconvenience, downtime, and loss of productivity; using these features, mobile users can roam quickly and easily between devices.

You can also log a user off of a session, disconnect a session, and configure session prelaunch and linger; see [Manage Delivery Groups](#).

## Session reliability

Session Reliability keeps sessions active and on the user's screen when network connectivity is interrupted. Users continue to see the application they are using until network connectivity resumes.

This feature is especially useful for mobile users with wireless connections. For example, a user with a wireless connection enters a railroad tunnel and momentarily loses connectivity. Ordinarily, the session is disconnected and disappears from the user's screen, and the user has to reconnect to the disconnected session. With Session Reliability, the session remains active on the machine. To indicate that connectivity is lost, the user's display freezes and the cursor changes to a spinning hourglass until connectivity resumes on the other side of the tunnel. The user continues to access the display during the interruption and can resume interacting with the application when the network connection is restored. Session Reliability reconnects users without reauthentication prompts.

Citrix Workspace app users cannot override the Controller setting.

You can use Session Reliability with Transport Layer Security (TLS). TLS encrypts only the data sent between the user device and Citrix Gateway.

Enable and configure Session Reliability with the following policy settings:

- The Session reliability connections policy setting allows or prevents session reliability.
- The Session reliability timeout policy setting has a default of 180 seconds, or three minutes. Although you can extend the amount of time Session Reliability keeps a session open, this feature is designed for user convenience and therefore does not prompt the user for reauthentication. As you extend the amount of time a session is kept open, chances increase that a user may get distracted and walk away from the user device, potentially leaving the session accessible to unauthorized users.
- Incoming session reliability connections use port 2598, unless you change the port number in the Session reliability port number policy setting.
- If you do not want users to be able to reconnect to interrupted sessions without having to reauthenticate, use the Auto Client Reconnect feature. You can configure the Auto client reconnect authentication policy setting to prompt users to reauthenticate when reconnecting to interrupted sessions.

If you use both Session Reliability and Auto Client Reconnect, the two features work in sequence. Session Reliability closes, or disconnects, the user session after the amount of time you specify in the Session reliability timeout policy setting. After that, the Auto Client Reconnect policy settings take effect, attempting to reconnect the user to the disconnected session.

## Auto Client Reconnect

With the Auto Client Reconnect feature, Citrix Workspace app can detect unintended disconnections of ICA sessions and reconnect users to the affected sessions automatically. When this feature is en-



abled on the server, users do not have to reconnect manually to continue working.

For application sessions, Citrix Workspace app attempts to reconnect to the session until there is a successful reconnection or the user cancels the reconnection attempts.

For desktop sessions, Citrix Workspace app attempts to reconnect to the session for a specified period of time, unless there is a successful reconnection or the user cancels the reconnection attempts. By default, this period of time is five minutes. To change this period of time, edit this registry on the user device:

```
HKLM\Software\Citrix\ICA Client\TransportReconnectRetryMaxTimeSeconds;  
DWORD;<seconds>
```

where `seconds` is the number of seconds after which no more attempts are made to reconnect the session.

Enable and configure Auto Client Reconnect with the following policy settings:

- **Auto client reconnect:** Enables or disables automatic reconnection by Citrix Workspace app after a connection has been interrupted.
- **Auto client reconnect authentication:** Enables or disables the requirement for user authentication after automatic reconnection.
- **Auto client reconnect logging:** Enables or disables logging of reconnection events in the event log. Logging is disabled by default. When enabled, the server's system log captures information about successful and failed automatic reconnection events. Each server stores information about reconnection events in its own system log; the site does not provide a combined log of reconnection events for all servers.

Auto Client Reconnect incorporates an authentication mechanism based on encrypted user credentials. When a user initially logs on, the server encrypts and stores the user credentials in memory, and creates and sends a cookie containing the encryption key to Citrix Workspace app. Citrix Workspace app submits the key to the server for reconnection. The server decrypts the credentials and submits them to Windows logon for authentication. When cookies expire, users must reauthenticate to reconnect to sessions.

Cookies are not used if you enable the auto client reconnection authentication setting. Instead, users are presented with a dialog box to users requesting credentials when Citrix Workspace app attempts to reconnect automatically.

For maximum protection of user credentials and sessions, use encryption for all communication between clients and the Site.

Disable Auto Client Reconnect on Citrix Workspace app for Windows by using the `icaclient.adm` file. For more information, see the documentation for your Citrix Workspace app for Windows version.

Settings for connections also affect Auto Client Reconnect:

- By default, Auto Client Reconnect is enabled through policy settings at the Site level, as described above. User reauthentication is not required. However, if a server's ICA TCP connection is configured to reset sessions with a broken communication link, automatic reconnection does not occur. Auto Client Reconnect works only if the server disconnects sessions when there is a broken or timed out connection. In this context, the ICA TCP connection refers to a server's virtual port (rather than an actual network connection) that is used for sessions on TCP/IP networks.
- By default, the ICA TCP connection on a server is set to disconnect sessions with broken or timed out connections. Disconnected sessions remain intact in system memory and are available for reconnection by Citrix Workspace app.
- The connection can be configured to reset or log off sessions with broken or timed-out connections. When a session is reset, attempting to reconnect initiates a new session; rather than restoring a user to the same place in the application in use, the application is restarted.
- If the server is configured to reset sessions, Auto Client Reconnect creates a new session. This process requires users to enter their credentials to log on to the server.
- Automatic reconnection can fail if Citrix Workspace app or the plug-in submits incorrect authentication information, which might occur during an attack or the server determines that too much time has elapsed since it detected the broken connection.

## ICA Keep-Alive

Enabling the ICA Keep-Alive feature prevents broken connections from being disconnected. When enabled, if the server detects no activity (for example, no clock change, no mouse movement, no screen updates), this feature prevents Remote Desktop Services from disconnecting that session. The server sends keep-alive packets every few seconds to detect if the session is active. If the session is no longer active, the server marks the session as disconnected.

### Important:

ICA Keep-Alive works only if you are not using Session Reliability. Session Reliability has its own mechanisms to prevent broken connections from being disconnected. Configure ICA Keep-Alive only for connections that do not use Session Reliability.

ICA Keep-Alive settings override keep-alive settings that are configured in Microsoft Windows Group Policy.

Enable and configure ICA Keep-Alive with the following policy settings:

- **ICA keep alive timeout:** Specifies the interval (1-3600 seconds) used to send ICA keep-alive messages. Do not configure this option if you want your network monitoring software to close inactive connections in environments where broken connections are so infrequent that allowing users to reconnect to sessions is not a concern.

The default interval is 60 seconds: ICA Keep-Alive packets are sent to user devices every 60 seconds. If a user device does not respond in 60 seconds, the status of the ICA sessions changes to disconnected.

- **ICA keep alives:** Sends or prevents sending ICA keep-alive messages.

## Workspace control

Workspace control lets desktops and applications follow a user from one device to another. This ability to roam enables a user to access all desktops or open applications from anywhere simply by logging on, without having to restart the desktops or applications on each device. For example, workspace control can assist health-care workers in a hospital who need to move quickly among different workstations and access the same set of applications each time they log on. If you configure workspace control options to allow it, these workers can disconnect from multiple applications at one client device and then reconnect to open the same applications at a different client device.

Workspace control affects the following activities:

- **Logging on:** By default, workspace control enables users to reconnect automatically to all running desktops and applications when logging on, bypassing the need to reopen them manually. Through workspace control, users can open disconnected desktops or applications, as well as any that are active on another client device. Disconnecting from a desktop or application leaves it running on the server. If you have roaming users who need to keep some desktops or applications running on one client device while they reconnect to a subset of their desktops or applications on another client device, you can configure the logon reconnection behavior to open only the desktops or applications that the user disconnected from previously.
- **Reconnecting:** After logging on to the server, users can reconnect to all of their desktops or applications at any time by clicking Reconnect. By default, Reconnect opens desktops or applications that are disconnected, plus any that are currently running on another client device. You can configure Reconnect to open only those desktops or applications that the user disconnected from previously.
- **Logging off:** For users opening desktops or applications through StoreFront, you can configure the Log Off command to log the user off from StoreFront and all active sessions together, or log off from StoreFront only.
- **Disconnecting:** Users can disconnect from all running desktops and applications at once, without needing to disconnect from each individually.

Workspace control is available only for Citrix Workspace app users who access desktops and applications through a Citrix StoreFront connection. By default, workspace control is disabled for virtual desktop sessions, but is enabled for hosted applications. Session sharing does not occur by default between published desktops and any published applications running inside those desktops.

User policies, client drive mappings, and printer configurations change appropriately when a user

moves to a new client device. Policies and mappings are applied according to the client device where the user is currently logged on to the session. For example, if a health care worker logs off from a client device in the emergency room of a hospital and then logs on to a workstation in the hospital's x-ray laboratory, the policies, printer mappings, and client drive mappings appropriate for the session in the x-ray laboratory go into effect at the session startup.

You can customize which printers appear to users when they change locations. You can also control whether users can print to local printers, how much bandwidth is consumed when users connect remotely, and other aspects of their printing experiences.

For information about enabling and configuring workspace control for users, see the StoreFront documentation.

## **Session roaming**

By default, sessions roam between client devices with the user. When the user launches a session and then moves to another device, the same session is used and applications are available on both devices. The applications follow, regardless of the device or whether current sessions exist. In many cases, printers and other resources assigned to the application also follow.

While this default behavior offers many advantages, it might not be ideal in all cases. You can prevent session roaming using the PowerShell SDK.

Example 1: A medical professional is using two devices, completing an insurance form on a desktop PC, and looking at patient information on a tablet.

- If session roaming is enabled, both applications appear on both devices (an application launched on one device is visible on all devices in use). This might not meet security requirements.
- If session roaming is disabled, the patient record does not appear on the desktop PC, and the insurance form does not appear on the tablet.

Example 2: A production manager launches an application on the PC in his office. The device name and location determine which printers and other resources are available for that session. Later in the day, he goes to an office in the next building for a meeting that will require him to use a printer.

- If session roaming is enabled, the production manager would probably be unable to access the printers near the meeting room, because the applications he launched earlier in his office resulted in the assignment of printers and other resources near that location.
- If session roaming is disabled, when he logs on to a different machine (using the same credentials), a new session is started, and nearby printers and resources will be available.

## Configure session roaming

To configure session roaming, use the following entitlement policy rule cmdlets with the “SessionReconnection” property. Optionally, you can also specify the “LeasingBehavior” property.

For desktop sessions:

```
Set-BrokerEntitlementPolicyRule <Delivery-Group-name> -SessionReconnection  
<value> -LeasingBehavior Allowed|Disallowed
```

For application sessions:

```
Set-BrokerAppEntitlementPolicyRule <Delivery-Group-name> -SessionReconnection  
<value> -LeasingBehavior Allowed|Disallowed
```

Where `value` can be one of the following:

- **Always:** Sessions always roam, regardless of the client device and whether the session is connected or disconnected. This is the default value.
- **DisconnectedOnly:** Reconnect only to sessions that are already disconnected; otherwise, launch a new session. (Sessions can roam between client devices by first disconnecting them, or using Workspace Control to explicitly roam them.) An active connected session from another client device is never used; instead, a new session is launched.
- **SameEndpointOnly:** A user gets a unique session for each client device they use. This completely disables roaming. Users can reconnect only to the same device that was previously used in the session.

The “LeasingBehavior” property is described below.

### Effects from other setting:

Disabling session roaming is affected by the application limit “Allow only one instance of the application per user” in the application’s properties in the Delivery Group.

- If you disable session roaming, then disable the “Allow only one instance ...” application limit.
- If you enable the “Allow only one instance ...” application limit, do not configure either of the two values that allow new sessions on new devices.

## Logon interval

If a virtual machine containing a desktop VDA closes before the logon process completes, you can allocate more time to the process. The default for 7.6 and later versions is 180 seconds (the default for 7.0-7.5 is 90 seconds).

On the machine (or the master image used in a machine catalog), set the following registry key:

Key: `HKLM\SOFTWARE\Citrix\PortICA`

- Value: `AutoLogonTimeout`

- Type: `DWORD`
- Specify a decimal time in seconds, in the range 0-3600.

If you change a master image, update the catalog.

This setting applies only to VMs with single-session desktop (workstation) VDAs. Microsoft controls the logon timeout on machines with multi-session server VDAs.

## Tags

July 19, 2021

### Introduction

Tags are strings that identify items such as machines, applications, desktops, delivery groups, application groups, and policies. After creating a tag and adding it to an item, you can tailor certain operations to apply to only items that have a specified tag.

- Tailor search displays in the Full Configuration management interface.

For example, to display only applications that have been optimized for testers, create a tag named “test” and then add (apply) it to those applications. You can now filter the search with the tag “test”.

- Publish applications from an application group or specific desktops from a delivery group, considering only a subset of the machines in selected delivery groups. This is called a *tag restriction*.

With tag restrictions, you can use your existing machines for more than one publishing task, saving the costs associated with deploying and managing more machines. A tag restriction can be thought of as subdividing (or partitioning) the machines in a delivery group. Its functionality is similar, but not identical, to worker groups in XenApp releases earlier than 7.x.

Using an application group or desktops with a tag restriction or can be helpful when isolating and troubleshooting a subset of machines in a delivery group.

Details and examples of using a tag restriction are described later in this article.

- Schedule periodic restarts for a subset of machines in a delivery group.

Using a tag restriction for machines enables you to use new PowerShell cmdlets to configure multiple restart schedules for subsets of machines in a delivery group. For examples and details, see [Manage delivery groups](#).

- Tailor the application (assignment) of Citrix policies to machines in delivery groups, delivery group types, or OUs that have (or don't have) a specified tag.

For example, if you want to apply a Citrix policy only to the more powerful workstations, add a tag named “high power” to those machines. Then, on the **Assign Policy** page of the Create Policy wizard, select that tag and the **Enable** check box. You can also add a tag to a delivery group and then apply a Citrix policy to that group. For details, see [Create policies](#).

You can apply tags to:

- Machines
- Applications
- Delivery groups
- Application groups

You can configure a tag restriction can be configured when creating or editing the following in the Full Configuration management interface:

- A desktop in a shared delivery group
- An application group

### **Tag restrictions for a desktop or an application group**


A tag restriction involves several steps:

- Create the tag and then add (apply) it to machines.
- Create or edit a group with the tag restriction (in other words, restrict launches to machines with tag x).

A tag restriction extends the broker’s machine selection process. The broker selects a machine from an associated delivery group, subject to access policy, configured user lists, zone preference, and launch readiness, plus the tag restriction (if present). For applications, the broker falls back to other delivery groups in priority order, applying the same machine selection rules for each considered delivery group.

#### **Example 1: Simple layout**

This example introduces a simple layout that uses tag restrictions to limit which machines are considered for certain desktop and application launches. There is one shared delivery group, one published desktop, and one application group configured with two applications.



- Tags have been added to each of the three machines (VDA 101-103).
- The desktop in the delivery group was created with a tag restriction named **Red**. So, that desktop can be launched only on machines in that delivery group that have the tag **Red**: VDA 101 and 102.
- The application group was created with the **Orange** tag restriction. So, each of its applications (**Calculator** and **Notepad**) can be launched only on machines in that delivery group that have the tag **Orange**: VDA 102 and 103.


Machine VDA 102 has both tags (**Red** and **Orange**), so it can be considered for launching the applications and the desktop.

### Example 2: More complex layout

This example contains several application groups that were created with tag restrictions. This results in the ability to deliver more applications with fewer machines than would otherwise be needed if you used only delivery groups.

How to configure example 2 shows the steps used to create and apply the tags, and then configure the tag restrictions in this example.





This example uses 10 machines (VDA 101-110), one delivery group (D01), and three application groups (A100, A200, A300). By applying tags to each machine and then specifying tag restrictions when creating each application group:

- Accounting users in the group can access the apps they need on five machines (VDA 101–105)
- CAD designers in the group can access the apps they need on five machines (VDA 106-110)
- Users in the group who need Office applications can access the Office apps on 10 machines (VDA 101-110)

Only 10 machines are used, with only one delivery group. Using delivery groups alone (without application groups) would require twice as many machines, because a machine can belong to only one delivery group.

### Manage tags and tag restrictions

Tags are created, added (applied), edited, and deleted from selected items through the **Manage Tags** action in the Full Configuration management interface.

(Exception: Tags used for policy assignments are created, edited, and deleted through the **Manage Tags** action. However, tags are applied (assigned) when you create the policy. See [Create policies](#) for details.)

Tag restrictions are configured when you create or edit desktops in delivery groups, and when you create and edit application groups.


### **Use the Manage Tags feature**

From **Manage > Full Configuration**, select the items to which you want to apply a tag. The items include:

- One or more machines
- One or more applications
- A desktop, a delivery group, or an application group
- A machine catalog

Then select **Manage Tags** in the action bar. The **Manage Tags** dialog box lists all existing tags, not just those for the items you selected.

- An enabled check box indicates that the tag has already been added to the selected items. (In the screen capture below, the selected machine has a tag named “Tag1” applied.)
- If you select more than one item, a check box containing a hyphen indicates that some but not all selected items have that tag added.



The following actions are available from the **Manage Tags** dialog box. Review Cautions when working with tags.

- **To create a tag:**

Select **Create**. Enter a name and description. Tag names must be unique and are not case-sensitive. Then select **OK**.

Creating a tag does not automatically apply it to any items you have selected. Use the check boxes to apply the tag.

- **To add (apply) one or more tags:**

Enable the check box next to the tag name. A check box containing a hyphen indicates that some, but not all selected items already have the tag applied. When you select multiple items and a tag's check box has a hyphen, changing it to a check mark affects all selected machines.

If you attempt to add a tag to machines, and that tag is used as a restriction in an application

group, you are warned that the action can make those machines available for launch. If that's what you intend, proceed.

- **To remove one or more tags:**

Clear the check box next to the tag name. A check box containing a hyphen indicates that some, but not all selected items already have the tag applied. When you select multiple items and a tag's check box has a hyphen, clearing the check box removes the tag from all selected machines.

If you try to remove a tag restriction from a machine, you're warned that the action can affect the machines considered for launch. If that's what you intend, proceed.

- **To edit a tag:**

Select a tag and then select **Edit**. Enter a new name, description, or both. You can edit only one tag at a time.

- **To delete one or more tags:**

Select the tags and then select **Delete**. The **Delete Tag** dialog box indicates how many items currently use the selected tags (for example "2 machines"). Select an item to display more information (for example, the names of the two machines that have the tag applied). Confirm whether you want to delete the tags.

You cannot delete a tag that is used as a restriction. First, edit the application group and remove the tag restriction or select a different tag.

When you're done in the **Manage Tags** dialog box, select **Save**.

To see if a machine has any tags applied: Select **Delivery Groups** in the left pane. Select a delivery group then select **View Machines** in the action bar. Select a machine and then select the **Tags** tab on the **Details** pane.

## Manage tag restrictions

Configuring a tag restriction is a multi-step process: You first create the tag and add/apply it to machines. Then, you add the restriction to the application group or the desktop.

- **Create and apply the tag:**

Create the tag and then add (apply) it to the machines that the tag restriction will affect, using **Manage Tags** actions.

- **To add a tag restriction to an application group:**

Create or edit the application group. On the **Delivery Groups** page, select **Restrict launches to machines with the tag** and then select the tag from the list.

- **To change or remove the tag restriction on an application group:**

Edit the group. On the **Delivery Groups** page, either select a different tag from the list or remove the tag restriction entirely by clearing **Restrict launches to machines with the tag**.

- **To add a tag restriction to a desktop:**

Create or edit a delivery group. Select **Add** or **Edit** on the **Desktops** page. In the **Add Desktop** dialog box, select **Restrict launches to machines with the tag** and then select the tag from the menu.

- **To change or remove the tag restriction on a delivery group:**

Edit the group. On the **Desktops** page, select **Edit**. In the dialog box, either select a different tag from the list or remove the tag restriction entirely by clearing **Restrict launches to machines with the tag**.

### **Cautions when working with tags**

A tag applied to an item can be used for different purposes. Keep in mind that adding, removing, and deleting a tag can have unintended effects. You can use a tag to sort machine displays when using search in the Full Configuration management interface. You can use the same tag as a restriction when configuring an application group or a desktop. That action limits launch consideration to only machines in specified delivery groups that have that tag.

If you add a tag to machines after that tag is configured as a desktop or application group tag restriction, you are warned that might make the machines available for launching more applications or desktops. If that's what you intend, proceed. If not, cancel the operation.

For example, let's say you create an application group with the **Red** tag restriction. Later, you add several other machines in the same delivery groups used by that application group. If you then try to add the **Red** tag to those machines, you see a message similar to: "The tag **Red** is used as a restriction on the following application groups. Adding this tag might make the selected machines available to launch applications in this application group." You can then confirm or cancel adding that tag to those additional machines.

Similarly, when a tag is used in an application group to restrict launches, you cannot delete the tag until you edit the group and remove it as a restriction. (If you were allowed to delete that tag, it might result in allowing applications to launch on all machines in the delivery groups associated with the application group.) The same prohibition against deleting a tag applies if the tag is being used as a restriction for desktop launches. After you edit the application group or desktops in the delivery group to remove that tag restriction, you can delete the tag.

All machines might not have the same sets of applications. A user can belong to more than one application group, each with a different tag restriction and different or overlapping sets of machines from delivery groups. The following table lists how machine considerations are decided.

When an application has been added to	These machines in the selected delivery groups are considered for launch
One application group with no tag restriction	Any machine.
One application group with tag restriction A	Machines that have tag A applied.
Two application groups, one with tag restriction A and the other with tag restriction B	Machines that have tag A and tag B. If none is available, then machines that have tag A or tag B.
Two application groups, one with tag restriction A and the other with no tag restriction	Machines that have tag A. If none is available, then any machine.

If you used a tag restriction in a machine restart schedule, any changes you make that affect tag applications or restrictions affect the next machine restart cycle. It does not affect any restart cycles that is in progress while the changes are being made.

## How to configure example 2

The following sequence shows the steps to create and apply tags, and then configure tag restrictions for the application groups illustrated in the earlier second example.

VDA and applications have already been installed on the machines and the delivery group has been created.

Create and apply tags to the machines:

1. From **Manage > Full Configuration**, select **Delivery Groups** in the left pane. Select delivery group **D01** and then select **View Machines** in the action bar.
2. Select machines VDA 101-105 and then select **Manage Tags** in the action bar.
3. In the **Manage Tags** dialog box, select **Create**. Create a tag named **CADApps**. Select **OK**.
4. Select **Create** again and create a tag named **OfficeApps**. Select **OK**.
5. Add (apply) the newly created tags to the selected machines by enabling the check boxes next to each tag's name (**CADApps** and **OfficeApps**). Then close the dialog box.
6. Select delivery group **D01**. Select **View Machines** in the action bar.
7. Select machines VDA 106-110 and then select **Manage Tags** in the action bar.
8. In the **Manage Tags** dialog box, select **Create**. Create a tag named **AcctgApps**. Select **OK**.
9. Apply the newly created **AcctgApps** tag and the **OfficeApps** tag to the selected machines by selecting the check boxes next to each tag's name. Then close the dialog box.

Create the application groups with tag restrictions.

1. From **Manage > Full Configuration**, select **Applications** in the left pane.

2. Select **Create Application Group** in the action bar. The wizard launches.
3. On the **Delivery Groups** page, select delivery group **D01**. Select **Restrict launches to machines with tag** and then select the **AcctgApps** tag from the list.
4. Complete the wizard, specifying the accounting users and the accounting applications. (When adding the application, choose the **From Start menu** source, which searches for the application on the machines that have the **AcctgApps** tag.) On the **Summary** page, name the group **A100**.
5. Repeat the preceding steps to create application group **A200**, specifying machines that have the **CADApps** tag, plus the appropriate users and applications.
6. Repeat steps to create application group **A300**, specifying machines that have the **OfficeApps** tag, plus the appropriate users and applications.

### Apply tags to machine catalogs

You can use **Manage > Full Configuration** or PowerShell to apply tags to machine catalogs.

- Using the management interface is described in [Manage tags](#). Catalog displays do not indicate whether tags are applied.
- To use PowerShell, see [Use PowerShell to apply tags to catalogs](#).

(If you're using **Manage > Legacy Cosnole**, you must use PowerShell to apply the tag to the catalog.)

Here are some examples of using tags with catalogs:

- A delivery group contains machines from several catalogs, but you want an operation (such as a restart schedule) to affect only the machines in a specific catalog. Applying a tag to that catalog accomplishes that goal.
- In an application group, you want to limit application sessions to machines in a specific catalog. Applying a tag to that catalog accomplishes that goal.

### Use PowerShell to apply tags to catalogs

The following PowerShell cmdlets are available:

- You can pass catalog objects to cmdlets such as **Add-BrokerTag** and **Remove-BrokerTag**.
- **Get-BrokerTagUsage** shows how many catalogs contain tags.
- **Get-BrokerCatalog** has a property named **Tags**.

For example, the following cmdlets add a previously created tag named **fy2018** to the catalog named **acctg**: `Get-BrokerCatalog -Name acctg | Add-BrokerTag fy2018`.

See the PowerShell cmdlet help for guidance and syntax.

### More information

Blog post: [How to assign desktops to specific servers](#).

## Zones

July 7, 2021

### Introduction

Citrix Virtual Apps and Desktops service deployments that span widely dispersed locations connected by a WAN can face challenges from network latency and reliability. Using zones can help users in remote regions connect to resources without necessarily forcing their connections to traverse large segments of the WAN. In a Citrix Virtual Apps and Desktops service environment, each resource location is considered a zone.

Zones can be helpful in deployments of all sizes. You can use zones to keep applications and desktops closer to users, which improves performance. Zones can be used for disaster recovery, geographically distant data centers, branch offices, a cloud, or an availability zone in a cloud.

Throughout this article, the term local refers to the zone being discussed. For example, “A VDA registers with a local Cloud Connector” means that a VDA registers with a Cloud Connector in the zone where the VDA is located.

### Differences from zones in on-premises Citrix Virtual Apps and Desktops environments

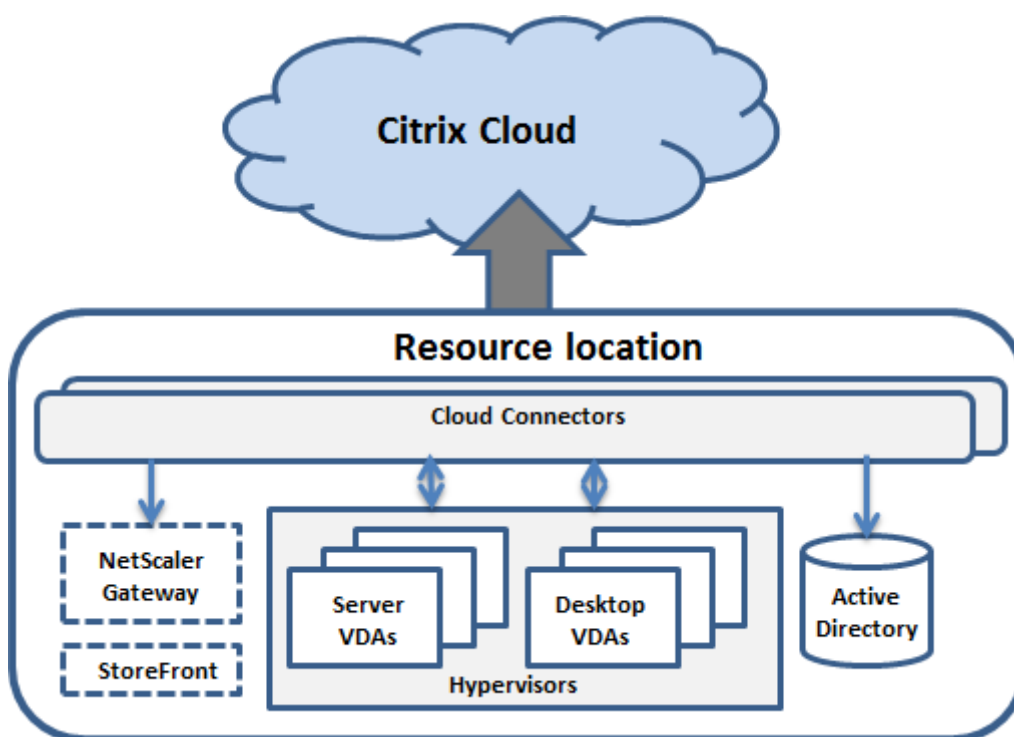
Zones in a Citrix Virtual Apps and Desktops service environment are similar, but not identical to zones in an on-premises Citrix Virtual Apps and Desktops deployment.

- In the Citrix Virtual Apps and Desktops service, zones are created automatically when you create a resource location and add a Cloud Connector to it. Unlike an on-premises deployment, a service environment does not classify zones as primary or satellite.
- In XenApp version 6.5 and earlier, zones included data collectors. The Citrix Virtual Apps and Desktops service does not use data collectors for zones. Also, failover and preferred zones work differently.

### What's in a zone

A zone is equivalent to a resource location. When you create a resource location and install a Cloud Connector, a zone is automatically created for you. Each zone can have a different set of resources, based on your unique needs and environment.





Each zone must always have at least one Cloud Connector, and preferably two or more, for redundancy.

You can place machine catalogs, hypervisors, host connections, users, and applications in a zone. A zone can also contain Citrix Gateway and StoreFront servers. To use the Local Host Cache feature, a zone must have a StoreFront server.

Zones support Workspace and the Citrix Gateway service.

Placing items in a zone affects how the service interacts with them and with other objects related to them.

- When a hypervisor connection is placed in a zone, it is assumed that all the hypervisors managed through that connection also reside in that zone.
- When a machine catalog is placed in a zone, it is assumed that all VDAs in the catalog are in the zone.
- Citrix Gateway instances can be added to zones. When you create a resource location, you are offered the option to add a Citrix Gateway. When a Citrix Gateway is associated with a zone, it is preferred for use when connections to VDAs in that zone are used.
- Ideally, Citrix Gateway in a zone is used for user connections coming into that zone from other zones or external locations. You can also use it for connections within the zone.
- After you create more resource locations and install Cloud Connectors in them (which automatically creates more zones), you can move resources between zones. This flexibility comes with the risk of separating items that work best in close proximity. For example, moving a catalog to a different zone than the connection (host) that creates the machines in the catalog, can affect

performance. So, consider potential unintended effects before moving items between zones. Keep a catalog and the host connection it uses in the same zone.

If the connection between a zone and Citrix Cloud fails, the Local Host Cache feature enables a Cloud Connector in the zone to continue brokering connections to VDAs in that zone. (The zone must have StoreFront installed.) For example, this is effective in an office where workers use the local StoreFront site to access their local resources, even if the WAN link connecting their office to the corporate network fails. For more information, see [Local Host Cache](#).

## Where VDAs register

VDAs must be minimum version 7.7 to use these zone registration features:

- A VDA in a zone registers with a local Cloud Connector.
  - As long as that Cloud Connector can communicate with Citrix Cloud, normal operations continue.
  - If that Cloud Connector is operational but cannot communicate with Citrix Cloud (and that zone has a local StoreFront), it enters Local Host Cache outage mode.
  - If a Cloud Connector fails, VDAs in that zone attempt to register with other local Cloud Connectors. A VDA in one zone never attempts to register with a Cloud Connector in another zone.
- If you add or remove a Cloud Connector in a zone (using the Citrix Cloud management console), and auto-update is enabled, VDAs in that zone receive updated lists of available local Cloud Connectors, so they know with whom they can register and accept connections from.
- If you move a machine catalog to another zone (using the Full Configuration management interface), the VDAs in that catalog re-register with Cloud Connectors in the zone where you moved the catalog. When you move a catalog, ensure you also move any associated host connection to the same zone.
- During an outage (when Cloud Connectors in a zone cannot communicate with Citrix Cloud), only the resources associated with machines that are registered in that zone are available.

## Zone preference

In a multi-zone Site, the zone preference feature offers the administrator more flexibility to control which VDA is used to launch an application or desktop.

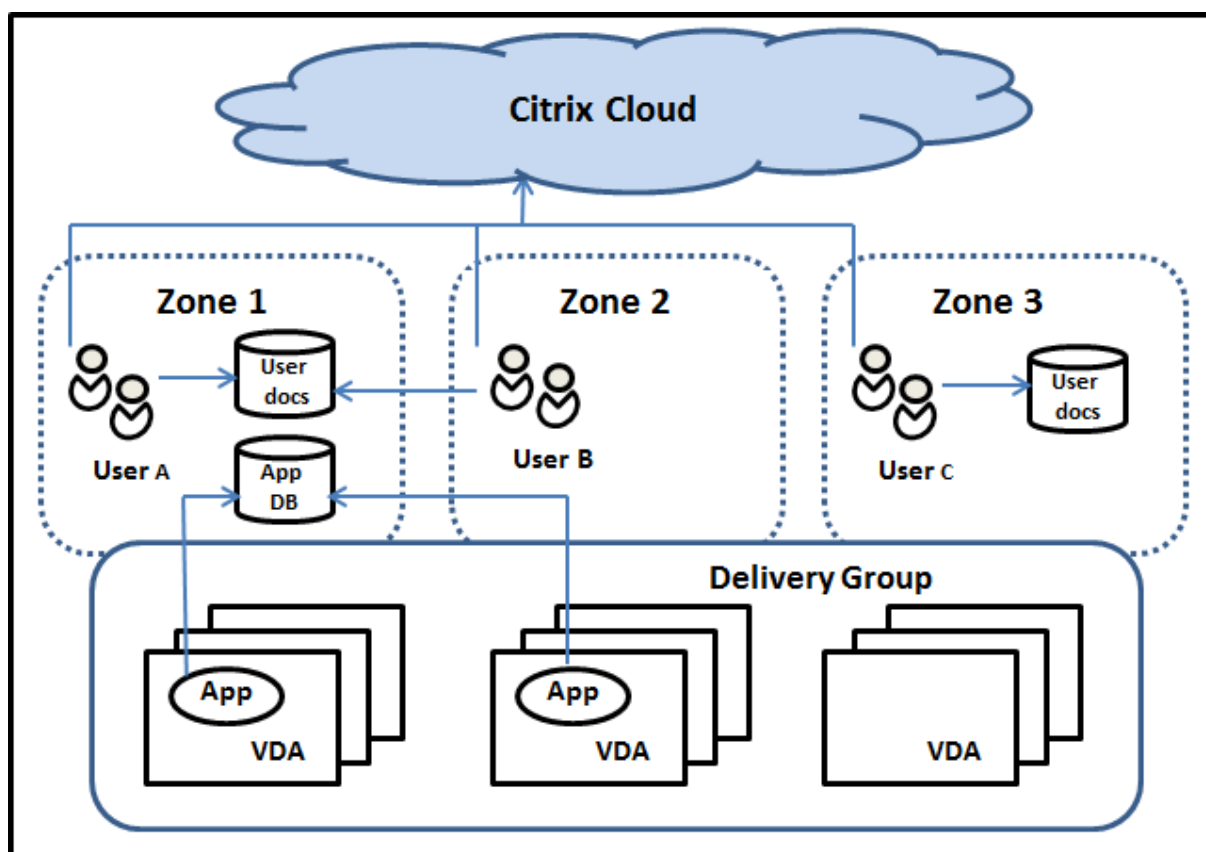
### How zone preference works

There are three forms of zone preference. You might prefer to use a VDA in a particular zone, based on:

- Where the application's data is stored. This is referred to as the application home.

- The location of the user's home data, such as a profile or home share. This is referred to as the user home.
- The user's current location (where the Citrix Workspace app is running). This is referred to as the user location. User location requires minimum StoreFront 3.7 and Citrix Gateway (formerly NetScaler Gateway) 11.0-65.x.

The following graphic shows an example multi-zone configuration.



In this example, VDAs are spread among three zones, but they are all in the same delivery group. Therefore, the Citrix Virtual Apps and Desktops service broker might have a choice which VDA to use for a user launch request. This example illustrates that users can be running their Citrix Workspace app endpoints at different locations. User A is using a device with Citrix Workspace app in zone 1. User B is using a device in zone 2. Similarly, a user's documents can be stored in different locations. Users A and B use a share located in zone 1. User C uses a share in zone 3. Also, one of the published applications uses a database located in zone 1.

You associate a user or application with a zone by configuring a home zone for the user or application. The broker then uses those associations to help select the zone where a session will be launched, if resources are available. You:

- Configure the home zone for a user by adding a user to a zone.
- Configure the home zone for an application by editing the application's properties.

A user or an application can have only one home zone at a time. (An exception for users can occur when multiple zone memberships occur because of user group membership. However, even in this case, the broker uses only one home zone.)

Although zone preferences for users and applications can be configured, the broker selects only one preferred zone for a launch. The default priority order for selecting the preferred zone is: application home > user home > user location. When a user launches an application:

- If that application has a configured zone association (an application home), then the preferred zone is the home zone for that application.
- If the application does not have a configured zone association, but the user does (a user home), then the preferred zone is the home zone for that user.
- If neither the application nor the user has a configured zone association, then the preferred zone is the zone where the user is running a Citrix Workspace app instance (the user location). If that zone is not defined, a random VDA and zone selection is used. Load balancing is applied to all VDAs in the preferred zone. If there is no preferred zone, load balancing is applied to all VDAs in the delivery group.

### Tailoring zone preference

When you configure (or remove) a home zone for a user or an application, you can also further restrict how zone preference is (or is not) used.

- **Mandatory user home zone use:** In a delivery group, you can specify “Launch the session in the user’s home zone (if the user has a home zone), with no failover to a different zone if resources are not available in the home zone.” This restriction is helpful if you want to avoid the risk of copying large profiles or data files between zones. In other words, you would rather deny a session launch than launch the session in a different zone.
- **Mandatory application home zone use:** Similarly, when you configure a home zone for an application, you can specify “launch the application only in that zone, with no failover to a different zone if resources are not available in the application’s home zone.”
- **No application home zone, and ignore configured user home zone:** If you do not specify a home zone for an application, you can also specify “do not consider any configured user zones when launching that application.” For example, use the user location zone preference if you want users to run a specific application on a VDA close to their machine, even though some users might have a different home zone.

### How preferred zones affect session use

When a user launches an application or desktop, the broker prefers using the preferred zone rather than using an existing session.

If the user launching an application or desktop already has a session that is suitable for the resource being launched (for example, can use session sharing for an application, or a session already running the resource being launched), but that session is on a VDA in a zone other than the preferred zone for the user/application, then the system might create a new session. This action satisfies launching in the correct zone (if it has available capacity), ahead of reconnecting to a session in a less-preferred zone for that user's session requirements.

To prevent an orphan session that can no longer be reached, reconnection is allowed to existing disconnected sessions, even if they are in a non-preferred zone.

The order of desirability for sessions to satisfy a launch is:

1. Reconnect to an existing session in the preferred zone.
2. Reconnect to an existing disconnected session in a non-preferred zone.
3. Start a new session in the preferred zone.
4. Reconnect to a connected existing session in a non-preferred zone.
5. Start a new session in a non-preferred zone.

### **Other zone preference considerations**

- If you configure a home zone for a user group (such as a security group), that group's users (through direct or indirect membership) are associated with the specified zone. However, a user can be a member of multiple security groups, and therefore might have a different home zone configured through other group membership. In such cases, determination of that user's home zone can be ambiguous.

If a user has a configured home zone that was not acquired through group membership, that zone is used for zone preference. Any zone associations acquired through group membership are ignored.

If the user has multiple different zone associations acquired solely through group membership, the broker chooses among the zones randomly. After the broker makes this choice, that zone is used for subsequent session launches, until the user's group membership changes.

- The user location zone preference requires detection of Citrix Workspace app on the endpoint device by the Citrix Gateway through which that device is connecting. The Citrix must be configured to associate ranges of IP addresses with particular zones. Discovered zone identity must be passed through StoreFront to the Citrix Virtual Apps and Desktops service.

Although written for on-premises use of zones, the [Zone Preference Internals](#) blog post contains relevant technical details.

## Permissions to manage zones

A Full Administrator can perform all supported zone management tasks. Moving items between zones does not require zone-related permissions (except zone read permission). However, you must have edit permission for the items you are moving. For example, to move a machine catalog from one zone to another, you must have edit permission for that catalog.

**If you use Citrix Provisioning:** The current Citrix Provisioning console is not aware of zones, so Citrix recommends using the **Manage > Full Configuration** interface to create machine catalogs that you want to place in specific zones. After you create the catalog, you can use the Citrix Provisioning console to provision machines in that catalog.

## Zone creation

When you create a resource location in Citrix Cloud and then add a Cloud Connector to that resource location, the Citrix Virtual Apps and Desktops service automatically creates and names a zone. You can optionally add a description later.

After you create more than one resource location (and the zones are created automatically), you can move resources from one zone to another.

Resource locations and zones are synchronized periodically, typically and approximately every five minutes. So, if you change a resource location's name in Citrix Cloud, that change is propagated to the associated zone within five minutes.

## Add or change a zone description

Although you cannot change a zone's name, you can add or change its description.

1. From **Manage > Full Configuration**, select **Zones** in the left pane.
2. Select a zone in the middle pane and then select **Edit Zone** in the action bar.
3. Add or change the zone description.
4. Select **OK** or **Apply**.

## Move resources from one zone to another zone

1. From **Manage > Full Configuration**, select **Zones** in the left pane.
2. Select a zone in the middle pane, and then select one or more items.
3. Either drag the items to the destination zone or select **Move Items** in the action bar, and then specify which zone to move them to. (Although you can select Cloud Connectors, you cannot actually move them to a different zone.)

A confirmation message lists the items you selected and asks if you are sure that you want to move all of them.

Remember: When a machine catalog uses a host connection to a hypervisor or cloud service, ensure that the catalog and the connection are in the same zone. Otherwise, performance can be affected. If you move one, move the other, too.

## Zone deletion

You cannot delete a zone. However, you can delete a resource location (after removing its Cloud Connectors). Deleting the resource location automatically deletes the zone.

- If the zone does not contain any items (such as catalogs, connections, applications, or users), the zone is deleted during the next synchronization between zones and resource locations. Synchronization occurs every five minutes.
- If the zone contains items, the zone is automatically deleted after all items are removed.

## Add a home zone for a user

Configuring a home zone for a user is also known as *adding a user to a zone*.

1. From **Manage > Full Configuration**, select **Zones** in the left pane.
2. Select a zone in the middle pane and then select **Add Users to Zone** in the action bar.
3. In the **Add Users to Zone** dialog box, select **Add**, and then select the users and user groups to add to the zone. If you specify users who already have a home zone, a message offers two choices: **Yes** = add only those users you specified who do not have a home zone; **No** = return to the user selection dialog.
4. Select **OK**.

For users with a configured home zone, you can require that sessions launch only from their home zone:

1. Create or edit a delivery group.
2. On the **Users** page, select the **Sessions must launch in a user's home zone, if configured** check box.

All sessions launched by a user in that delivery group must launch from machines in that user's home zone. If a user in the delivery group does not have a configured home zone, this setting has no effect.

## Remove a home zone for a user

This procedure is also known as removing a user from a zone.

1. From **Manage > Full Configuration**, select **Zones** in the left pane.
2. Select a zone in the middle pane and then select **Remove Users from Zone** in the action bar.

3. In the **Add Users to Zone** dialog box, select **Remove**, and then select the users and groups to remove from the zone. This action removes the users only from the zone. Those users remain in the delivery groups to which they belong.
4. Confirm the removal when prompted.

## Manage home zones for applications

Configuring a home zone for an application is also known as adding an application to a zone. By default, in a multi-zone environment, an application does not have a home zone.

An application's home zone is specified in the application's properties. You can configure application properties when you add the application to a group or later.

- When [creating a delivery group](#) or [adding applications to existing groups](#), select **Properties** on the **Applications** page of the wizard.
- To change an application's properties after the application is added, select **Zones** in the left pane. Select an application and then select **Properties** in the action bar.

On the **Zones** page of the application's properties/settings:

- If you want the application to have a home zone:
  - Select the **Use the selected zone to decide** radio button and then select the zone.
  - If you want the application to launch only from the selected zone (and not from any other zone), select the check box under the zone selection.
- If you do not want the application to have a home zone:
  - Select the **Do not configure a home zone** radio button.
  - If you do not want the broker to consider any configured user zones when launching this application, select the check box under the radio button. In this case, neither application nor user home zones are used to determine where to launch this application.

## Other actions that include specifying zones

If you have more than one zone, you can specify a zone when you add a host connection or create a catalog. Zones are listed alphabetically in selection lists. By default, the first alphabetical name is selected.

## Licenses

June 10, 2021

This article covers tasks and resources for Microsoft licenses and Citrix licenses.



## Configure a Microsoft RDS License Server for Windows Server workloads

This information applies when you are delivering Windows Server workloads.

This service accesses Windows Server remote session capabilities when delivering a Windows Server workload, such as Windows 2019. This typically requires a Remote Desktop Services client access license (RDS CAL). The VDA must be able to contact an RDS license server to request RDS CALs.

Install and activate the license server. For more information, see the Microsoft document [Activate the Remote Desktop Services License Server](#). For proof of concept environments, you can use the grace period provided by Microsoft.

With this method, you can have this service apply the license server settings. You can configure the license server and per user mode in the RDS console on the image. You can also configure the license server using Microsoft Group Policy settings. For more information, see the Microsoft document [License your RDS deployment with client access licenses \(CALs\)](#).

To configure the RDS license server using Microsoft Group Policy settings:

1. Install a Remote Desktop Services License Server on an available VM. The VM must always be available. The Citrix service workloads must be able to reach this license server.
2. Specify the license server address and per-user license mode using Microsoft Group Policy. For details, see the Microsoft document [Specify the Remote Desktop Licensing Model for an RD Session Host Server](#).

Windows 10 workloads require appropriate Windows 10 license activation. We recommend that you follow Microsoft documentation to activate Windows 10 workloads.

## Citrix license usage

For information about Citrix license usage, see:

- [Monitor license and active usage for cloud services](#)
- [Monitor license and active usage for Citrix Virtual Apps and Desktops service](#)

## User access

July 10, 2020

There are two primary components that provide access to applications and desktops in a Citrix Virtual Apps and Desktops service deployment:

- **Citrix Workspace:** Citrix Workspace is a complete digital solution that allows you to deliver secure access to the information, apps, and other content that are relevant to a person's role

in your organization. Users subscribe to the services you make available and can access them from anywhere, on any device. Citrix Workspace helps you organize and automate the most important details your users need to collaborate, make better decisions, and focus fully on their work.

There is zero effort to deploy Workspace, and it is kept evergreen by Citrix. Workspace is recommended for new and existing customers, previews, and proofs-of-concept.

- **An on-premises StoreFront:** Customers can also use an existing StoreFront to aggregate applications and desktops in Citrix Cloud. This use case offers greater security, including support for two-factor authentication, and prevents users from entering their password into the cloud service. It also allows customers to customize their domain names and URLs. This deployment type is recommended for any Citrix Virtual Apps and Desktops customers who already have StoreFront deployed.

See also Local Host Cache and StoreFront.

When users connect from outside the corporate firewall, Citrix Cloud can use Citrix Gateway (formerly NetScaler Gateway) technology to secure these connections with SSL. Citrix Gateway or the Citrix VPX virtual appliance is an SSL VPN appliance that is deployed in the demilitarized zone (DMZ). It provides a single secure point of access through the corporate firewall.

## Using Citrix Workspace

Access to Workspace is through <https://<customername>.cloud.com>. If needed, you can customize the <customername> portion of the workspace URL. You can then configure the connectivity for each resource location you want to use, so that end-users can access the resources in their workspace. End-users access their workspace using the latest version of Citrix Workspace app.

For more information about using Workspace, see:

- [Configure workspaces](#): For configuring access and customizations.
- [Secure workspaces](#): For configuring authentication.
- [Manage your Workspace experience](#): For understanding how end-users access their workspace and how it appears.

To provide remote access for end-users through Workspace, you can use either Citrix Gateway service or your own Citrix Gateway.

- To use the Citrix Gateway service:
  1. In **Citrix Cloud > Resource Locations**, select **Gateway** for the resource location you want to use.
  2. Select **Gateway Service** and then click **Save**.
  3. In **Citrix Cloud > Workspace Configuration > Service Integrations**, locate the Gateway service and select **Enable** from the ellipsis menu.

- To use your own Citrix Gateway:
  1. Set up Citrix Gateway as an ICA Proxy (No authentication or session policies are needed).
  2. Configure a resource location to use Citrix Gateway:
    - a) In **Citrix Cloud > Resource Locations**, select **Gateway** for the resource location you want to use.
    - b) Select **Traditional Gateway** and enter the external FQDN. Do not add a protocol. Ports are optional. Combination remote and internal access is not supported in Workspace.
  3. Bind Citrix Cloud Connectors as Secure Ticket Authority (STA) servers to Citrix Gateway. For details, see [CTX232640](#).

For more information about the Citrix Gateway service and Citrix Gateway, see [Citrix Gateway](#).

### Using an on-premises StoreFront

For information about configuring an on-premises StoreFront, see the [StoreFront documentation](#).

One benefit of using an existing StoreFront is that the Citrix Cloud Connector provides encryption of user passwords. The Cloud Connector encrypts credentials using AES-256, using a random-generated one-time key. This key is returned directly to Citrix Workspace app and never sent to the cloud. Citrix Workspace app then supplies it to the VDA during session launch to decrypt the credentials and provide a single sign-on experience into Windows.

- For transport, select HTTP and port 80. The StoreFront machine must be able to directly access the Cloud Connector through the FQDN (fully qualified domain name) provided. The Cloud Connector must be able to reach the Cloud NFuse/STA URL at (<https://<customername>.xendesktop.net/Scripts/wpnbr.dll> and [ctxsta.dll](#)).
- Add Cloud Connectors as Delivery Controllers for high availability.

Use the most recent version of StoreFront.

### External access

To provide external access through Citrix Gateway and on-premises StoreFront:

- Set up Citrix Gateway as usual, with authentication and session policies. See the [Citrix Gateway documentation](#) for details.
- Point your on-premises StoreFront store's Delivery Controllers to the Citrix Cloud Connectors. Bind Cloud Connectors as STA servers to Citrix Gateway.
- The Citrix Gateway must use the same STA URLs as StoreFront. If the gateway is not already configured to use the STA of an existing Citrix Virtual Apps and Desktops environment, Cloud Connectors can be used as a STA.

### Internal access

To provide internal access through an on-premises StoreFront, point the on-premises StoreFront store's Delivery Controllers to the Citrix Cloud Connectors.

### External and internal access

To provide external and internal access through Citrix Gateway and on-premises StoreFront:

- Set up Citrix Gateway as usual, with authentication and session policies. See the [Citrix Gateway documentation](#) for details.
- Bind Cloud Connectors as STA servers to Citrix Gateway.
- Point your on-premises StoreFront store's Delivery Controllers to the Cloud Connectors.

### Local Host Cache and StoreFront

Local Host Cache enables connection brokering operations in a Citrix Virtual Apps and Desktops service deployment to continue when Cloud Connectors cannot communicate with Citrix Cloud.

The Local Host Cache feature works only in resource locations containing a customer-deployed on-premises StoreFront. Local Host Cache does not support Workspace.

Each resource location must have a customer-deployed on-premises StoreFront. Verify that the resource location contains a local StoreFront that points to all the Cloud Connectors in that resource location.

For more information, see [Local Host Cache](#).

## Monitor

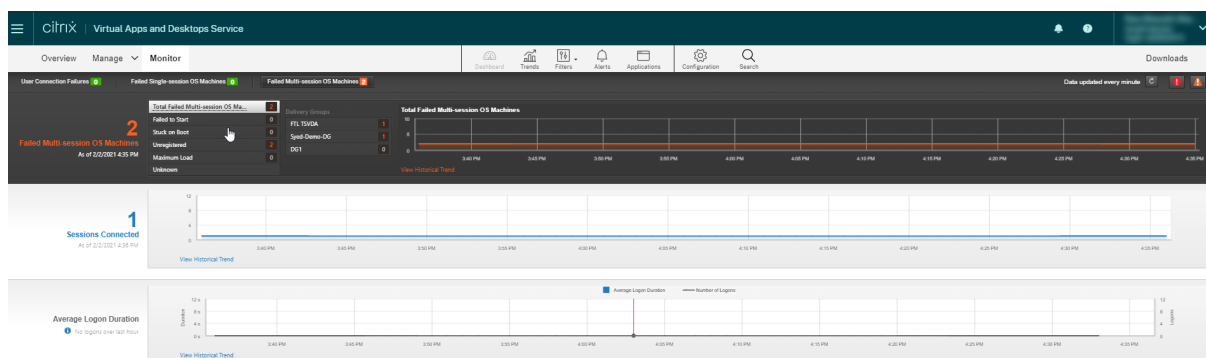
February 5, 2021

Administrators and help-desk personnel can monitor Citrix Virtual Apps and Desktops service from **Monitor**, the monitoring and troubleshooting console. The **Monitor** tab displays a dashboard to monitor, troubleshoot, and perform support tasks for subscribers.

#### Note:

Monitor is available as the Director console to monitor and troubleshoot Citrix Virtual Apps and Desktops [Current Release](#) and [LTSR](#) deployments.

To access **Monitor**, sign in to [Citrix Cloud](#). In the upper left menu, select **My Services > Virtual Apps and Desktops**. Click **Monitor**.



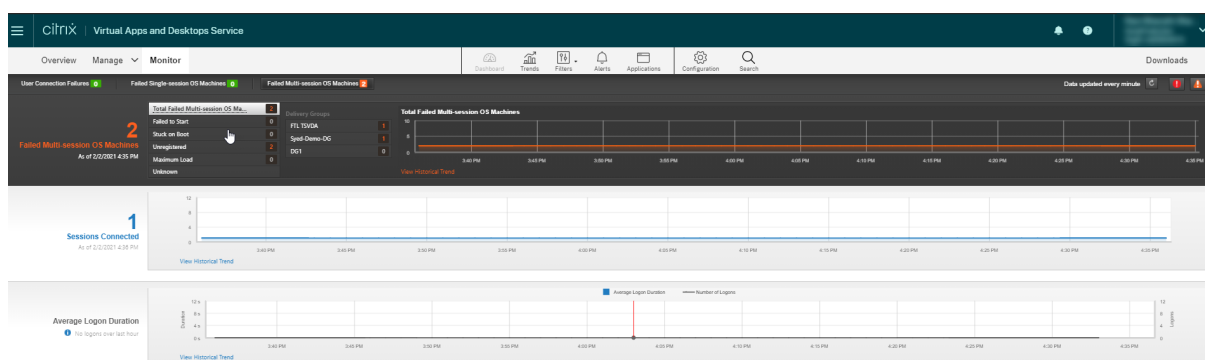
Monitor provides:

- Real-time data from the Broker Agent using a unified console integrated with Analytics and Performance Manager.
- Analytics includes performance management for health and capacity assurance, and historical trending to identify bottlenecks in your Citrix Virtual Apps or Desktops service environment.
- Historical data stored in the Monitor database to access the Configuration Logging database.
- Gain visibility into the end-user experience for virtual applications, desktops, and users for Citrix Virtual Apps or Desktops service.
- Monitor uses a troubleshooting dashboard that provides real-time and historical health monitoring of the Citrix Virtual Apps and Desktops service. This feature allows you to see failures in real time, providing a better idea of what the end users are experiencing.

## Site Analytics

September 14, 2020

The Monitor dashboard provides a centralized location to monitor the health and usage of a site.



If there are currently no failures and no failures have occurred in the past 60 minutes, panels stay collapsed. When there are failures, the specific failure panel automatically appears.

Panel	Description
User Connection Failures	Connection failures over the last 60 minutes. Click the categories next to the total number to view metrics for that type of failure. In the adjacent table, that number is broken out by delivery groups. Connection failures include failures caused by application limits being reached. For more information on application limits, see <a href="#">Applications</a> .
Failed Single session OS Machines or Failed Multi-session OS Machines	Total failures in the last 60 minutes broken out by delivery groups. Failures broken out by types, including failed to start, stuck on boot, and unregistered. For Multi-session OS machines, failures also include machines reaching maximum load.
Sessions Connected	Connected sessions across all delivery groups for the last 60 minutes.
Average Logon Duration	Log on data for the last 60 minutes. The large number on the left is the average logon duration across the hour. Log on data for VDAs earlier than XenDesktop 7.0 is not included in this average. For more information, see <a href="#">Diagnose user logon issues</a> .

**Note:**

If no icon appears for a particular metric, this indicates that this metric is not supported by the type of host you are using. For example, no health information is available for System Center Virtual Machine Manager (SCVMM) hosts, AWS and CloudStack.

Continue to troubleshoot issues using these options (which are documented below):

- [Control user machine power](#)
- [Prevent connections to machines](#)

**Monitor sessions**

If a session becomes disconnected, it is still active and its applications continue to run, but the user device is no longer communicating with the server.

Action	Description
View a user's currently connected machine or session	From the Activity Manager and User Details views, view the user's currently connected machine or session and a list of all machines and sessions to which this user has access. To access this list, click the session switcher icon in the user title bar. For more information, see <a href="#">Restore sessions</a> .
View the total number of connected sessions across all delivery groups	From the Dashboard, in the <b>Sessions Connected</b> pane, view the total number of connected sessions across all delivery groups for the last 60 minutes. Then click the large total number, which opens the Filters view, where you can display graphical session data based on selected delivery groups and ranges and usage across delivery groups.
End idle sessions	The Sessions Filters view displays data related to all active sessions. Filter the sessions based on Associated User, delivery group, Session State, and Idle Time greater than a threshold time period. From the filtered list, select sessions to log off or disconnect. For more information, see <a href="#">Troubleshoot applications</a> .

---

Action	Description
View data over a longer period	On the Trends view, select the <b>Sessions</b> tab to drill down to more specific usage data for connected and disconnected sessions over a longer period of time (that is, session totals from earlier than the last 60 minutes). To view this information, click <b>View historical trends</b> .

---

**Note:**

If the user device is running a legacy Virtual Delivery Agent (VDA), such as a VDA earlier than version 7, or a Linux VDA, Monitor cannot display complete information about the session. Instead, it displays a message that the information is not available.

**Desktop Assignment Rules limitation:**

Citrix Studio allows assignment of multiple Desktop Assignment Rules (DAR) for different users or user groups to a single VDA in the delivery group. StoreFront displays the assigned desktop with the corresponding **Display Name** as per the DAR for the logged in user. However, Monitor does not support DARs and displays the assigned desktop using the delivery group name regardless of the logged in user. As a result, you cannot map a specific desktop to a machine in Monitor.

You can map the assigned desktop displayed in StoreFront to the delivery group name displayed in Monitor using the following PowerShell command. Run the PowerShell command using Remote PowerShell SDK as described in the [blog](#).

```
1 Get-BrokerDesktopGroup | Where-Object {
2   $_.Uid -eq (Get-BrokerAssignmentPolicyRule | Where-Object {
3     $_.PublishedName -eq "<Name on StoreFront>" }
4   ).DesktopGroupUid }
5 | Select-Object -Property Name, Uid
```

**Disable the visibility of running applications in the Activity Manager**

By default, the Activity Manager displays a list of all running applications for a user's session. This information can be viewed by all administrators that have access to the Activity Manager feature. For Delegated Administrator roles, this includes Full Administrator, delivery group Administrator, and Help Desk Administrator.

To protect the privacy of users and the applications they are running, you can disable the Applications tab to list running applications. To do this, on the VDA, modify the registry key at HKEY\_LOCAL\_MACHINE\Software\Citrix\Director\TaskManagerDataDisplayed. By default, the key is



set to 1. Change the value to 0, which means the information is not collected from the VDA and hence not displayed in the Activity Manager.

**Warning:**

Editing the registry incorrectly can cause serious problems that might require you to reinstall your operating system. Citrix cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

**Session transport protocol**

View the transport protocol in use for the HDX connection type for the current session in the **Session Details** panel. This information is available for sessions launched on VDAs Version 7.13 or later.

The screenshot shows the Citrix Activity Manager interface. At the top right, there is a tab labeled 'Activity Manager' and a refresh icon. Below this, the 'Session Details' section is visible. It contains three buttons: 'Session Control' (with a dropdown arrow), 'Shadow', and 'Send Message'. The main area displays a list of session properties:

<b>ID</b>	7
<b>Session State</b>	Active
<b>Application State</b>	Desktop
<b>Anonymous</b>	No
<b>Time in state</b>	0 minutes
<b>Endpoint name</b>	XXXXXXXXXX
<b>Endpoint IP</b>	10.146.1.13
<b>Connection type</b>	HDX
<b>Protocol</b>	TCP
<b>Citrix Workspace App Version</b>	18.12.0.12
<b>ICA RTT</b>	n/a
<b>ICA Latency</b>	284 ms
<b>Launched via</b>	n/a
<b>Connected via</b>	XXXXXXXXXX

Below the session details, there are three tabs: 'Policies' (selected), 'Hosted Applications', and 'SmartAccess Filters'. Under the 'Policies' tab, two policies are listed: 'Policy1' and 'Policy0'.

- For **HDX** Connection type,
  - The Protocol is displayed as **UDP**, if EDT is used for the HDX connection.
  - The Protocol is displayed as **TCP**, if TCP is used for the HDX connection.
- For **RDP** Connection type, the Protocol is displayed as **n/a**.

When adaptive transport is configured, the session transport protocol dynamically switches between EDT (over UDP) and TCP, based on the network conditions. If the HDX session cannot be established

using EDT, it falls back to the TCP protocol.

For more information about adaptive transport configuration, see [Adaptive Transport](#).

## Export reports

You can export trends data to generate regular usage and capacity management reports. Export supports PDF, Excel, and CSV report formats. Reports in PDF and Excel formats contain trends represented as graphs and tables. CSV format reports contain tabular data that can be processed to generate views or can be archived.

To export a report:

1. Go to the **Trends** tab.
2. Set filter criteria and time period and click **Apply**. The trend graph and table are populated with data.
3. Click **Export** and enter name and format of the report.

Monitor generates the report based on the filter criteria you select. If you change the filter criteria, click **Apply** before you click **Export**.

### Note:

Export of a large amount of data causes a significant increase in memory and CPU consumption on the Monitor server, the Delivery Controller, and the SQL servers. The supported number of concurrent export operations and the amount of data that can be exported is set to default limits to achieve optimal export performance.

## Supported export limits

Exported PDF and Excel reports contain complete graphical charts for the selected filter criteria. However, tabular data in all report formats is truncated beyond the default limits on the number of rows or records in the table. The default number of records supported is defined based on the report format.

Report format	Default number of records supported
PDF	500
Excel	100,000
CSV	100,000 (10,000,000 in <b>Sessions</b> tab)

## Error Handling

Errors that you might encounter during an Export operation:

- **Director has timed out:** This error can occur due to network issues or high resource usage on the Director server or with the Monitor Service.
- **Monitor has timed out:** This error could occur due to network issues or high resource usage with the Monitor Service or on the SQL server.
- **Max concurrent Export or Preview operations ongoing:** Only one instance of Export or Preview can run at a specific time. If you get the **Max concurrent Export or Preview operations ongoing** error, try the next operation again later.

## Monitor hotfixes

To view the hotfixes installed on a specific machine VDA (physical or VM), choose the **Machine Details** view.

## Control user machine power states

To control the state of the machines that you select in Monitor, use the Power Control options. These options are available for Single session OS machines, but might not be available for Multi-session OS machines.

**Note:**

This functionality is not available for physical machines or machines using Remote PC Access.

Command	Function
<b>Restart</b>	Performs an orderly (soft) shutdown of the VM and all running processes are halted individually before restarting the VM. For example, select machines that appear in Monitor as “failed to start,” and use this command to restart them.
<b>Force Restart</b>	Restarts the VM without first performing any shut-down procedure. This command works in the same way as unplugging a physical server and then plugging it back in and turning it back on.
<b>Shut Down</b>	Performs an orderly (soft) shutdown of the VM. All running processes are halted individually.

---

Command	Function
<b>Force Shutdown</b>	Shuts down the VM without first performing any shut-down procedure. This command works in the same way as unplugging a physical server. It might not always shut down all running processes, and you risk losing data if you shut down a VM in this way.
<b>Suspend</b>	Suspends a running VM in its current state and stores that state in a file on the default storage repository. This option allows you to shut down the VM's host server and later, after rebooting it, resume the VM, returning it to its original running state.
<b>Resume</b>	Resumes a suspended VM and restores its original running state.
<b>Start</b>	Starts a VM when it is off (also called a cold start).

---

If power control actions fail, hover the mouse over the alert, and a pop-up message appears with details about the failure.

## Prevent connections to machines

Use maintenance mode to prevent new connections temporarily while the appropriate administrator performs maintenance tasks on the image.

When you enable maintenance mode on machines, no new connections are allowed until you disable it. If users are currently logged on, maintenance mode takes effect as soon as all users are logged off. For users who do not log off, send a message informing them that machines will be shut down at a certain time, and use the power controls to force the machines to shut down.

1. Select the machine, such as from the User Details view, or a group of machines in the Filters view.
2. Select **Maintenance Mode**, and turn on the option.

If a user tries to connect to an assigned desktop while it is in maintenance mode, a message appears indicating that the desktop is currently unavailable. No new connections can be made until you disable maintenance mode.

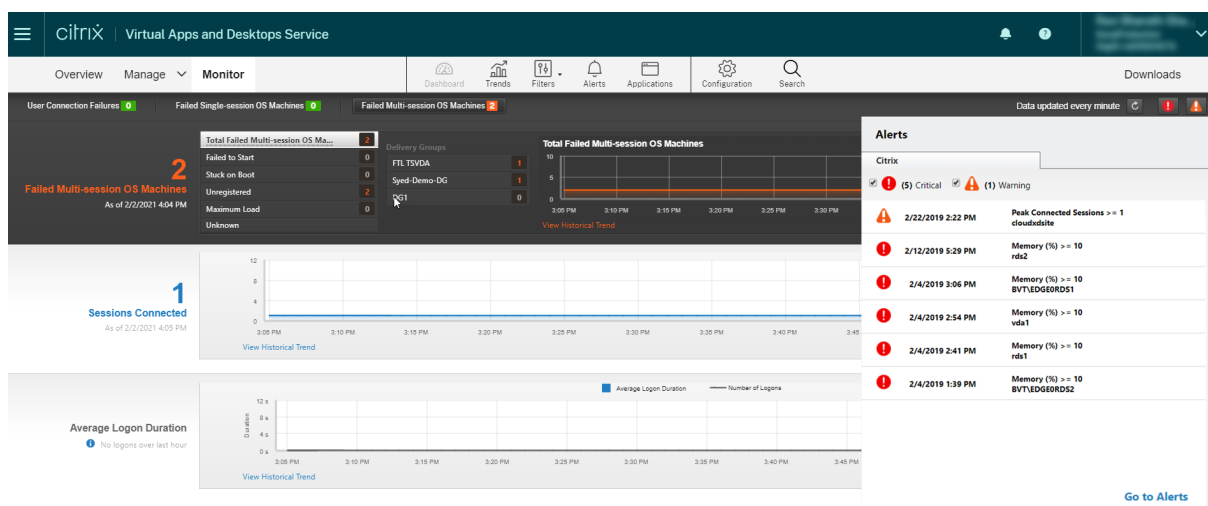
## Application Analytics

The **Applications** tab displays application-based analytics in a single, consolidated view to help analyze and manage application performance efficiently. You can gain valuable insight into the health and usage information of all applications published on the site. It shows metrics such as the probe results, number of instances per application, and faults and errors associated with the published applications. For more information, see the [Application Analytics](#) section in **Troubleshooting Applications**.

## Alerts and notifications

November 10, 2020

Alerts are displayed in Monitor on the dashboard and other high level views with warning and critical alert symbols. Alerts update automatically every minute; you can also update alerts on demand.

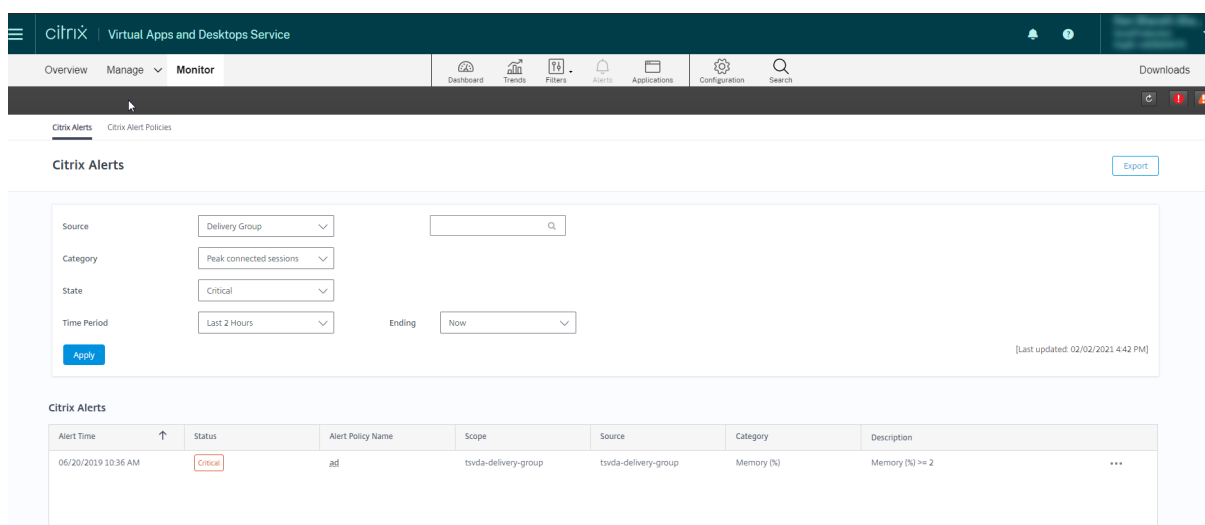


A warning alert (amber triangle) indicates that the warning threshold of a condition has been reached or exceeded.

A critical alert (red circle) shows that the critical threshold of a condition has been reached or exceeded.

You can view more detailed information on alerts by selecting an alert from the sidebar, clicking the **Go to Alerts** link at the bottom of the sidebar or by selecting **Alerts** from the top of the Monitor page.

In the Alerts view, you can filter and export alerts. For example, Failed Multi-session OS machines for a specific delivery group over the last month, or all alerts for a specific user. For more information, see [Export reports](#).



## Citrix alerts

Citrix alerts are the ones that originate from Citrix components. You can configure Citrix alerts within Monitor in **Alerts > Citrix Alerts Policy**. As part of the configuration, you can set notifications to be sent by email to individuals and groups when alerts exceed the thresholds you have set up. For more information on setting up Citrix Alerts, see [Create alerts policies](#).

## Smart alert policies

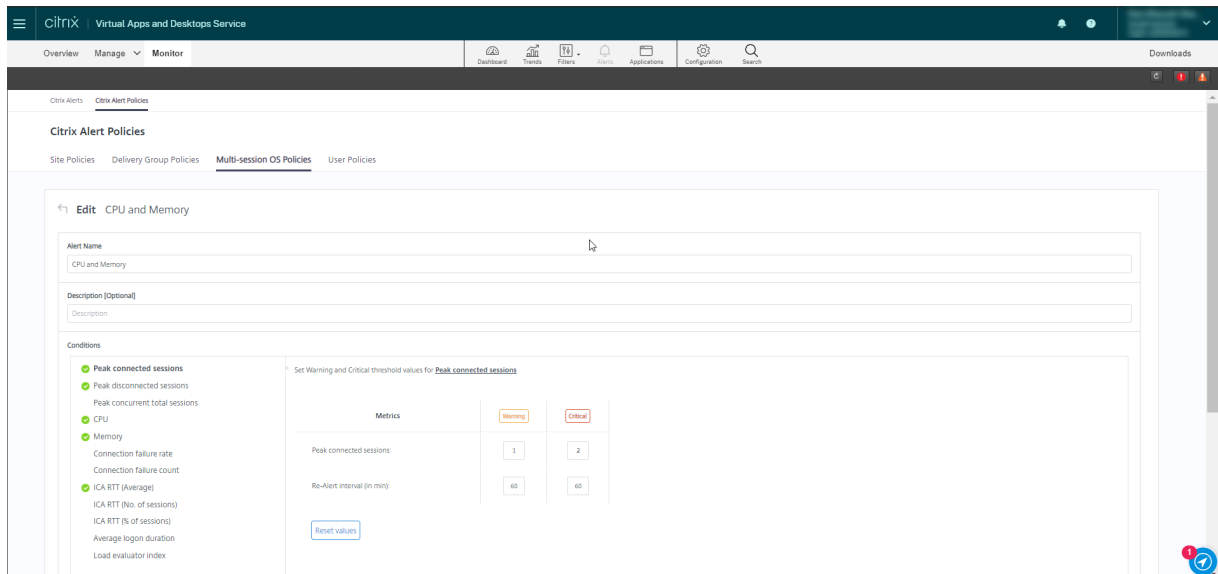
A set of built-in alert policies with predefined threshold values are available for delivery groups and Multi-session OS VDAs scope. You can modify the threshold parameters of the built-in alert policies in **Alerts > Citrix Alerts Policy**.

These policies are created when there is at least one alert target - a delivery group or a Multi-session OS VDA defined in your site. Additionally, these built-in alerts are automatically added to a new delivery group or a Multi-session OS VDA.

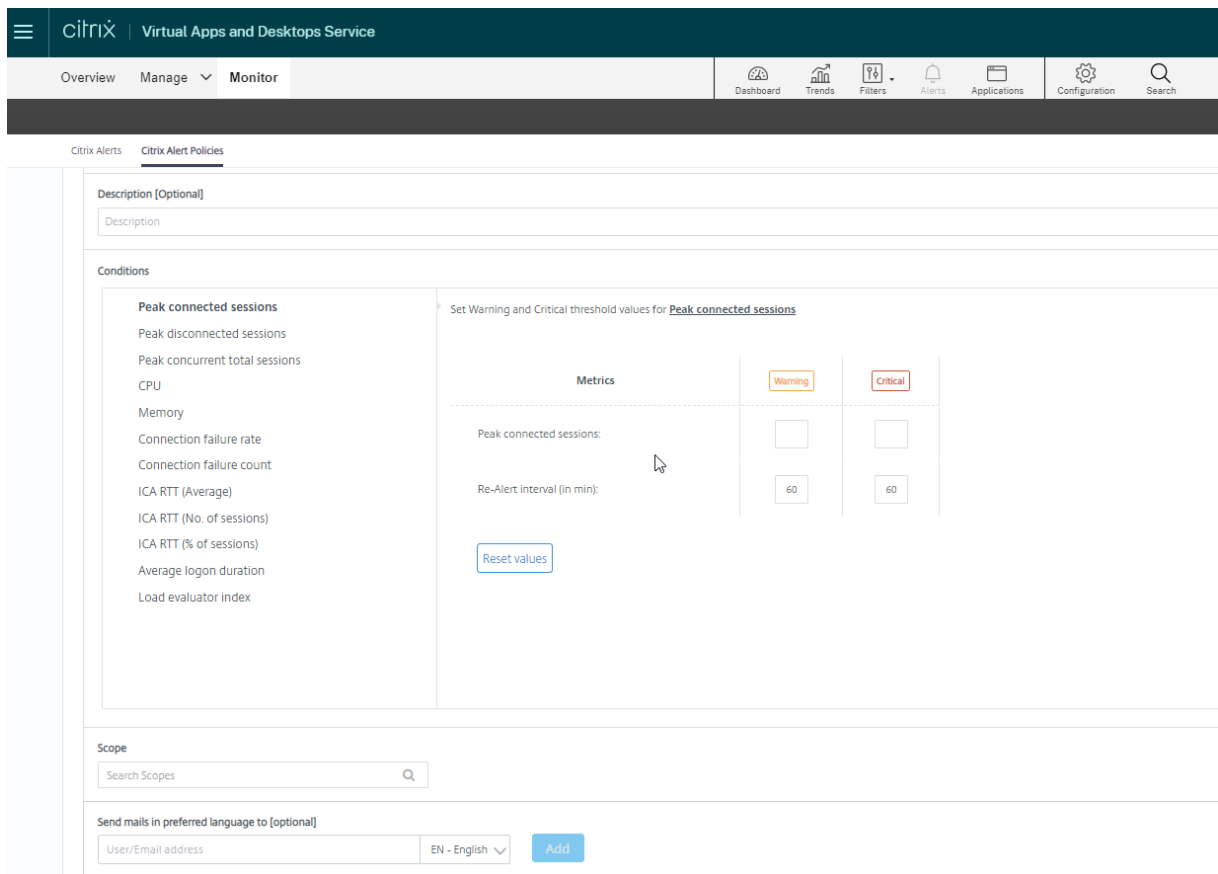
Built-in alert policies are created only if no corresponding alert rules exist in the Monitor database.

For the threshold values of the built-in alert policies, see the Alerts policies conditions section.

# Citrix Virtual Apps and Desktops service



## Create alerts policies



To create a new alerts policy, for example, to generate an alert when a specific set of session count criteria are met:



1. Go to **Alerts** > **Citrix Alerts Policy** and select, for example, Multi-session OS Policy.
2. Click **Create**.
3. Name and describe the policy, then set the conditions that have to be met for the alert to be triggered. For example, specify Warning and Critical counts for Peak Connected Sessions, Peak Disconnected Sessions, and Peak Concurrent Total Sessions. Warning values must not be greater than Critical values. For more information, see [Alerts policies conditions](#).
4. Set the Re-alert interval. If the conditions for the alert are still met, the alert is triggered again at this time interval and, if set up in the alert policy, an email notification is generated. A dismissed alert does not generate an email notification at the re-alert interval.
5. Set the Scope. For example, set for a specific delivery group.
6. In Notification preferences, specify who should be notified by email when the alert is triggered. Email notifications are sent via SendGrid. Ensure that the email address 'donotreplynotifications@citrix.com' is white-listed in your email setup.
7. Click **Save**.

Creating a policy with 20 or more delivery groups defined in the Scope might take approximately 30 seconds to complete the configuration. A spinner is displayed during this time.

Creating more than 50 policies for up to 20 unique delivery groups (1000 delivery group targets in total) might result in an increase in response time (over 5 seconds).

Moving a machine containing active sessions from one delivery group to another might trigger erroneous delivery group alerts that are defined using machine parameters.

## Alerts policies conditions

Find below the alert categories, recommended actions to mitigate the alert, and built-in policy conditions if defined. The built-in alert policies are defined for alert and realert intervals of 60 minutes.

### Peak Connected Sessions

- Check Monitor Session Trends view for peak connected sessions.
- Check to ensure that there is enough capacity to accommodate the session load.
- Add new machines if needed

### Peak Disconnected Sessions

- Check Monitor Session Trends view for peak disconnected sessions.
- Check to ensure that there is enough capacity to accommodate session load.
- Add new machines if needed.
- Log off disconnected sessions if needed

### Peak Concurrent Total Sessions

- Check Monitor Session Trends view in Monitor for peak concurrent sessions.
- Check to ensure that there is enough capacity to accommodate session load.
- Add new machines if needed.
- Log off disconnected sessions if needed

### CPU

Percentage of CPU usage indicates the overall CPU consumption on the VDA, including that of the processes. You can get more insight into the CPU utilization by individual processes from the **Machine details** page of the corresponding VDA.

- Go to **Machine Details > View Historical Utilization > Top 10 Processes**, identify the processes consuming CPU. Ensure that process monitoring policy is enabled to initiate collection of process level resource usage statistics.
- End the process if necessary.
- Ending the process causes unsaved data to be lost.
- If all is working as expected, add additional CPU resources in the future.

**Note:**

The policy setting, **Enable resource monitoring** is allowed by default for the monitoring of CPU and memory performance counters on machines with VDAs. If this policy setting is disabled, alerts with CPU and memory conditions are not triggered. For more information, see [Monitoring policy settings](#).

**Smart policy conditions:**

- **Scope:** Delivery group, Multi-session OS scope
- **Threshold values:** Warning - 80%, Critical - 90%

### Memory

Percentage of Memory usage indicates the overall memory consumption on the VDA, including that of the processes. You can get more insight into the memory usage by individual processes from the **Machine details** page of the corresponding VDA.

- Go to **Machine Details > View Historical Utilization > Top 10 Processes**, identify the processes consuming memory. Ensure that process monitoring policy is enabled to initiate collection of process level resource usage statistics.
- End the process if necessary.

- Ending the process causes unsaved data to be lost.
- If all is working as expected, add additional memory in the future.

**Note:**

The policy setting, **Enable resource monitoring**, is allowed by default for the monitoring of CPU and memory performance counters on machines with VDAs. If this policy setting is disabled, alerts with CPU and memory conditions are not triggered. For more information, see [Monitoring policy settings](#).

**Smart policy conditions:**

- **Scope:** Delivery group, Multi-session OS scope
- **Threshold values:** Warning - 80%, Critical - 90%

### Connection Failure Rate

Percentage of connection failures over the last hour.

- Calculated based on the total failures to total connections attempted.
- Check Monitor Connection Failures Trends view for events logged from the Configuration log.
- Determine if applications or desktops are reachable.

### Connection Failure Count

Number of connection failures over the last hour.

- Check Monitor Connection Failures Trends view for events logged from the Configuration log.
- Determine if applications or desktops are reachable.

### ICA RTT (Average)

Average ICA round-trip time.

- Check Citrix ADM for a breakdown of the ICA RTT to determine the root cause. For more information, see [Citrix ADM](#) documentation.
- If Citrix ADM is not available, check the Monitor User Details view for the ICA RTT and Latency, and determine if it is a network problem or an issue with applications or desktops.

### ICA RTT (No. of Sessions)

Number of sessions that exceed the threshold ICA round-trip time.

- Check Citrix ADM for the number of sessions with high ICA RTT. For more information, see [Citrix ADM](#) documentation.

- If Citrix ADM is not available, contact the network team to determine the root cause.

**Smart policy conditions:**

- **Scope:** Delivery group, Multi-session OS scope
- **Threshold values:** Warning - 300 ms for 5 or more sessions, Critical - 400 ms for 10 or more sessions

**ICA RTT (% of Sessions)**

Percentage of sessions that exceed the average ICA round-trip time.

- Check Citrix ADM for the number of sessions with high ICA RTT. For more information, see [Citrix ADM](#) documentation.
- If Citrix ADM is not available, contact the network team to determine the root cause.

**ICA RTT (User)**

ICA round-trip time that is applied to sessions launched by the specified user. The alert is triggered if ICA RTT is greater than the threshold in at least one session.

**Failed Machines (Single session OS)**

Number of failed Single session OS machines. Failures can occur for various reasons as shown in the Monitor Dashboard and Filters views.

- Run Citrix Scout diagnostics to determine the root cause. For more information, see [Troubleshoot user issues](#).

**Smart policy conditions:**

- **Scope:** Delivery group scope
- **Threshold values:** Warning - 1, Critical - 2

**Failed Machines (Multi-session OS)**

Number of failed Multi-session OS machines. Failures can occur for various reasons as shown in the Monitor Dashboard and Filters views.

- Run Citrix Scout diagnostics to determine the root cause.

**Smart policy conditions:**

- **Scope:** Delivery group, Multi-session OS scope
- **Threshold values:** Warning - 1, Critical - 2

### Average Logon Duration

Average logon duration for logons that occurred over the last hour.

- Check the Monitor Dashboard to get up-to-date metrics regarding the logon duration. A large number of users logging in during a short timeframe can increase the logon duration.
- Check the baseline and break down of the logons to narrow down the cause. For more information, see [Diagnose user logon issues](#).

#### Smart policy conditions:

- **Scope:** Delivery group, Multi-session OS scope
- **Threshold values:** Warning - 45 seconds, Critical - 60 seconds

### Logon Duration (User)

Logon duration for logons for the specified user that occurred over the last hour.

### Load Evaluator Index

Value of the Load Evaluator Index over the last 5 minutes.

- Check Monitor for Multi-session OS Machines that might have a peak load (Max load). View both Dashboard (failures) and Trends Load Evaluator Index report.

#### Smart policy conditions:

- **Scope:** Delivery group, Multi-session OS scope
- **Threshold values:** Warning - 80%, Critical - 90%

### Hypervisor Alerts Monitoring

Monitor displays alerts to monitor hypervisor health. Alerts from Citrix Hypervisor and VMware vSphere help monitor hypervisor parameters and states. The connection status to the hypervisor is also monitored to provide an alert if the cluster or pool of hosts is rebooted or unavailable.

To receive hypervisor alerts, ensure that a hosting connection is created in the Manage tab. For more information, see [Connections and resources](#). Only these connections are monitored for hypervisor alerts. The following table describes the various parameters and states of Hypervisor alerts.

Alert	Supported Hypervisors	Triggered by	Condition	Configuration
CPU usage	Citrix Hypervisor, VMware vSphere	Hypervisor	CPU usage alert threshold is reached or exceeded	Alert thresholds must be configured in the Hypervisor.
Memory usage	Citrix Hypervisor, VMware vSphere	Hypervisor	Memory usage alert threshold is reached or exceeded	Alert thresholds must be configured in the Hypervisor.
Network usage	Citrix Hypervisor, VMware vSphere	Hypervisor	Network usage alert threshold is reached or exceeded	Alert thresholds must be configured in the Hypervisor.
Disk usage	VMware vSphere	Hypervisor	Disk usage alert threshold is reached or exceeded	Alert thresholds must be configured in the Hypervisor.
Host connection or power state	VMware vSphere	Hypervisor	Hypervisor Host has been rebooted or is unavailable	Alerts are prebuilt in VMware vSphere. No additional configurations are needed.
Hypervisor connection unavailable	Citrix Hypervisor, VMware vSphere	Delivery Controller	Connection to the hypervisor (pool or cluster) is lost or powered down or rebooted. This alert is generated every hour as long as the connection is unavailable.	Alerts are prebuilt with the Delivery Controller. No additional configurations are needed.

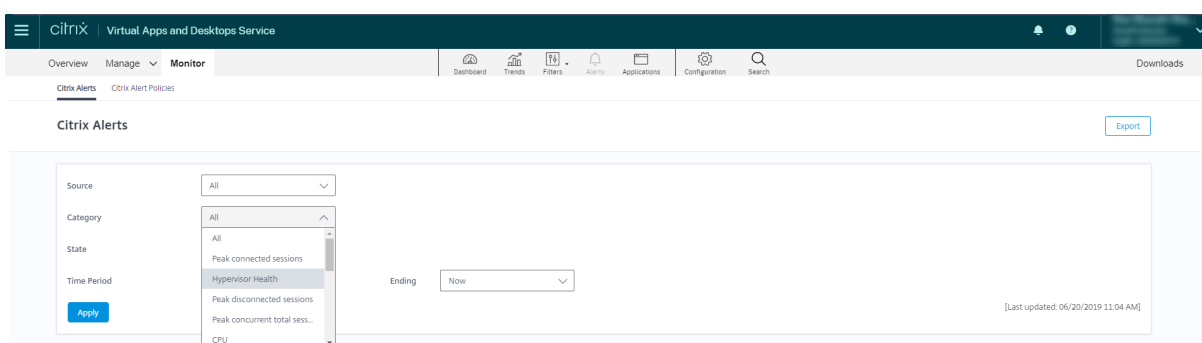
**Note:**

For more information about configuring alerts, see [Citrix XenCenter Alerts](#) or [VMware vCenter Alerts](#).

Email notification preference can be configured under **Citrix Alerts Policy > Site Policy > Hypervisor Health**. The threshold conditions for Hypervisor alert policies can be configured, edited, disabled, or deleted from the hypervisor only and not from Monitor. However, modifying email preferences and dismissing an alert can be done in Monitor.

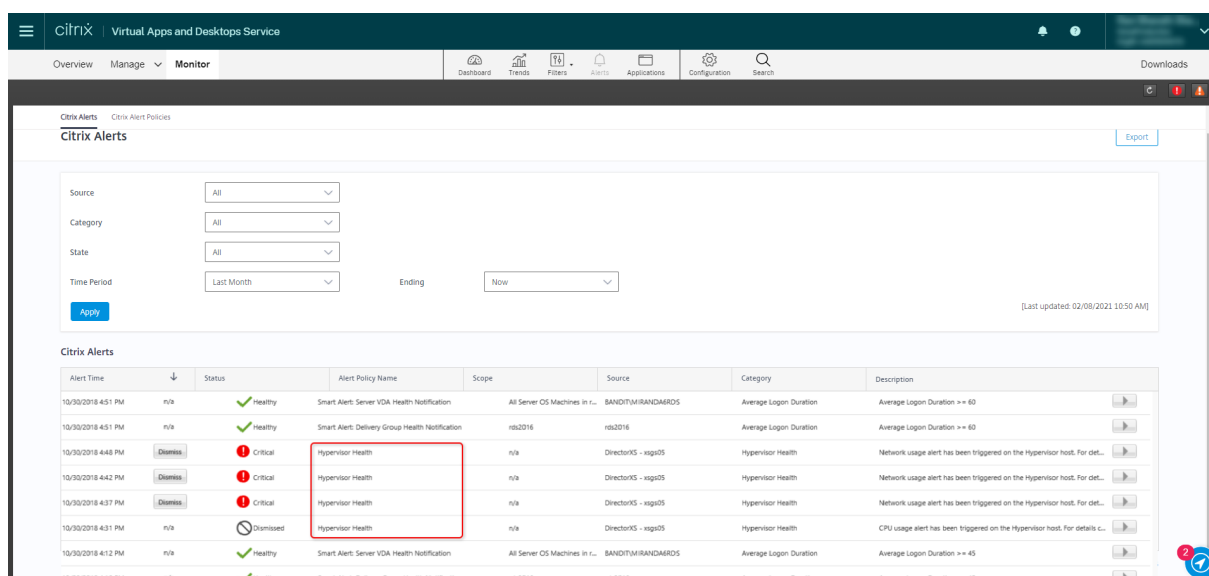
**Important:**

- All hypervisor alerts older than a day are automatically dismissed.
- Alerts triggered by the Hypervisor are fetched and displayed in Monitor. However, changes in the life cycle/state of the Hypervisor alerts are not reflected in Monitor.
- Alerts that are healthy or dismissed or disabled in the Hypervisor console will continue to appear in Monitor and have to be dismissed explicitly.
- Alerts that are dismissed in Monitor are not dismissed automatically in the Hypervisor console.



A new Alert category called **Hypervisor Health** has been added to enable filtering only the hypervisor alerts. These alerts are displayed once the thresholds are reached or exceeded. Hypervisor alerts can be:

- **Critical**—critical threshold of the hypervisor alarm policy reached or exceeded
- **Warning**—warning threshold of the hypervisor alarm policy reached or exceeded
- **Dismissed**—alert no longer displayed as an active alert



## Filter data to troubleshoot failures

September 14, 2020

When you click numbers on the Dashboard or select a predefined filter from the **Filters** menu, the Filters view opens to display data based on the selected machine or failure type.

Predefined filters cannot be edited, but you can save a predefined filter as a custom filter and then modify it. Also, you can create custom filtered views of machines, connections, sessions, and application instances across all delivery groups.

1. Select a view:

- **Machines.** Select Single session OS Machines or Multi-session OS Machines. These views show the number of configured machines. The Multi-session OS Machines tab also includes the load evaluator index, which indicates the distribution of performance counters and tool tips of the session count if you hover over the link.
- **Sessions.** You can also see the session count from the Sessions view. Use the idle time measurements to identify sessions that are idle beyond a threshold time period.
- **Connections.** Filter connections by different time periods, including last 60 minutes, last 24 hours, or last 7 days.
- **Application Instances.** This view displays the properties of all application instances on VDAs of Multi-session and Single session OS. The session idle time measurements are available for Application instances on VDAs of Multi-session OS.

2. For **Filter by**, select the criteria.

3. Use the additional tabs for each view, as needed, to complete the filter.



4. Select extra columns, as needed, to troubleshoot further.
5. Save and name your filter.
6. To open the filter later, from the **Filters** menu, select the filter type (Machines, Sessions, Connections, or Application Instances), and then select the saved filter.
7. Click **Export** to export the data to CSV format files. Data of up to 100,000 records can be exported.
8. If needed, for **Machines** or **Connections** views, use power controls for all the machines you select in the filtered list. For the Sessions view, use the session controls or option to send messages.
9. In the **Machines** and **Connections** views, click the **Failure Reason** of a failed machine or connection to get a detailed description of the failure and actions recommended to troubleshoot the failure. The failure reasons and the recommended actions for Machine and Connection failures are available in the [Citrix Director Failure Reasons Troubleshooting Guide](#).
10. In the **Machines** view, click a machine name link to go to the corresponding **Machine Details** page. This page displays the details of the machine, provides power controls, displays the CPU, memory, disk monitoring, and GPU monitoring graphs. Also, click **View Historical Utilization** to see the resource utilization trends for the machine. For more information, see [Troubleshoot machines](#).
11. In the **Application Instances** view, sort or filter based on **Idle Time** greater than a threshold time period. Select the idle application instances to end. Log off or Disconnect of an application instance ends all active application instances in the same session. For more information, see [Troubleshoot applications](#). The Application Instances filter page and idle time measurements in the Sessions filter pages are available if VDAs are version 7.13 or later.

**Note:**

Citrix Studio allows assignment of multiple Desktop Assignment Rules (DAR) for different users or user groups to a single VDA in the delivery group. StoreFront displays the assigned desktop with the corresponding Display Name as per the DAR for the logged in user. However, Monitor does not support DARs and displays the assigned desktop using the delivery group name regardless of the logged in user. As a result, you cannot map a specific desktop to a machine in Monitor. To map the assigned desktop displayed in StoreFront to the delivery group name displayed in Monitor, use the following PowerShell command. Run the PowerShell command using Remote PowerShell SDK as described in the [blog](#).

```
1 Get-BrokerDesktopGroup | Where-Object {  
2   $_.Uid -eq (Get-BrokerAssignmentPolicyRule | Where-Object {  
3     $_.PublishedName -eq "<Name on StoreFront>" }  
4   }).DesktopGroupId }
```

```
5 | Select-Object -Property Name, Uid
6 <!--NeedCopy-->
```

## Monitor historical trends across a site

March 29, 2021

The Trends view accesses historical trend information of each site for the following parameters:

- sessions
- connection failures
- machine failures
- logon performance
- load evaluation
- capacity management
- machine usage
- resource utilization

To locate this information, click the **Trends** menu.

The zoom-in drill down feature lets you navigate through trend charts by zooming in on a time period (clicking a data point in the graph) and drilling down to see the details associated with the trend. This feature enables you to better understand the details of who or what has been affected by the trends being displayed.

To change the default scope of each graph, apply a different filter to the data.

### Note:

- Sessions, failures, and logon performance trend information are available as graphs and tables when the time period is set to Last month (**Ending now**) or shorter. When the time period is chosen as Last month with a custom ending date or as Last year, the trend information is available as graphs but not as tables.
- Citrix Virtual Apps and Desktops service supports historical data retention only for 90 days. Hence, one-year trends and reports in Monitor show the last 90 days of data.

## Available trends

**View trends for sessions:** From the Sessions tab, select the delivery group and time period to view more detailed information about the concurrent session count.

The **Session Auto Reconnect** column displays the number of auto reconnects in a session. Auto reconnect is enabled when the Session Reliability or the Auto Client Reconnect policies are in effect.

When there is a network interruption on the endpoint, the following policies come into effect:

- Session reliability comes into effect (by default for 3 minutes) where the Citrix Receiver or Citrix Workspace app tries to connect to the VDA.
- Auto client reconnect comes into effect between 3 and 5 minutes where the client tries to connect to the VDA.

Both these reconnects are captured and displayed to the user. This information can take a maximum time of 5 minutes to appear on the Director UI after the reconnect occurs.

The auto reconnect information helps you view and troubleshoot network connections having interruptions, and to analyze networks having a seamless experience. You can view the number of reconnects for a specific delivery group or time period selected in the Filters.

A drilldown provides additional information like Session Reliability or Auto Client Reconnect, time stamps, Endpoint IP, and Endpoint Name of the machine where Workspace app is installed.

By default, logs are sorted by the event time stamps in descending order. This feature is available for Citrix Workspace app for Windows, Citrix Workspace app for Mac, Citrix Receiver for Windows, and Citrix Receiver for Mac. This feature requires VDAs 1906 or later.

For more information about session reconnections, see [Sessions](#). For more information about policies, see [Auto client reconnect policy settings](#) and [Session reliability policy settings](#).

Sometimes, the auto reconnect data might not appear in Monitor for the following reasons:

- Workspace app is not sending auto reconnect data to VDA.
- VDA is not sending data to monitor service.

**Note:**

Sometimes, the client IP address might not be obtained correctly if certain Citrix Gateway policies are set.

**View trends for connection failures:** From the Failures tab, select the connection, machine type, failure type, delivery group, and time period to view a graph containing more detailed information about the user connection failures across your site.

**View trends for machine failures:** From the Single session OS Machine Failures tab or Multi-session OS Machines tab, select the failure type, delivery group, and time period to view a graph containing more detailed information about the machine failures across your site.

**View trends for logon performance:** From the Logon Performance tab, select the delivery group and time period to view a graph containing more detailed information about the duration of user logon times across your site and whether the number of logons affects the performance. This view also shows the average duration of the logon phases, such as brokering duration and VM start time. This data is specifically for user logons and does not include users trying to reconnect from disconnected sessions.

The table below the graph shows Logon Duration by User Session. You can choose the columns to display and sort the report by any of the columns.

For more information, see [Diagnose user logon issues](#).

**View trends for load evaluation:** From the Load Evaluator Index tab, view a graph containing more detailed information about the load that is distributed among Multi-session OS machines. The filter options for this graph include the delivery group or Multi-session OS machine in a delivery group, Multi-session OS machine (available only if Multi-session OS machine in a delivery group was selected), and range. The Load Evaluator Index is displayed as percentages of Total CPU, Memory, Disk or Sessions and is shown in comparison with the number of connected users in the last interval.

**View hosted applications usage:** From the Capacity Management tab, select the Hosted Applications Usage tab, select the delivery group and time period to view a graph displaying peak concurrent usage and a table displaying application based usage. From the Application Based Usage table, you can choose a specific application to see details and a list of users who are using, or have used, the application. You can see the predicted peak concurrent application instances values chosen future time period with Application instance prediction. For more information, see the [Application instance prediction section](#).

**View single and multi-session OS usage:** The Trends view shows the usage of Single session OS by site and by delivery group. When you select site, usage is shown per delivery group. When you select delivery group, usage is shown per User.

The Trends view also shows the usage of Multi-session OS by site, by delivery group, and by Machine. When you select site, usage is shown per delivery group. When you select delivery group, usage is shown per Machine and per User. When Machine is selected usage is shown per User.

**View virtual machine usage:** From the Machine Usage tab, select Single session OS Machines or Multi-session OS Machines to obtain a real-time view of your VM usage. The page displays the number of Autoscale enabled Multi-session and Single session OS machines that are powered on for a selected delivery group and time period. Also available is the estimated savings achieved by enabling Autoscale in the selected delivery group, this percentage is calculated using the per machine costs.

The usage trends of Autoscale enabled machines indicate the actual usage of the machines, enabling you to quickly assess your site's capacity needs.

- Single session OS availability - displays the current state of Single session OS machines (VDIs) by availability for the entire site or a specific delivery group.
- Multi-session OS availability - displays the current state of Multi-session OS machines by availability for the entire site or a specific delivery group.

**Note:**

The grid below the chart displays the delivery group based machine usage data in real-time. The data includes machine availability of all machines independent of Autoscale enablement. The

number of machines displayed in the Available Counter column in the grid includes machines in maintenance mode.

The monitoring data consolidation depends on the time period you select.

- Monitoring data for the one day and one week time periods is consolidated per hour.
- Monitoring data for the one month time period is consolidated per day.

The machine status is read at the time of consolidation and any changes during the period in between is not considered. For the consolidation period, refer to the [Monitor API documentation](#).

For more information on monitoring Autoscale enabled machines see the [Autoscale](#) article.

**View resource utilization:** From the Resource Utilization tab, select Single session OS Machines or Multi-session OS Machines to obtain insight into historical trends data for CPU and memory usage, and IOPS and disk latency for each VDI machine for better capacity planning.

This feature requires VDAs **version 7.11** or later.

Graphs show data for average CPU, average memory, average IOPS, disk latency, and peak concurrent sessions. You can drill down to the machine, and view data and charts for the top 10 processes consuming CPU. Filter by delivery group and Time period. CPU, memory usage, and peak concurrent sessions graphs are available for the last 2 hours, 24 hours, 7 days, month, and year. The average IOPS and disk latency graphs are available for the last 24 hours, month, and year.

**Note:**

- The Monitoring policy setting, [Enable Process Monitoring](#), must be set to "Allowed" to collect and display data in the Top 10 Processes table on the Historic Machine Utilization page. The policy is set to "Prohibited" by default. All resource utilization data is collected by default. This can be disabled using the [Enable Resource Monitoring](#) policy setting. The table below the graphs shows the resource utilization data per machine.
- Average IOPS shows the daily averages. Peak IOPS is calculated as the highest of the IOPS averages for the selected time range. (An IOPS average is the hourly average of IOPS collected during the hour on the VDA).

**View application failures:** The Application Failures tab displays failures associated with the published applications on the VDAs.

This feature requires VDAs **version 7.15** or later. Single session OS VDAs running Windows Vista and later, and Multi-session OS VDAs running Windows Server 2008 and later are supported.

For more information, see [Historical application failure monitoring](#).

By default, only application faults from Multi-session OS VDAs are displayed. You can set the monitoring of application failures by using Monitoring policies. For more information, see [Monitoring policy settings](#).

**View application probe results:** The Application Probe Results tab displays the results of probe for applications that have been configured for probing in the Configuration page. Here, the stage of

launch during which the application launch failure occurred is recorded.

This feature requires VDAs **version 7.18** or later. For more information see [Application probing](#).

**Create customized reports:** The Custom Reports tab provides a user interface for generating Custom Reports containing real-time and historical data from the Monitoring database in tabular format.

From the list of previously saved Custom Report queries, you can click **Run and download** to export the report in CSV format, click **Copy OData** to copy and share the corresponding OData query, or click **Edit** to edit the query.

You can create a Custom Report query based on machines, connections, sessions, or application instances. Specify filter conditions based on fields such as machine, delivery group, or time period. Specify extra columns required in your Custom Report. Preview displays a sample of the report data. Saving the Custom Report query adds it to the list of saved queries.

You can create a Custom Report query based on a copied OData query. To do this, select the OData Query option and paste the copied OData query. You can save the resultant query for execution later.

**Note:**

The column names in Preview and Export report generated using OData queries are not localized, but appear in English.

The flag icons on the graph indicate significant events or actions for that specific time range. Hover the mouse over the flag and click to list events or actions.

**Note:**

- HDX connection logon data is not collected for VDAs earlier than 7. For earlier VDAs, the chart data is displayed as 0.
- Delivery groups deleted in Citrix Studio are available for selection in the Trends filters until data related to them are groomed out. Selecting a deleted delivery group displays graphs for available data until retention. However, the tables don't show data.
- Moving a machine containing active sessions from one delivery group to another causes the **Resource Utilization and Load Evaluator Index** tables of the new delivery group to display metrics consolidated from the old and new delivery groups.

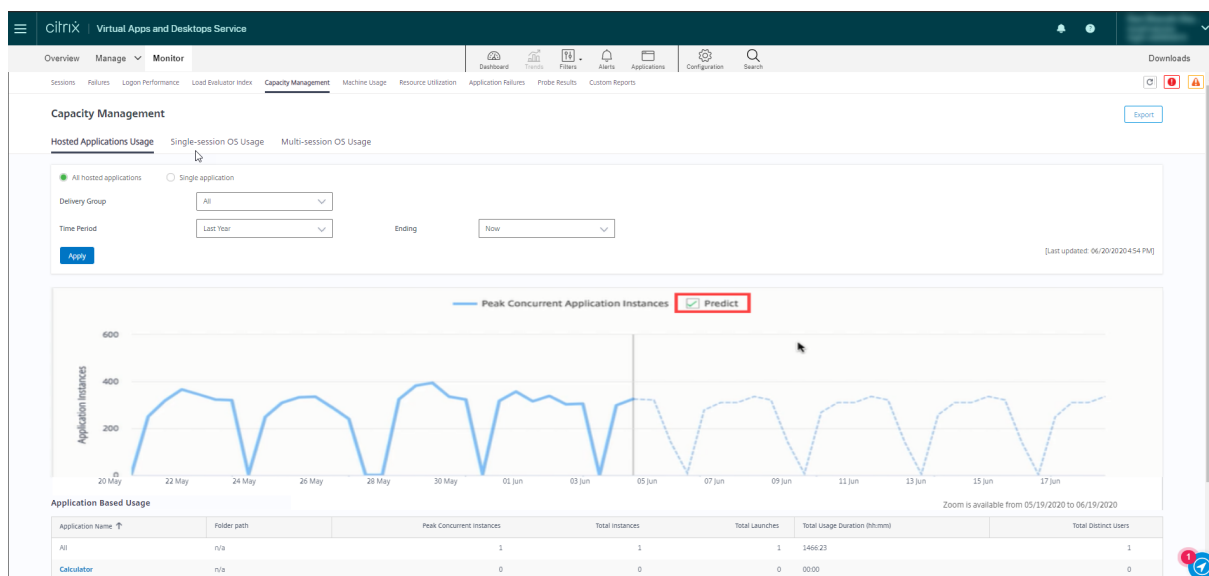
## Application instance prediction

Predictive analytics gives you the ability to predict future resource usage. This feature is especially useful for administrators to organize required resources and licenses on each resource.

The first predictive analysis feature, Application instance prediction predicts the number of hosted application instances likely to be launched per site or delivery group over time.

Application instance prediction is available in the **Trends > Capacity Management** tab that displays

the hosted application usage for the chosen time period. The historical graph contains the peak concurrent application instances values plotted for the chosen period.



To get the predicted graph, select the Predict check box. A dotted line prediction graph is displayed as an extension of the historical graph. The predicted peak concurrent application instances values are plotted with the time line extended into the future for the chosen time period.

You can predict the application instances for time periods of the next 7 days, 1 month, or one year. Custom ending dates are not supported.

Prediction is done using machine learning algorithms that are based on data models created with existing historical data. The predictions are therefore as accurate as the quality of the existing data.

The accuracy of prediction is indicated by the tolerance level that is displayed as a tool tip over the predicted graph. It indicates the amount of possible variation of the actual values from the predicted values.

The tolerance level can be high if either the available data does not follow a regular pattern or is missing for certain periods or is insufficient.

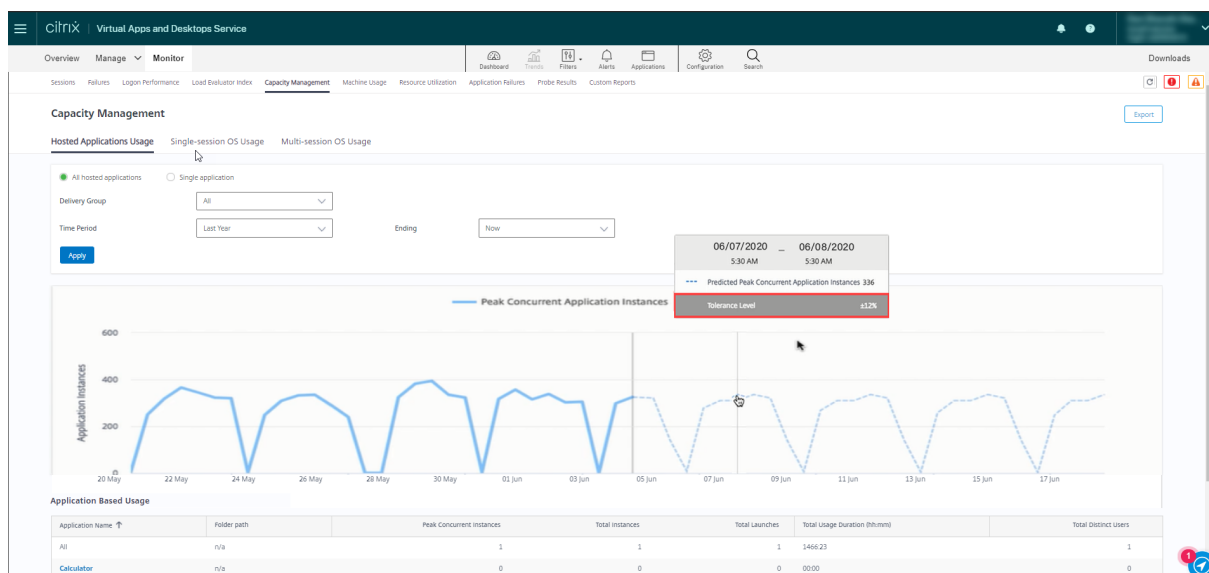
Prediction for a year captures the monthly and quarterly patterns coupled with the overall trend for the year. Similarly, monthly prediction captures the daily and weekly patterns along with weekly trends such as reduced activity over weekends.

Sufficient historical data must be available for prediction as follows:

- 14 days data for 7 days' prediction
- 35 days data for one month's prediction
- 84 days data for one year's prediction

**Note:**

You can export only the historical graph, but not the predicted graph.



## Monitor Autoscale-managed machines

August 10, 2020

Autoscale is a power management feature that enables proactive power management of all registered Multi-session and Single session OS machines in a delivery group. You can configure Autoscale for a selected delivery group from the **Manage** tab. For more information, see [Autoscale](#).

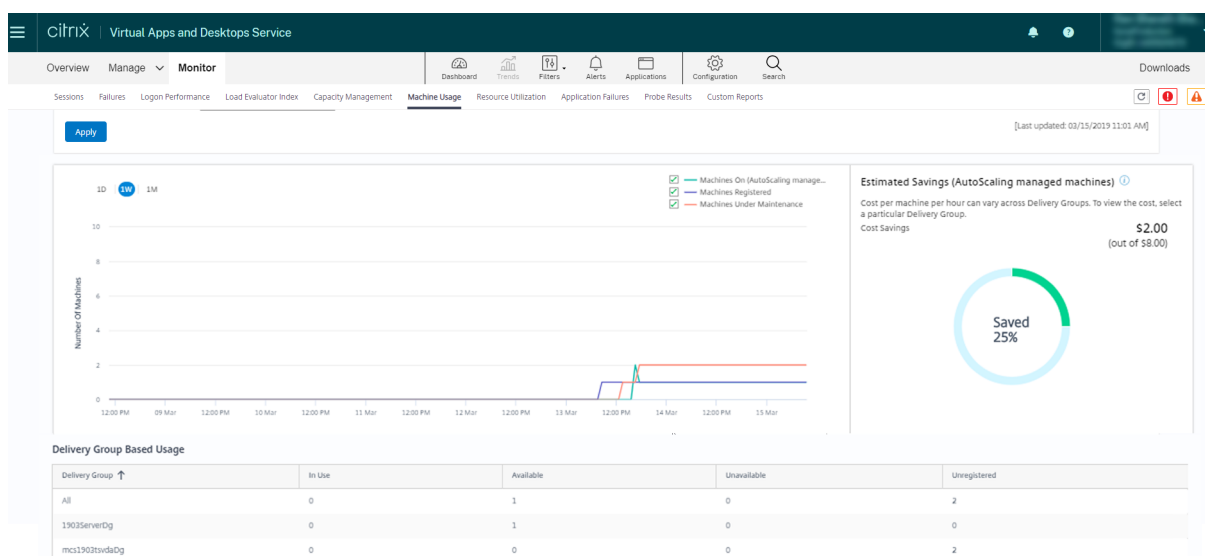
You can monitor the key metrics of Autoscale enabled machines from the **Monitor** tab.

### Machine Usage

The **Monitor > Trends > Machine Usage** page displays the total number of Autoscale enabled Multi-session and Single session OS machines that are powered on for a selected delivery group and time period. This metric indicates the actual usage of machines in the delivery group.

From the **Single session OS Machines** or the **Multi-session OS Machines** tab, select the Delivery group and the time period.






The chart plots the following metrics:

- **Machines On** - the number of Autoscale enabled machines that are powered on
- **Machines Registered** - the number of registered Multi-session or Single session OS machines
- **Machines under Maintenance** - the number of Multi-session or Single session OS machines with maintenance mode switched on

### Estimated Savings

The **Monitor > Trends > Machine Usage** page also displays the estimated cost savings achieved by enabling Autoscale in the selected delivery group.



Estimated Savings is calculated as the percentage of savings per machine per hour (in US \$) as con-

figured in **Manage > Edit Delivery Group > Autoscale**. For more information about configuring the savings per machine, see [Autoscale](#).

When you select all Delivery groups, the average value of Estimated Savings across all the delivery groups is displayed.

The estimated savings help administrators consolidate the existing infrastructure and plan the capacity to achieve maximum savings and utilization.

## Alert notifications for machines and sessions

The Monitor Dashboard displays alert notifications that can be further drilled down. Alert details are displayed on the **Monitor > Alerts** page.

- To create an alert policy in a delivery group, go to **Monitor > Alerts > Citrix Alerts Policy > Delivery Group Policy**.
- Here, you can set the following Warning and Critical thresholds:
  - Failed Machines (Single session OS) and Failed Machines (Multi-session OS),
  - Peak Connected Sessions, Peak Disconnected Sessions and Peak Concurrent Total Sessions in the delivery group.
- Alerts are generated when the corresponding metric in the delivery group reaches the threshold.

For more details regarding the alert policy conditions and creation of new alert policies, see [Alerts and notifications](#).

## Machine status

- **Monitor > Filters > Machines** displays the power state of all machines in a tabular format. You can filter by a specific delivery group.
- **Monitor > Filters > Sessions** displays filter by the Machine name to see the associated sessions and their real-time status.
- In **Monitor > Trends > Sessions**, select your delivery group and time period to see the trend of the sessions and their associated metrics.

For more information, see [Filter data to troubleshoot failures](#).

## Load Evaluation trends

The **Monitor > Trends > Load Evaluator Index** page displays a graph with detailed information about the load that is distributed among the Multi-session OS machines. The filter options for this graph include the delivery group or Multi-session OS machine in a delivery group, Multi-session OS machine (available only if Multi-session OS machine in a delivery group was selected), and range. The Load Evaluator Index is displayed as percentages of Total CPU, Memory, Disk, or Sessions and is shown in comparison with the number of connected users in the last interval.

## Troubleshoot deployments

June 22, 2020

As a help desk administrator, you can search for the user reporting an issue and display details of sessions or applications associated with that user.

Similarly, you can search for machines or endpoints where issues are reported. Issues can be quickly resolved by monitoring the relevant metrics and performing suitable actions.

The following actions are available:

- ending an unresponsive application or process
- shadowing operations on the user's machine
- logging off an unresponsive session
- restarting the machine
- putting a machine into maintenance mode
- resetting the user profile

## Troubleshoot applications

September 14, 2020

### Application Analytics

The **Applications** view displays application-based analytics in a single, consolidated view to help analyze and manage application performance efficiently. You can gain valuable insight into the health and usage information of all applications published on the site. The default view helps identify the top running applications.

This feature requires VDAs Version 7.15 or later.



## Real-time application monitoring

You can troubleshoot applications and sessions by using the idle time metric to identify instances that are idle beyond a specific time limit.

Typical use cases for application-based troubleshooting are in the healthcare sector, where employees share application licenses. There, you must end idle sessions and application instances to purge the Citrix Virtual Apps and Desktops environment, to reconfigure poorly performing servers, or to maintain and upgrade applications.

The **Application Instances** filter page lists all application instances on VDAs of Multi-session and Single session OS. The associated idle time measurements are displayed for application instances on VDAs of Multi-session OS that have been idle for at least 10 minutes.

### Note:

The Application Instances metrics are available on sites of all license editions.

Use this information to identify the application instances that are idle beyond a specific time period and log off or disconnect them as appropriate. To do this, select **Filters > Application Instances** and select a pre-saved filter or choose **All Application Instances** and create your own filter.

The screenshot shows the 'Filters - All Application Instances' page in the Citrix Virtual Apps and Desktops service. The page has a navigation bar with 'Overview', 'Manage', and 'Monitor' tabs. Below the navigation bar, there are icons for 'Dashboard', 'Trends', 'Filters', 'Alerts', 'Applications', 'Configuration', and 'Search'. The 'Filters' icon is selected. The page title is 'Filters - All Application Instances\*'. Below the title, there is a 'View:' section with radio buttons for 'Machines', 'Sessions', 'Connections', and 'Application Instances'. The 'Application Instances' radio button is selected. Below the 'View:' section, there is a 'Filter by:' section with two dropdown menus. The first dropdown menu is 'Published Name' and the second is 'Associated User'. The first dropdown menu is set to 'contains' and the second is set to 'contains'. The first dropdown menu has 'US' selected and the second has 'a' selected. Below the 'Filter by:' section, there are buttons for 'Save', 'Save As...', 'Delete', and 'Clear'. Below the filter section, there is a table titled '10 Application Sessions'. The table has columns for 'Published Name', 'Login Time', 'Idle Time (hh:mm)', 'Associated User', 'Anonymous', 'Machine Name', 'IP Address', 'Endpoint Name', and 'Endpoint IP'. The table contains 10 rows of data. The 'Idle Time' column shows 'n/a' for all entries. The 'Associated User' column shows various usernames like 'james@us...' and 'james@us...'. The 'Machine Name' column shows 'XENDESKTOP-us-7616-' and 'XENDESKTOPUS19020'. The 'IP Address' column shows '192.168.200...' and '172.16.128.200'. The 'Endpoint Name' column shows 'XENDESKTOPUS19020' and 'XENDESKTOPUS19020'. The 'Endpoint IP' column shows '192.168.200...' and '192.168.200...'. The table has a 'Session Control' dropdown menu and a 'Send Message' button. The 'Session Control' dropdown menu is set to 'Logoff'. The 'Send Message' button is disabled. The table has an 'Export' button and a 'Choose Columns' button. The page footer shows 'Displaying 1 - 8 of 10' and a help icon.

An example of a filter would be as follows. As **Filter by** criteria, choose **Published Name** (of the application) and **Idle Time**. Then, set **Idle Time** to **greater than or equal to** a specific time limit and save the filter for reuse. From the filtered list, select the application instances. Select option to send messages or from the **Session Control** drop-down, choose **Logoff** or **Disconnect** to end the instances.

### Note:

Logging off or disconnecting an application instance logs off or disconnects the current session, thereby ending all application instances that belong to the same session.

You can identify idle sessions from the **Sessions** filter page using the session state and the session idle

time metric. Sort by the **Idle Time** column or define a filter to identify sessions that are idle beyond a specific time limit. Idle time is listed for sessions on VDAs of Multi-session OS that have been idle for at least 10 minutes.

The screenshot shows the Citrix Virtual Apps and Desktops service interface. The top navigation bar includes 'Overview', 'Manage', and 'Monitor'. The 'Monitor' tab is active, showing a 'Filters - All Sessions\*' section with options to view 'Machines', 'Sessions', 'Connections', and 'Application Instances'. Below this is a 'Filter by:' section with a dropdown menu and buttons for 'Save', 'Save As...', 'Delete', and 'Clear'. The main content area displays a table titled '6 Sessions' with the following data:

Associated User	Session State	Session Start Time	Anonymous	Endpoint Name	Endpoint IP	Citrix Workspace App...	Machine Name	IP Address	Idle Time (hh:mm)
Administrator	Disconnected	2/12/2019 12:07 PM	No	[redacted]	107.255.1	n/a	[redacted]	10.40.0.2	242:38
Administrator	Disconnected	2/4/2019 10:43 AM	No	[redacted]	10.182.0.12	18.12.0.12	[redacted]	10.40.0.2	268:51
Administrator	Disconnected	2/4/2019 10:48 AM	No	[redacted]	10.182.0.12	n/a	[redacted]	10.40.0.2	432:06
administrator0	Disconnected	2/6/2019 2:27 PM	No	[redacted]	10.182.0.12	18.12.0.12	[redacted]	10.40.0.2	n/a
Administrator	Active	2/4/2019 4:16 PM	No	[redacted]	107.255.1	n/a	[redacted]	10.40.0.2	n/a
administrator0	Active	2/22/2019 2:20 PM	No	[redacted]	10.182.0.12	18.12.0.12	[redacted]	10.40.0.2	n/a

The **Idle time** is displayed as **N/A** when the session or application instance

- has not been idle for more than 10 minutes,
- is launched on a VDA of Single session OS, or
- is launched on a VDA running Version 7.12 or earlier.

## Historical application failure monitoring

The **Trends -> Application Failures** tab displays failures associated with the published applications on the VDAs.

Application failure trends are available for the last 2 hours, 24 hours, 7 days, and month for Premium and Advanced licensed sites. They are available for the last 2 hours, 24 hours, and 7 days for other license types. The application failures that are logged to the Event Viewer with source “Application Errors” are monitored. Click **Export** to generate reports in CSV, Excel, or PDF formats

The screenshot shows the Citrix Virtual Apps and Desktops Service interface. The main content area is titled 'Application Failures' and includes a search and filter section. Below this is a table of 'Application Fault Details'. A tooltip is displayed over the first row of the table, providing detailed error information for a 'ThrowException.exe' failure.

Time	Application Name	Process Name	Version	Machine Name
01/17/2019 11:53 AM	ThrowException	ThrowException.exe	1.0.0.0	BVT/INXRD52
01/17/2019 11:53 AM	PassArguments	PassArguments.exe	1.0.0.0	BVT/INXRD52
01/17/2019 11:52 AM	Unknown	CdeEngine.exe	7.21.101.0	BVT/INXRD52

The failures are displayed as **Application Faults** or **Application Errors** based on their severity. The Application Faults tab displays failures associated with loss of functionality or data. Application Errors indicate problems that are not immediately relevant; they signify conditions that might cause future problems.

You can filter the failures based on **Published Application Name**, **Process Name** or **Delivery Group**, and **Time Period**. The table displays the fault or error code and a brief description of the failure. The detailed failure description is displayed as a tooltip.

#### Note:

The Published Application name is displayed as “Unknown” when the corresponding application name cannot be derived. This typically occurs when a launched application fails in a desktop session or when it fails due to an unhandled exception caused by a dependent executable.

By default, only faults of applications hosted on Multi-session OS VDAs are monitored. You can modify the monitoring settings through the Monitoring Group Policies: Enable monitoring of application failures, Enable monitoring of application failures on Single session OS VDAs, and List of applications excluded from failure monitoring. For more information, see [Policies for application failure monitoring](#) in Monitoring policy settings.

The **Trends > Application Probe Results** page displays the results of application probing executed in the site for the last 24 hours and 7 days. For more details on how to configure application probes, see [Application Probing](#).

## Application probing

March 26, 2021

Application probing automates the process of checking the health of Citrix Virtual Apps that are published in a site. The results of application probing are available in the **Monitor** tab of Citrix Virtual Apps and Desktops service.

Ensure that the endpoint machines running probe agents are Windows machines with Citrix Receiver for Windows Version 4.8 or later, or Citrix Workspace app for Windows (formerly Citrix Receiver for Windows) Version 1808 or later. Workspace app for Unified Windows Platform (UWP) is not supported.

Requirements:

- Endpoint machines running probe agents are Windows machines with Citrix Receiver for Windows Version 4.8 or later, or Citrix Workspace app for Windows (formerly Citrix Receiver for Windows) Version 1906 or later. Workspace app for Unified Windows Platform (UWP) is not supported.
- Monitor and Workspace support the default form-based authentication only. Monitor does not support Single sign-on (SSO) authentication.
- Ensure that Microsoft .NET Framework version 4.7.2 or later is installed on the endpoint machine where you want to install the Probe Agent.

User accounts/permissions required to run Application Probing are as follows:

- A unique Workspace user to probe on each endpoint machine. The Workspace user is not required to be an administrator; the probes can run in a non-admin context.
- User accounts with Windows administrator permissions to install and configure the Citrix Probe Agent on the endpoint machines
- A full administrator user account with the following permissions. Reusing existing user accounts for application probing might log off from the users' active sessions.
  - Delivery group permissions:
    - \* Read-only
  - Director permissions:
    - \* Create/Edit/Remove Probe Configurations
    - \* View Configurations page
    - \* View Trends page

## Configure Application Probing

Configure your application probes to run during off-peak hours across multiple geographies. The comprehensive probe results can help to troubleshoot issues related to the applications, hosting machine or connection before the users experience them.

Citrix Probe Agent version 2103 supports [site aggregation](#). Applications and desktops can be enumerated and launched from aggregated sites. When you configure the probe agent, select the **Workspace (StoreFront) Site Aggregation Enabled** option to enable enumeration of applications and desktops from aggregated sites. The following combinations of sites are supported:



- Multiple on-premises sites having one StoreFront URL.
- On-premises and cloud sites having either a StoreFront or Workspace URL.
- Multiple cloud sites having one Workspace URL.

**Note:**

You must create separate administrators or users to configure probes that have access to only one site.

**Step 1: Install and configure the Citrix Probe Agent**

The Citrix Probe Agent is a Windows executable that simulates the actual application launch by the user through Citrix Workspace. It tests application launches as configured in Monitor and reports back the results to Monitor.

1. Identify endpoint machines from where you want to run application probing.
2. Users with administrative privileges can install and configure the Citrix Probe Agent on the endpoint machine. Download the Citrix Probe Agent executable available at <https://www.citrix.com/downloads/citrix-virtual-apps-and-desktops/components/app-probe-agent.html>
3. Start the agent and configure your Citrix Workspace credentials. Configure a unique Workspace user on each endpoint machine. The credentials are encrypted and stored securely.

**Notes:**

- To access the site to be probed from outside the network, type the login URL for Citrix Gateway in the **Workspace URL** field. Citrix Gateway automatically routes the request to the corresponding site Workspace URL.
- Use NetBIOS as the domain name in the user name field. For example, NetBIOS/user-name.
- App probing supports Citrix Content Collaboration service using Workspace authentication (AD only).

The screenshot shows the 'Citrix Probe Agent' configuration window. On the left, a sidebar lists three steps: '1. Configure Workspace Credentials' (selected), '2. Configure to Display Probe Result', and '3. View Summary'. The main area has a toggle for 'Workspace (StoreFront) Site Aggregation Enabled' which is turned on. Below it are three input fields: 'Workspace URL (StoreFront URL in case of on-premises Site)', 'User name', and 'Password'. A green note below the password field reads 'Provide unique Workspace user credentials on each probe machine'. A 'Next' button is located at the bottom right.

4. On the **Configure To Display Probe Result** tab, enter credentials to access the Citrix Virtual Apps and Desktops service. You can find the Customer Name or Customer ID, Client ID, and Secret Key from the API Access page in the Citrix Cloud console.

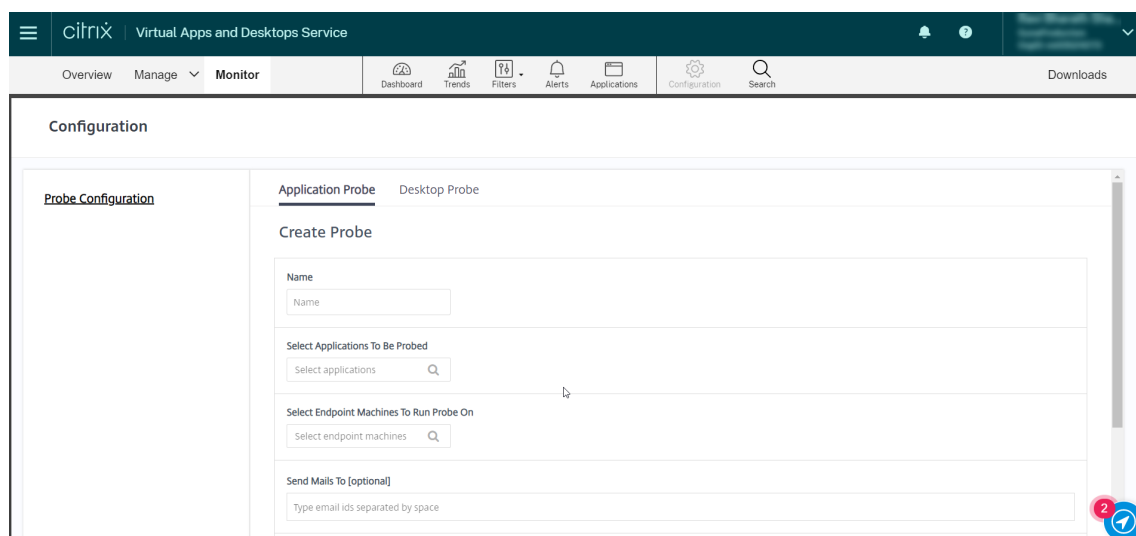
The screenshot shows the 'Citrix Application Probe Agent' configuration window. The sidebar now highlights '2. Configure to Display Probe Result'. The main area features a toggle for 'VIEW THE PROBE RESULT ON CITRIX CLOUD:' set to 'Yes'. Below this are three input fields: 'Client ID', 'Secret Key', and 'Customer Name'. A 'Validate' button is positioned to the right of the 'Customer Name' field. A 'Next' button is at the bottom right.

## Step 2: Configure Application probing in the Monitor tab

1. In the Citrix Virtual Apps and Desktops service, go to **Configuration > Application Probe Configuration**.
2. Create a probe and choose:

- the applications to be probed,
- the endpoint machines on which the probe must run,
- the email addresses to which the failure probe results are sent,
- the time of the day at which the probe must run (as per the local time zone of the endpoint machine).

After configuration in the **Monitor** tab, the agent takes 10 minutes before it is ready to start probing. Then, it runs configured probes starting the next hour.



### Step 3: Probe execution

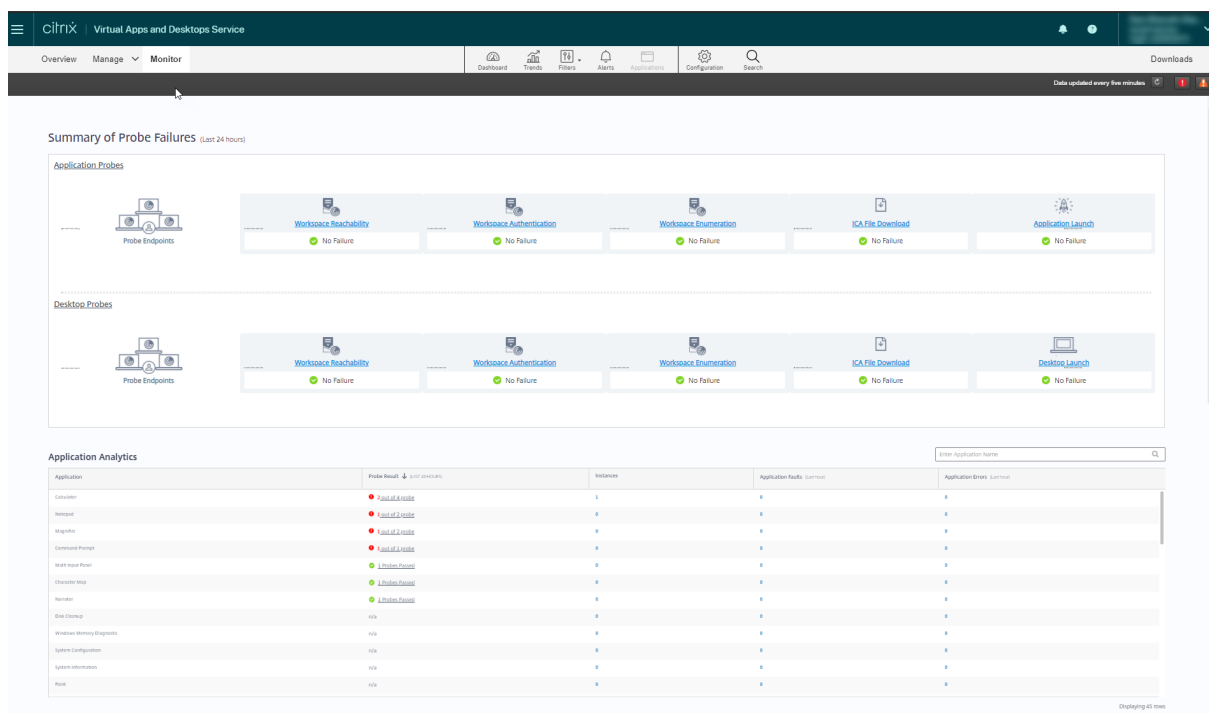
The agent runs application probing as per the probe configuration it fetches from Monitor every hour. It launches selected applications serially using Workspace. The agent reports the results back to Monitor via the Monitor database. Failures are reported in five specific stages:

- **Workspace Reachability** - configured Workspace URL is not reachable.
- **Workspace Authentication** - configured Workspace credentials are invalid.
- **Workspace Enumeration** - Workspace Enumerate applications list does not contain the application to be probed.
- **ICA download** - the ICA file is not available.
- **Application launch** - the application cannot be launched.

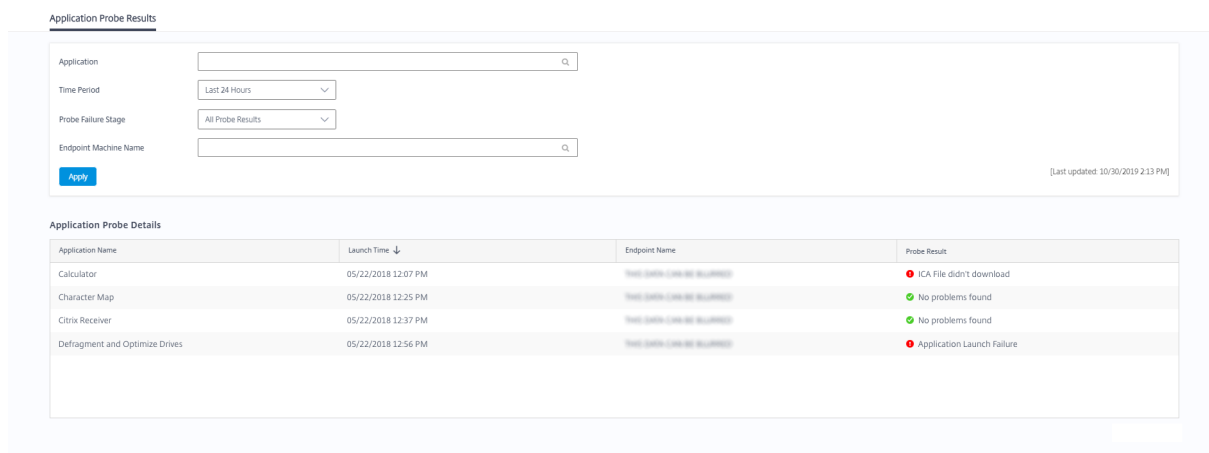
### Step 4: View probe results

You can view the latest probe results in the Citrix Virtual Apps and Desktops service > **Applications** page.

# Citrix Virtual Apps and Desktops service



To troubleshoot further, click the probe result link to see more details on the **Trends > Application Probe Results** page.



The consolidated probe results data is available for the last 24 hours or last 7 days time periods on this page. You can see the stage in which the probe failed. You can filter the table for a specific application, probe failure stage, or endpoint machine.

## Desktop probing

May 17, 2021

Desktop probing automates the process of checking the health of Citrix Virtual Desktops that are published in a site. The results of desktop probing are available in Monitor.

In Monitor's Configuration page, configure the desktops to be probed, the endpoint machines to run the probe on, and the probe time. The agent tests the launch of selected desktops using Workspace and reports the results back to Monitor. The probe results are displayed on the Monitor UI – the last 24-hours' data on the Applications page and historical probe data on the **Trends > Probe Results > Desktop Probe Results** page.

Here, you can see the stage when the probe failure occurred - Workspace Reachability, Workspace Authentication, Workspace Enumeration, ICA download, or Desktop launch. The failure report is sent to the configured email addresses.

You can schedule your desktop probes to run during off-peak hours across multiple geographies. The comprehensive results can help to proactively troubleshoot issues related to provisioned desktops, hosting machines or connections before the users experience them.

Desktop probing is available for Premium licensed sites. This feature requires Probe Agent 1903 or later.

Requirements:

- Endpoint machines running probe agents are Windows machines with Citrix Receiver for Windows Version 4.8 or later, or Citrix Workspace app for Windows (formerly Citrix Receiver for Windows) Version 1906 or later. Workspace app for Unified Windows Platform (UWP) is not supported.
- Monitor and Workspace support the default form-based authentication only. Monitor does not support Single sign-on (SSO) authentication.
- Ensure that Microsoft .NET Framework version 4.7.2 or later is installed on the endpoint machine where you want to install the Probe Agent.

User accounts or permissions required to run Desktop probing:

- A unique Workspace user to probe on each endpoint machine. The Workspace user need not be an administrator; the probes can run in a non-admin context.
- User accounts with Windows administrator permissions to install and configure the Citrix Probe Agent on the endpoint machines
- A full administrator user account or a custom role with the following permissions. Reusing normal user accounts for desktop probing might log off the users' active sessions.
  - Delivery group permissions:
    - \* Read-only
  - Monitor permissions:
    - \* Create, Edit, Remove Alert Email Server Configuration - if the email server is not already configured
    - \* Create, Edit, Remove Probe Configurations

- \* View Configurations page
- \* View Trends page

## Configure desktop probing

You can schedule your desktop probes to run during off-peak hours across multiple geographies. The comprehensive probe results can help to troubleshoot issues related to the desktops, hosting machine or connection before the users experience them.

Citrix Probe Agent version 2103 supports [site aggregation](#). Applications and desktops can be enumerated and launched from aggregated sites. When you configure the probe agent, select the **Workspace (StoreFront) Site Aggregation Enabled** option to enable enumeration of applications and desktops from aggregated sites. The following combinations of sites are supported:

- Multiple on-premises sites having one StoreFront URL.
- On-premises and cloud sites having either a StoreFront or Workspace URL.
- Multiple cloud sites having one Workspace URL.

### Note:

You must create separate administrators or users to configure probes that have access to only one site.

## Step 1: Install and configure the Citrix Probe Agent

The Citrix Probe Agent is a Windows executable that simulates the actual desktop launch by the user through Workspace. It tests desktop launches as configured in Monitor and reports back the results to Monitor.

1. Identify endpoint machines from where you want to run desktop probing.
2. Users with administrative privileges can install and configure the Citrix Probe Agent on the endpoint machine. Download the Citrix Probe Agent executable available at <https://www.citrix.com/downloads/citrix-virtual-apps-and-desktops/components/app-probe-agent.html>
3. Start the agent and configure your Workspace Receiver for Web credentials. Configure a unique Workspace user on each endpoint machine. The credentials are encrypted and stored securely.

### Notes:

- To access the site to be probed from outside the network, type the Citrix Gateway login page URL on the Workspace URL field. Citrix Gateway automatically routes the request to the corresponding site Workspace URL. This feature is available for Citrix Gateway version 12.1 or later.
- Use NetBIOS as the domain name in the user name field. For example, NetBIOS/user-

name.

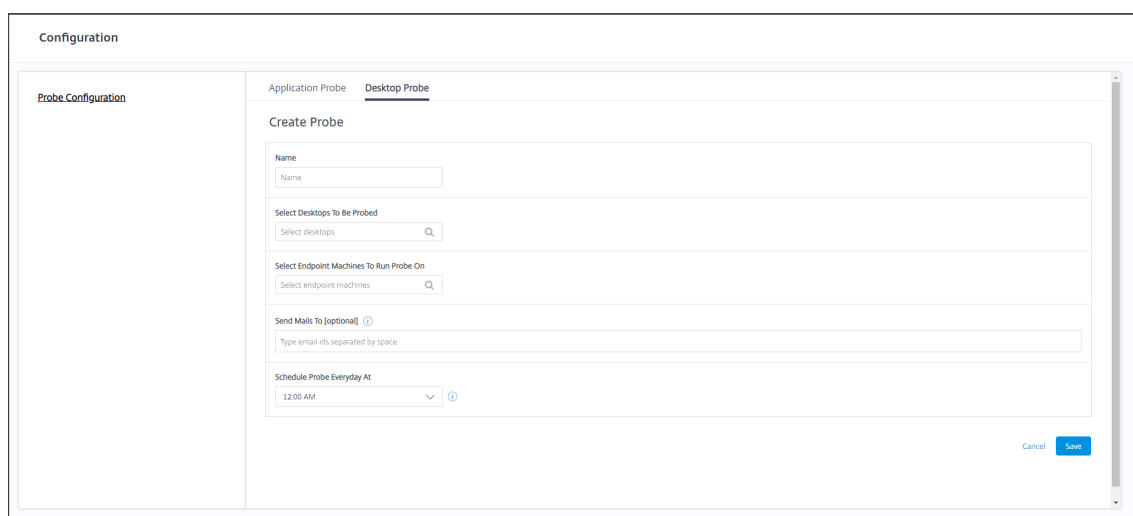
- Desktop probing supports Citrix Content Collaboration service using Workspace authentication (AD only).
- You must enable Interactive Logon for the configured unique StoreFront user.

4. On the **Configure To Display Probe Result** tab, enter your Monitor credentials. You can find the Customer Name or Customer ID, Client ID, and Secret Key from the API Access page in the Citrix Cloud console.

The screenshot shows the 'Citrix Probe Agent' configuration window. On the left, a sidebar lists three steps: '1. Configure Workspace Credentials', '2. Configure to Display Probe Result' (which is selected), and '3. View Summary'. The main area is titled 'VIEW THE PROBE RESULT ON CITRIX CLOUD:' with a toggle switch set to 'Yes'. Below this, there are three input fields: 'Client ID' (containing a masked value), 'Secret Key' (filled with dots), and 'Customer Name' (containing a masked value). To the right of the 'Customer Name' field is a 'Validate' button. At the bottom right of the main area is a 'Next' button.

## Step 2: Configure desktop probing in Monitor

1. Go to **Configuration > Desktop Probe Configuration**.
2. To create a probe, enter the details and click **Save**.



The screenshot shows a web interface for configuring a Desktop Probe. The page is titled "Configuration" and has a sidebar with "Probe Configuration" selected. The main content area is divided into "Application Probe" and "Desktop Probe" tabs, with "Desktop Probe" active. Under "Create Probe", there are several fields: "Name" (text input), "Select Desktops To Be Probed" (dropdown menu), "Select Endpoint Machines To Run Probe On" (dropdown menu), "Send Mails To (optional)" (text input with a help icon), and "Schedule Probe Everyday At" (dropdown menu set to "12:00 AM" with a help icon). "Cancel" and "Save" buttons are at the bottom right.

### Note:

Configure your email server in **Alerts > Email Server Configuration**.

After desktop probing configuration is complete, the agent takes 10 minutes before it is ready to start probing. Then, it runs the configured probes starting the next hour.

### Step 3: Probe execution

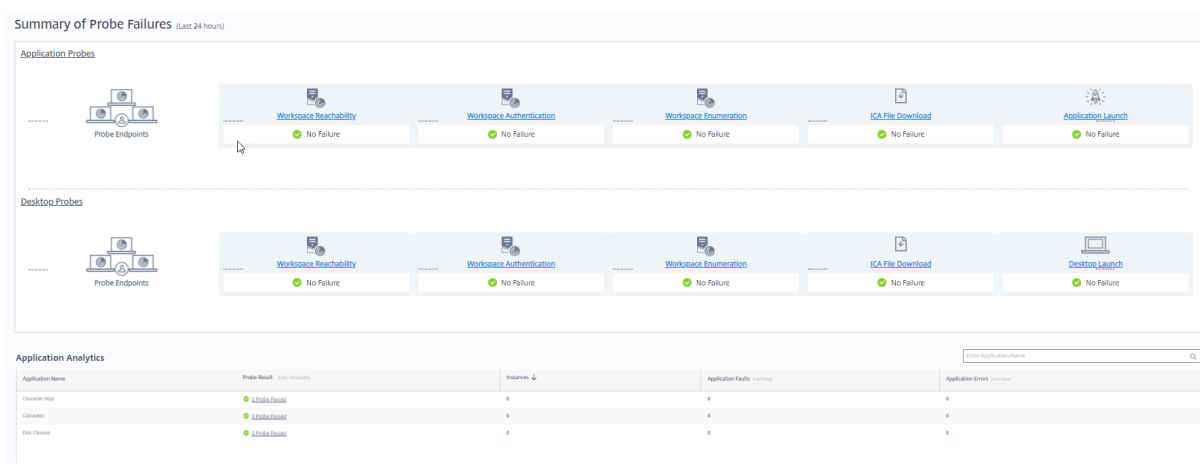
The agent runs desktop probing as per the probe configuration it fetches from Monitor periodically. It launches selected desktops serially using Workspace. The agent reports the results back to Monitor via the Monitor database. Failures are reported in five specific stages:

- **Workspace Reachability** - configured Workspace URL is not reachable.
- **Workspace Authentication** - configured Workspace credentials are invalid.
- **Workspace Enumeration** - Workspace Enumerate desktops list does not contain the desktop to be probed.
- **ICA download** - the ICA file is not available.
- **Desktop launch** - the desktop cannot be launched.

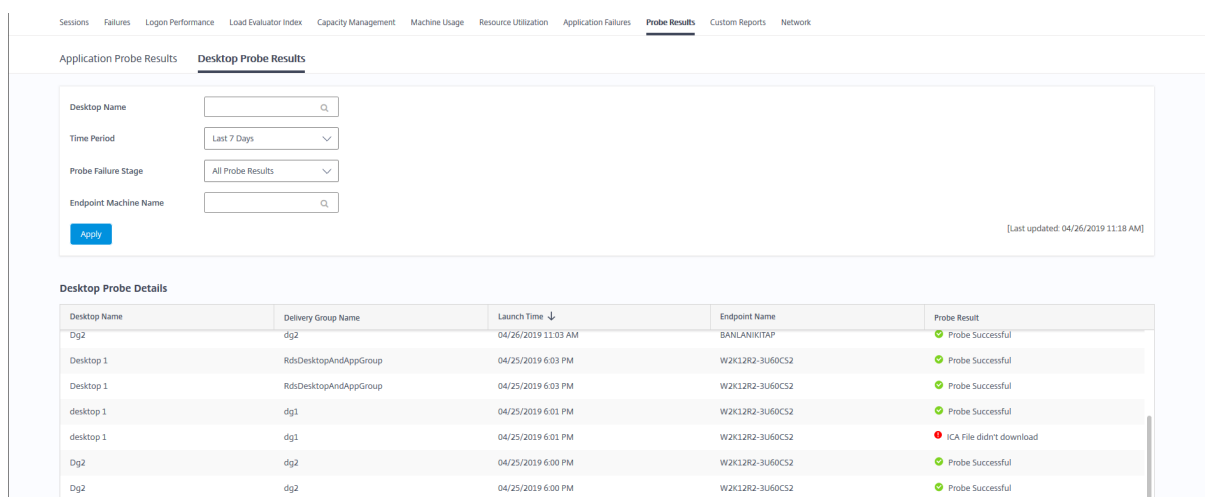
### Step 4: View probe results

You can view the latest probe results on the **Desktops** page.





To troubleshoot further, click the probe result link to see more details on the **Trends > Probe Results > Desktop Probe Results** page.



The consolidated probe results data is available for the last 24 hours or last 7 days' time periods on this page. You can see the stage in which the probe failed. You can filter the table for a specific desktop, probe failure stage, or endpoint machine.

## Troubleshoot machines

November 10, 2020

### Note:

**Citrix Health Assistant** is a tool to troubleshoot configuration issues in unregistered VDAs. The tool automates a number of health checks to identify possible root causes for VDA registration failures and issues in session launch and time zone redirection configuration. The Knowledge

Center article, [Citrix Health Assistant - Troubleshoot VDA Registration and Session Launch](#) contains the **Citrix Health Assistant** tool download and usage instructions.

The **Filters > Machines** view in the Monitor tab displays the machines configured in the site. The Multi-session OS Machines tab includes the load evaluator index, which indicates the distribution of performance counters and tooltips of the session count if you hover over the link.

Click the **Failure Reason** column of a failed machine to get a detailed description of the failure and actions recommended to troubleshoot the failure. The failure reasons and the recommended actions for machine and connection failures are available in the [Citrix Director Failure Reasons Troubleshooting Guide](#).

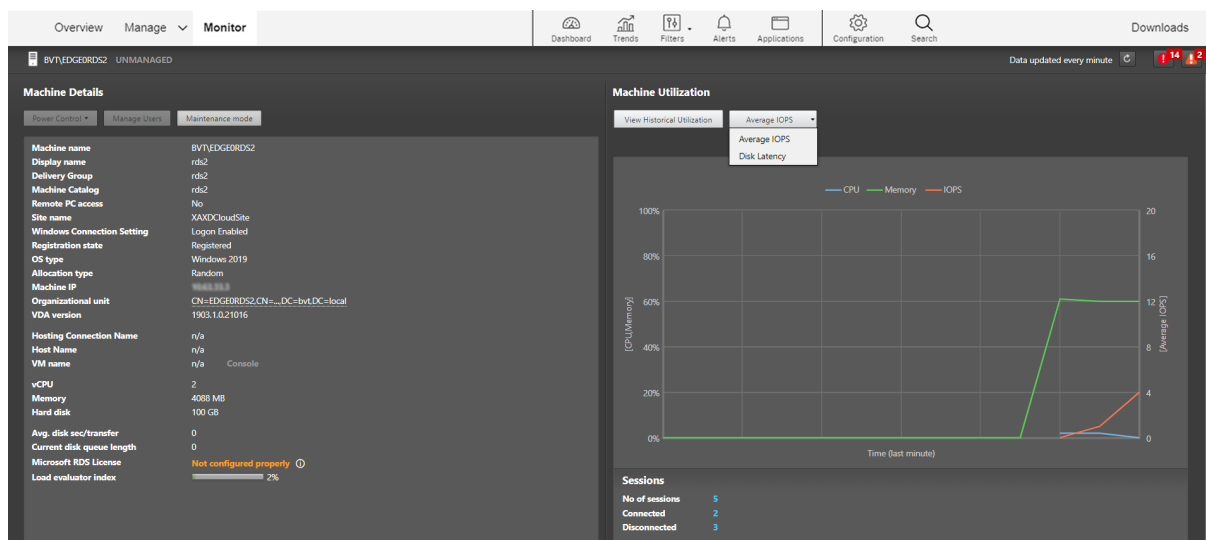
Click the machine name link to go to the **Machine Details** page.

The Machine Details page lists the machine details, infrastructure details, and details of the hotfixes applied on the machine.

## Machine-based real-time resource utilization

The **Machine Utilization** panel displays graphs showing real-time utilization of CPU and memory. In addition, disk and GPU monitoring graphs are available for sites with VDA versions **7.14** or later.

Disk monitoring graphs, average IOPS, and disk latency are important performance measurements that help you monitor and troubleshoot issues related to VDA disks. The Average IOPS graph displays the average number of reads and writes to a disk. Select **Disk Latency** to see a graph of the delay between a request for data and its return from the disk, measured in milliseconds.



Select **GPU Utilization** to see percentage utilization of the GPU, the GPU memory, and of the Encoder and the Decoder to troubleshoot GPU-related issues on Multi-session or Single session OS VDAs. The GPU Utilization graphs are available only for VDAs running 64-bit Windows with NVIDIA Tesla M60 GPUs, and running Display Driver version 369.17 or later.

The VDAs must have HDX 3D Pro enabled to provide GPU acceleration. For more information, see GPU acceleration for Windows Single session OS and GPU acceleration for Windows Multi-session OS. When a VDA accesses more than one GPU, the utilization graph displays the average of the GPU metrics collected from the individual GPUs. The GPU metrics are collected for the entire VDA and not for individual processes.

### **Machine-based historical resource utilization**

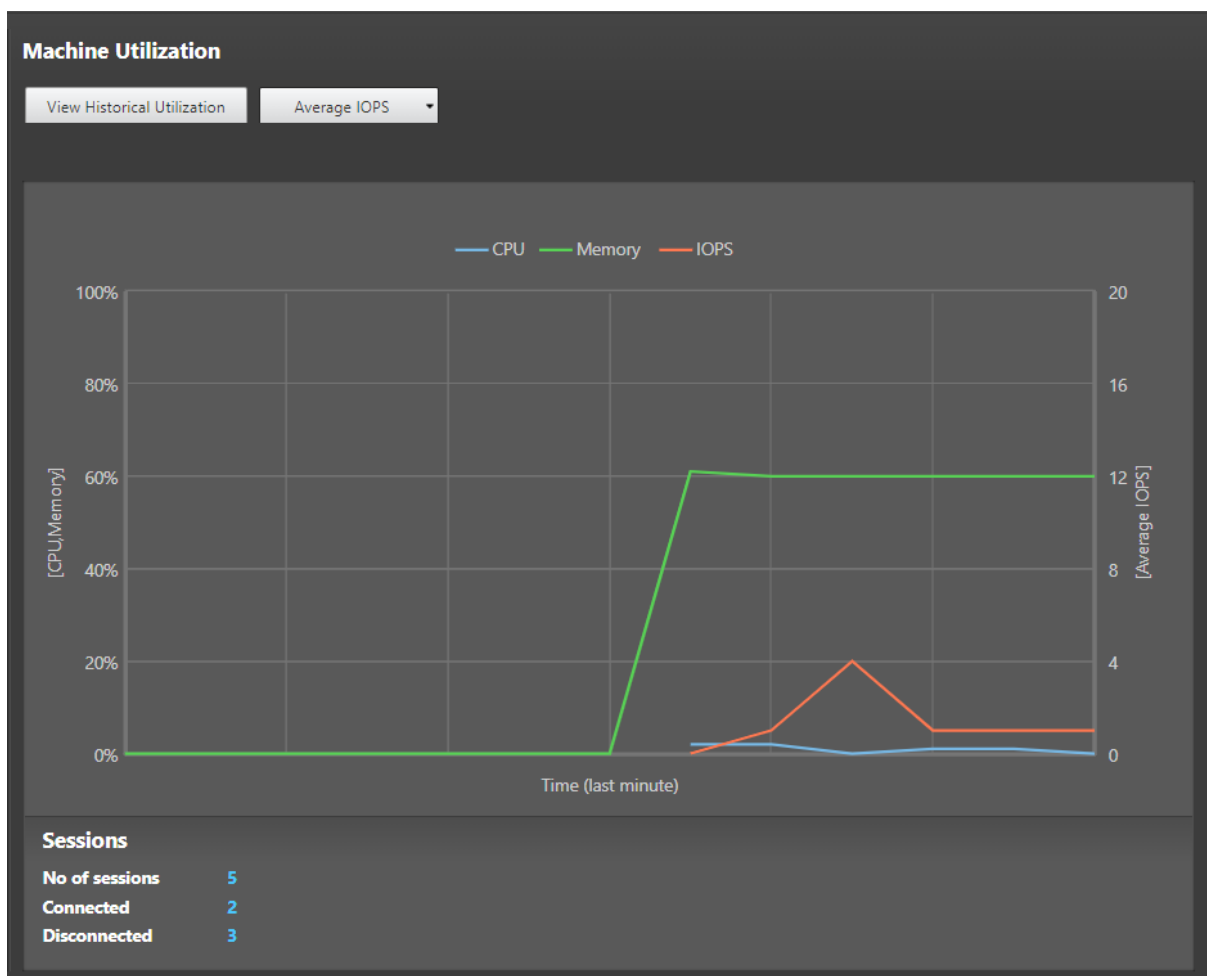
In the **Machine Utilization** panel, click **View Historical Utilization** to view the historical usage of resources on the selected machine.

The utilization graphs include critical performance counters of CPU, memory, peak concurrent sessions, average IOPS, and disk latency.

**Note:**

The Monitoring policy setting, **Enable Process Monitoring**, must be set to Allowed to collect, and display data in the Top 10 Processes table on the Historic Machine Utilization page. The collection is prohibited by default.

The CPU and memory utilization, average IOPS, and disk latency data is collected by default. You can disable the collection by using the **Enable Resource Monitoring** policy setting.

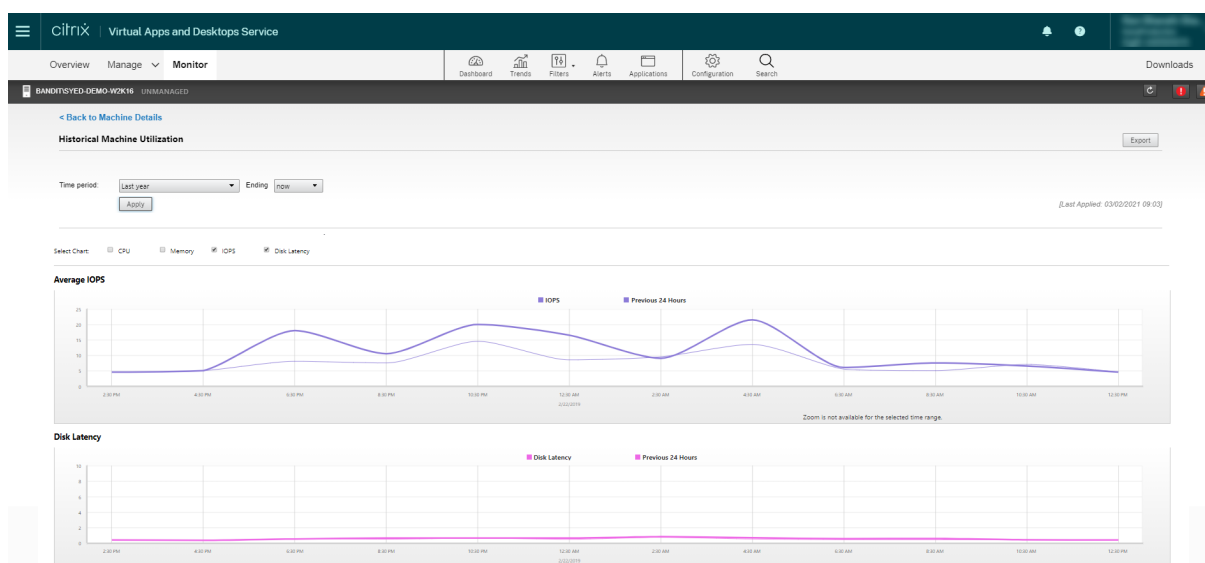


1. From the **Machine Utilization** panel in the **Machine Details** view, select **View Historical Utilization**.
2. In the **Historical Machine Utilization** page, set **Time Period** to view usage for the last 2 hours, 24 hours, 7 days, month, or year.

**Note:**

Average IOPS and disk latency usage data are available only for the last 24 hours, month, and year ending now. Custom end time is not supported.

3. Click **Apply** and select the required graphs.
4. Hover over different sections of the graph to view more information for the selected time period.



For example, if you select **Last 2 hours**, the baseline period is the 2 hours prior to the selected time range. View the CPU, memory, and session trend over the last 2 hours and the baseline time. If you select **Last month**, the baseline period is the previous month. Select to view the Average IOPS and disk latency over the last month and the baseline time.

1. Click **Export** to export the resource utilization data for the selected period. For more information, see [Export reports](#) section in Monitor Deployments.
2. Below the graphs, the table lists the top 10 processes based on CPU or memory utilization. You can sort by any of the columns, which show Application Name, User Name, Session ID, Average CPU, Peak CPU, Average Memory, and Peak Memory over the selected time range. The IOPS and Disk Latency columns cannot be sorted.

**Note:**

The session ID for system processes is displayed as “0000”.

3. To view the historical trend on the resource consumption of a particular process, drill into any of the Top 10 processes.

### Machine Console access

You can access the consoles of Desktop and Multi-session OS machines hosted on XenServer Version 7.3 and later directly from Monitor. This way, you don’t require XenCenter to troubleshoot issues on XenServer hosted VDAs. For this feature to be available, the XenServer hosting the machine must be of Version 7.3 or later and must be accessible from the Monitor.



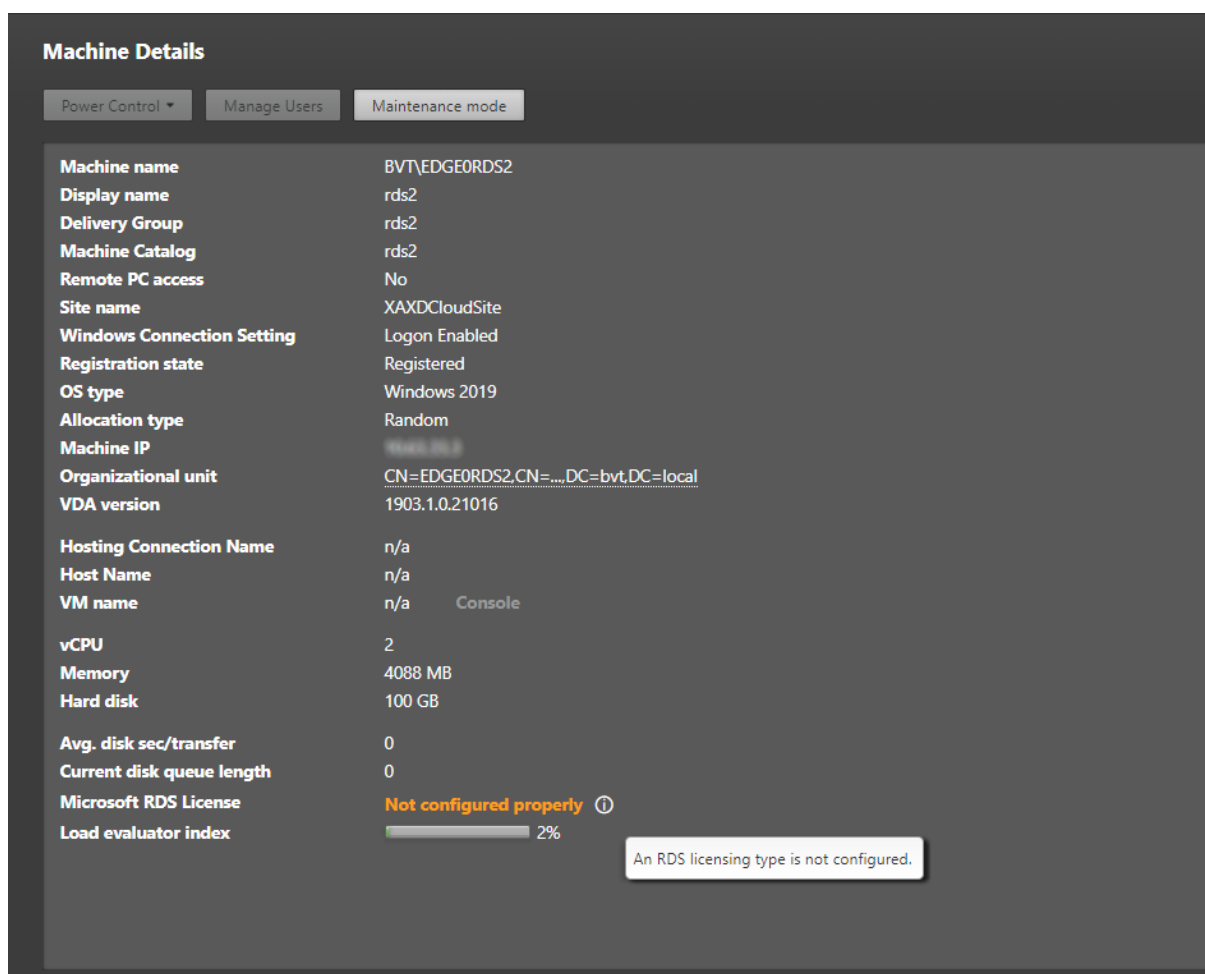
To troubleshoot a machine, click the **Console** link in the corresponding Machine Details panel. After authentication of the host credentials you provide, the machine console opens in a separate tab using noVNC, a web-based VNC client. You now have keyboard and mouse access the console.

**Note:**

- This feature is not supported on Internet Explorer 11.
- If the mouse pointer on the machine console is misaligned, see [CTX230727](#) for steps to fix the issue.
- Console access is launched on a new tab, ensure that your browser settings allow pop-ups.
- For security reasons, Citrix recommends that you install SSL certificates on your browser.

### Microsoft RDS license health

You can view the status of Microsoft RDS license in the Machine Details panel in the **Machine Details** and the **User Details** page for Multi-session OS machines.



One of the following messages is displayed:

- License available
- Not configured properly (warning)
- License error (error)
- Incompatible VDA version (error)

**Note:**

The RDS license health status for machines under grace period with valid license displays a **License available** message in green. Renew your license before they expire.

For warning and error messages, hover over the info icon to view additional information as given in the following table.

Message Type	Messages in Monitor
Error	Available for VDAs version 7.16 and later.
Error	New RDS connections are not allowed.

Message Type	Messages in Monitor
Error	RDS licensing has exceeded its grace period.
Error	A License Server is not configured for the required OS level with the Per Device Client Access licensing type.
Error	The configured License Server is incompatible with the RDS Host OS level with the Per Device Client Access licensing type.
Warning	Personal Terminal Server is not a valid RDS licensing type in a Citrix Virtual Apps and Desktops deployment.
Warning	Remote Desktop for Administration is not a valid licensing type in a Citrix Virtual Apps and Desktops deployment.
Warning	An RDS licensing type is not configured.
Warning	The Domain Controller or License Server is unreachable with the Per User Client Access RDS licensing type.
Warning	With the Per Device Client Access licensing type, the Client Device license could not be determined since the license server for the required OS level is unreachable.

**Note:**

This feature is applicable only for Microsoft RDS CAL (Client Access License).

**PVS target device metrics**

You can view the status of PVS target devices for single-session and multi-session OS machines on the **Machine Details** page in Director. Several metrics for **Network**, **Boot**, and **Cache** are available on this panel. These metrics help you monitor and troubleshoot PVS target devices to ensure that they are up and running.



PVS Target Device Metrics					
Network		Boot		Cache	
NIC Bandwidth Utilization (%)	12	Boot Bytes Read MB	231	Write Cache Type	Device RAM with overflow on local har...
Server Reconnect Count	5	Boot Bytes Written MB	0	Write Cache Volume Drive Letter	D:
Total UDP Retry Count	7	Boot From	vDisk	Write Cache Volume Size MB	6142
		Boot Retry Count	0	Cache File Size MB	1058
		Boot Time (sec)	31	Ram Cache Usage MB	62.3125
		Target Software Version	7.23.0		
		vDisk Name	v10vDisk.vhdx		

### Network:

- **Network Bandwidth Utilization:** Average bandwidth utilization across all NICs.
- **Server Reconnect Count:** Number of times the server has reconnected due to network issues or server rebalancing or shutdowns and restarts of the Citrix Provisioning Stream Service.
- **Total UDP Retry Count:** Number of times the provisioning target device has tried to reconnect to the provisioning server using UDP. This metric helps you to know if there are any network issues in the Citrix Provisioning Stream Service (for example, bad switch configurations).

### Boot:

- **Boot Bytes Read MB:** Bytes read while booting.
- **Boot Bytes Written MB:** Bytes written while booting.
- **Boot From:** Boot medium (vDisk, local disk, and so on).
- **Boot Retry Count:** Number of retries to boot the machine.
- **Boot Time:** Time taken to boot the machine, in seconds. By default, there is a 5 second delay between retries. If this delay grows into double digits, there is a significant increase in boot time. Check your provisioning configuration to resolve this issue.
- **Target Software Version:** Version of the Provisioning target device software.
- **vDisk Name:** vDisk from which the Provisioning target device is booting.

### Cache:

- **Write Cache Type:** vDisk can be set to different types of cache. For more information, see Knowledge Center article [CTX119469](#).
- **Write Cache Volume Drive Letter:** Drive letter for write cache types involving drives.
- **Write Cache Volume Size MB:** Total configured volume size for write cache.
- **Cache File Size MB:** Current cache file size (cache on device RAM with overflow on hard disk).
- **Ram Cache Usage MB:** Current RAM cache size (cache on device RAM with overflow on hard disk). Use Overflow to disk only if necessary. This metric is useful when setting or optimizing the proper size of RAM cache.

For more information, see [Using the Status Tray on a target device](#).

Provisioning target device metrics is available only on:

- Provisioning machines.
- Provisioning target device version 7.19 and later.
- VDA version 2003 and later.

**Note:**

Metrics for Server Reconnect Count and UDP Retry Count are available only for Provisioning target version 1912 CU2 and later.

## Troubleshoot user issues

September 14, 2020

Use the Monitor's **Help Desk** view (**Activity Manager** page) to view information about the user:

- Check for details about the user's logon, connection, and applications.
- Shadow the user's machine.
- Troubleshoot the issue with the recommended actions in the following table, and, if needed, escalate the issue to the appropriate administrator.

### Troubleshooting tips

---

User issue	Suggestions
Logon takes a long time or fails intermittently or repeatedly	<a href="#">Diagnose user logon issues</a>
Session startup takes a long time or fails intermittently or repeatedly	<a href="#">Diagnose session startup issues</a>
Application is slow or won't respond	<a href="#">Resolve application failures</a>
Connection failed	<a href="#">Restore desktop connections</a>
Session is slow or not responding	<a href="#">Restore sessions</a>
Video is slow or poor quality	<a href="#">Run HDX channel system reports</a>

---

**Note:**

To make sure that the machine is not in maintenance mode, from the User Details view, review the Machine Details panel.

## Search tips

Search for username is conducted across all configured Active Directories.

When you type a multiuser machine name in a Search field, the Machine Details for the specified machine is displayed.

When you type an endpoint name in a Search field, the unauthenticated (anonymous) and authenticated sessions that are connected to a specific endpoint are listed. This enables troubleshooting unauthenticated sessions. Ensure that endpoint names are unique to enable troubleshooting of unauthenticated sessions.

The search results also include users who are not currently using or assigned to a machine.

- Searches are not case-sensitive.
- Partial entries produce a list of possible matches.
- After you type a few letters of a two-part name (username, family name and first name, or display name), separated by a space, the results include matches for both strings. For example, if you type jo rob, the results might include strings such as “John Robertson” or Robert, Jones.

To return to the landing page, click the Monitor tab.

## Diagnose session startup issues

September 14, 2020

In addition to the logon process phases mentioned in the [Diagnose user logon issues](#) section, Monitor displays the session startup duration. This duration is divided into the Workspace App Session Startup duration and the VDA Session Startup duration on the **User Details** and **Endpoint Details** pages. These two durations further contain individual phases whose startup durations are also displayed. This data helps you to understand and troubleshoot high session startup duration. Further, the time duration for each phase involved in the session startup helps in troubleshooting issues associated with individual phases. For example, if the Drive Mapping time is high, you can check to see whether all the valid drives are mapped correctly in the GPO or script.

## Prerequisites

Ensure that the following prerequisites are met for session startup duration data to be displayed:

- VDA 1903 or later.
- Citrix End User Experience Monitoring (EUEM) service must be running on the VDA.

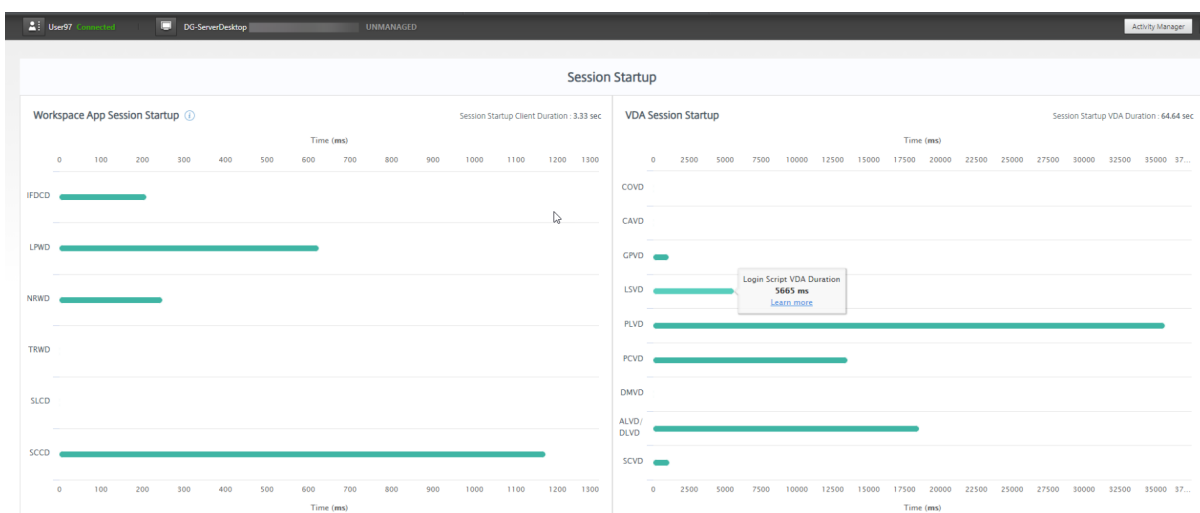
## Limitations

The following limitations apply when Monitor displays the session startup duration data:

- Session startup duration is available only for HDX sessions.
- For session launches from iOS and Android OS, only VDA Startup Duration is available.
- IFDCD is available only when Workspace App is detected while launching from a browser.
- For session launches from macOS, IFDCD is available for Workspace App 1902 and later only.
- For session launches from Windows OS, IFDCD is available for Workspace app 1902 and later. For earlier versions, IFDCD is displayed for only app launches from a browser with Workspace app detected.

### Notes:

- If you face issues in the sessions startup duration display after the prerequisites are met, view the Monitor server and VDA logs as described in [CTX130320](#).  
For shared sessions (multiple applications launched in the same session), the Workspace App Startup metrics are displayed for the latest connection or the latest application launch.
- Some metrics in VDA Session Startup are not applicable on reconnects. In such cases, a message is displayed.



## Workspace App session startup phases

### Session Startup Client Duration (SSCD)

When this metric is high, it indicates a client-side issue that is causing long start times. Review subsequent metrics to determine the probable root cause of the issue. SSCD starts as close as possible

to the time of the request (mouse click) and ends when the ICA connection between the client device and VDA has been established. For a shared session, this duration is much smaller, as much of the setup costs associated with the creation of a new connection to the server are not incurred. At the next level down, there are several detailed metrics available.

### **ICA File Download Duration (IFDCD)**

IFDCD is the time taken for the client to download the ICA file from the server. The overall process is as follows:

1. The user clicks a resource (application or desktop) on the Workspace Application.
2. A request from the user is sent to StoreFront through the Citrix Gateway (if configured), which sends the request to the Delivery Controller.
3. The Delivery Controller finds an available machine for the request and sends the machine information and other details to StoreFront. Also, StoreFront requests and receives a one-time ticket from the Secure Ticket Authority.
4. StoreFront generates an ICA File and sends it to the user via Citrix Gateway (if configured).

IFDCD represents the time it takes for the complete process (steps 1–4). The IFDCD duration stops counting when the client receives the ICA file.

LPWD is the StoreFront component of the process.

If IFDCD is high (but LPWD is normal), the server-side processing of the launch was successful, but there were communication issues between the client device and the StoreFront. This results from network issues between the two machines. So you can troubleshoot potential network issues first.

### **Launch Page Web Server Duration (LPWD)**

This is the time taken to process the launch page (launch.aspx) on the StoreFront. If LPWD is high, there might be a bottleneck on the StoreFront.

Possible causes include:

- High load on the StoreFront. Try to identify the cause of slowdown by checking the Internet Information Services (IIS) logs and monitoring tools, Task Manager, Performance Monitor and so on.
- StoreFront is having issues communicating with other components such as Delivery Controller. Check if the network connection between StoreFront and Delivery Controller is slow or some Delivery Controllers are down or overloaded.

### **Name Resolution Web Server Duration (NRWD)**

This is the time taken by the Delivery Controller to resolve the name of a published application/desktop to a VDA Machine IP Address.

When this metric is high, it indicates that the Delivery Controller is taking a long time to resolve the name of a published application to an IP address. Possible causes include:

- a problem on the client
- issues with the Delivery Controller, such as the Delivery Controller being overloaded, or a problem with the network link between them

### **Ticket Response Web Server Duration (TRWD)**

This duration indicates the time it takes to get a ticket (if necessary) from the Secure Ticket Authority (STA) Server or Delivery Controller. When this duration is high, it indicates that the STA server or the Delivery Controller are overloaded.

### **Session Look-up Client Duration (SLCD)**

This duration represents the time taken to query every session to host the requested published application. The check is performed on the client to determine whether an existing session can handle the application launch request. The method used depends on whether the session is new or shared.

### **Session Creation Client Duration (SCCD)**

This duration represents the time taken to create a session, from the moment wfica32.exe (or a similar equivalent file) is launched to the time when the connection is established.

## **VDA session startup phases**

### **Session Startup VDA Duration (SSVD)**

This duration is the high-level server-side connection start-up metric that indicates the time VDA takes to perform the entire start-up operation. When this metric is high, it indicates that there is a VDA issue increasing session start times. This includes the time spent on the VDA performing the entire start-up operation.

### **Credentials Obtention VDA Duration (COVD)**

The time taken for the VDA to obtain the user credentials.

This duration can be artificially inflated if a user fails to provide credentials in a timely manner, and thus, not included in the VDA Startup Duration. This time is likely to be a significant only if manual login is being used and the server side credentials dialog is displayed (or if a legal notice is displayed before login commences).

#### **Credentials Authentication VDA Duration (CAVD)**

This is the time taken by the VDA to authenticate the user's credentials against the authentication provider, which can be Kerberos, Active Directory, or a Security Support Provider Interface (SSPI).

#### **Group Policy VDA Duration (GPVD)**

This duration is the time taken to apply group policy objects during logon.

#### **Login Script Execution VDA Duration (LSVD)**

This is the time taken by the VDA to run the user's login scripts.

You can make the user or group's login scripts asynchronous. Optimize any application compatibility scripts or use environment variables instead.

#### **Profile Load VDA Duration (PLVD)**

This is the time taken by the VDA to load the user's profile.

If this duration is high, review your User Profile configuration. Roaming profile size and location contribute to slow session starts. When a user logs on to a session where Terminal Services roaming profiles and home folders are enabled, the roaming profile contents and access to that folder are mapped during logon, which takes extra resources. Sometimes, this can consume significant amount of the CPU usage. Use the **Terminal Services home** folders with redirected personal folders to mitigate this problem. In general, use Citrix Profile Management to manage user profiles in Citrix environments. If you are using Citrix Profile Management and have slow logon times, check if your antivirus software is blocking the Citrix Profile Management tool.

#### **Printer Creation VDA Duration (PCVD)**

This is the time taken for the VDA to map the user's client printers synchronously. If the configuration is set for printer creation to be performed asynchronously, no value is recorded for PCVD as it does not impact completion of the session startup.

Excessive time spent in mapping printers is often the result of the printer auto creation policy settings. The number of printers added locally on the users' client devices and your printing configuration can

directly affect your session start times. When a session starts, Citrix Virtual Apps and Desktops have to create every locally mapped printer on the client device. Reconfigure your printing policies to reduce the number of printers that get created, specifically when users have many local printers. To do this, edit the Printer Auto creation policy in Delivery Controller and Citrix Virtual Apps and Desktops.

### **Drive Mapping VDA Duration (DMVD)**

This is the time taken by the VDA to map the user's client drives, devices, and ports.

Ensure that your base policies include settings to disable unused virtual channels, such as audio or COM port mapping, to optimize the ICA protocol and improve overall session performance.

### **Application/Desktop Launch VDA Duration (ALVD/DLVD)**

This phase is a combination of userinit and Shell duration. When a user logs on to a Windows machine, Winlogon runs userinit.exe. Userinit.exe runs logon scripts, re-establishes network connections, and then starts explorer.exe, the Windows User interface. userinit represents the duration between the start of userinit.exe to the start of the user interface for the virtual desktop or application. The Shell duration is the time between the initialization of the user interface to the time the user receives keyboard and mouse control.

### **Session Creation VDA Duration (SCVD)**

This time includes miscellaneous delays in session creation on VDA.

## **Diagnose user logon issues**

November 10, 2020

Use Logon Duration data to troubleshoot user logon issues.

Logon duration is measured only for initial connections to a desktop or app using HDX. This data does not include users trying to connect with Remote Desktop Protocol or reconnect from disconnected sessions. Specifically, logon duration is not measured when a user initially connects using a non-HDX protocol and reconnects using HDX.

In the User Details view, the duration is displayed as a number value below which the time the logon occurred is displayed and a graph of the phases of the logon process.

As users logon to Citrix Virtual Apps and Desktops, the Monitor Service tracks the phases of the logon process from the time the user connects from Citrix Workspace app to the time when the desktop is ready to use.



The large number on the left is the total logon time and is calculated by combining the time spent establishing the connection and obtaining a desktop from the Delivery Controller with the time spent to authenticate and log on to a virtual desktop. The duration information is presented in seconds (or fractions of seconds).

## Prerequisites

Ensure that the following prerequisites are met for logon duration data and drilldowns to appear:

1. Install **Citrix User Profile Manager** and **Citrix User Profile Manager WMI Plugin** on the VDA.
2. Ensure that the Citrix Profile Management Service is running.
3. For XenApp and XenDesktop sites 7.15 and earlier, disable the GPO setting, **Do not process the legacy run list**.
4. Audit process tracking must be enabled for Interactive Session drilldown.
5. For GPO drilldown, increase the size of Group Policy operational logs.

### Note:

Logon duration is supported only on the default Windows shell (explorer.exe) and not on custom shells.

## Steps to troubleshoot user logon issues

1. From the **User Details** view, troubleshoot the logon state using the Logon Duration panel.
  - If the user is logging on, the view reflects the process of logging on.
  - If the user is logged on, the Logon Duration panel displays the time it took for the user to log on to the current session.
2. Examine the phases of the logon process.

## Logon process phases

### Brokering

Time taken to decide which desktop to assign to the user.

### VM start

If the session required a machine start, this is the time taken to start the virtual machine.

### HDX connection

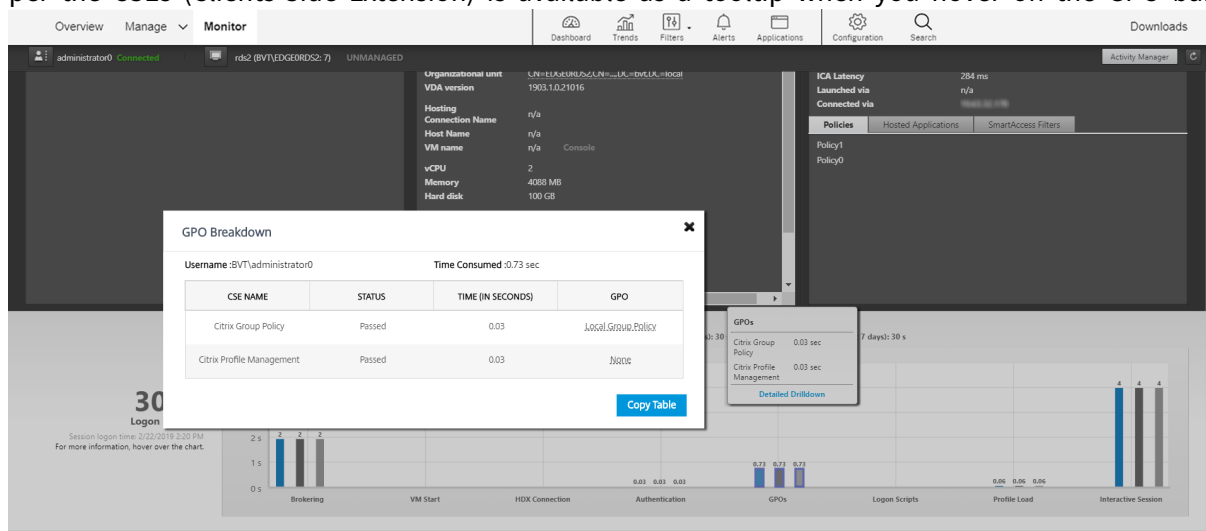
Time taken to complete the steps required in setting up the HDX connection from the client to the virtual machine.

## Authentication

Time taken to complete authentication to the remote session.

## GPOs

If Group Policy settings are enabled on the virtual machines, this is the time taken to apply group policy objects during logon. The drill-down of the time taken to apply each policy as per the CSEs (Clients-Side Extension) is available as a tooltip when you hover on the GPO bar.



Click **Detailed Drilldown** to see a table with the policy status, and the corresponding GPO name. The time durations in the drilldown represent the CSE processing time only and do not add up to the total GPO time. You can copy the drill-down table for further troubleshooting or use in reports. The GPO time for the policies is retrieved from Event Viewer logs. The logs can get overwritten depending on the memory allocated for the operational logs (default size is 4 MB). For more information about increasing the log size for the operational logs, see the Microsoft TechNet article [https://docs.microsoft.com/en-us/previous-versions/tn-archive/dd277416\(v=technet.10\)](https://docs.microsoft.com/en-us/previous-versions/tn-archive/dd277416(v=technet.10)).

## Logon scripts

If logon scripts are configured for the session, this is the time taken for the logon scripts to be run.

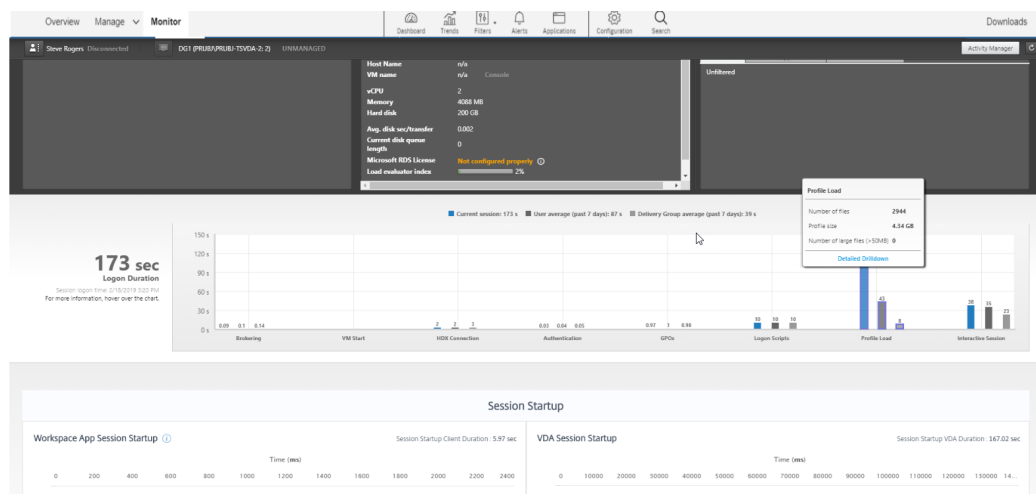
## Profile load

If profile settings are configured for the user or the virtual machine, this is the time taken for the profile to load.

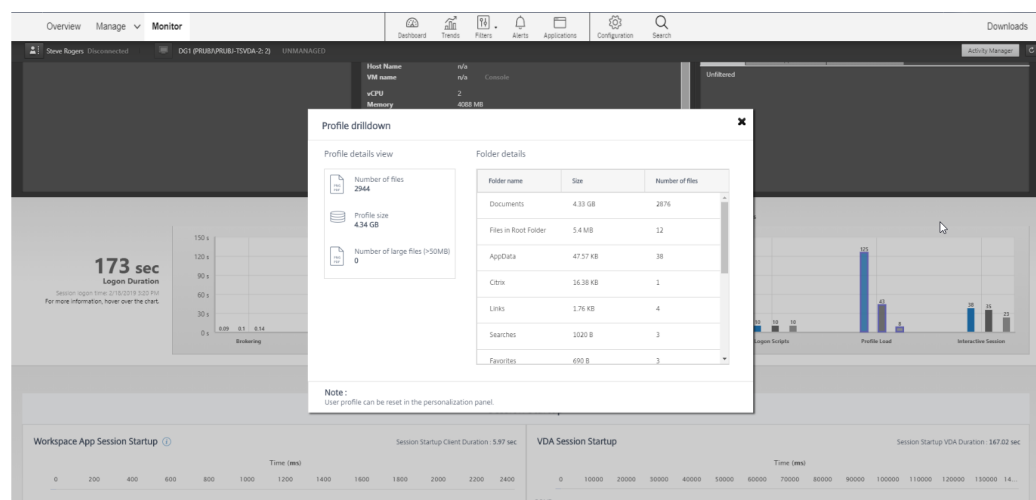
If Citrix Profile Management is configured, the Profile Load bar includes the time taken by Citrix Profile Management to process user profiles. This information helps administrators to troubleshoot high profile processing duration issues. When Profile Management is configured, the Profile Load bar displays

an increased duration. This increase is caused by this enhancement and does not reflect a performance degradation. This enhancement is available on VDAs 1903 and later.

Hovering over the Profile Load bar displays a tooltip showing the user profile details for the current session. This additional information can help troubleshoot high profile load issues.



Click **Detailed Drilldown** to drilldown further into each individual folder in the profile root folder (for instance, C:/Users/username), its size, and the number of files (including files inside nested folders).



Profile drilldown is available on VDAs 1811 and later. Using the profile drilldown information, you can resolve issues involving a high profile load time. You can:

- Reset the user profile
- Optimize the profile by removing unwanted large files
- Reduce the number of files to reduce the network load
- Use profile streaming

By default, all folder names are visible. To hide the folder names, edit the registry values on the VDA machine using the following steps:

**Warning:**

Adding and editing the registry incorrectly can cause serious problems that might require you to reinstall your operating system. Citrix does not guarantee that problems resulting from the incorrect use of the Registry Editor can be solved. Use Registry Editor at your own risk. Be sure to back up the registry before you edit it.

1. On the VDA, add a new registry value **ProfileFoldersNameHidden** at HKEY\_LOCAL\_MACHINE\Software\Citrix
2. Set the value to 1. This value must be a DWORD (32-bit) value. Folder names visibility is now disabled.
3. To make the folder names visible again, set the value to 0.

**Note:**

You can use GPO or PowerShell to apply the registry value change on multiple machines. For more information about using GPO to deploy registry changes, see the [blog](#).

**Additional information**

- Profile drilldown does not consider redirected folders.
- NTUser.dat files in the root folder might not be visible to end users. However, they are included in the profile drilldown and displayed in the list of files in **Root Folder**.
- There are some hidden files in AppData folder which are not included in profile drilldown.
- Number of files and profile size data might not match with the data in the Personalization panel due to certain Windows limitations.

**Interactive Session**

This is the time taken to “hand off” keyboard and mouse control to the user after the user profile has been loaded. It is normally the longest duration out of all the phases of the logon process and is calculated as **Interactive Session duration = Desktop Ready Event Timestamp (EventId 1000 on VDA) - User Profile Loaded Event Timestamp (EventId 2 on VDA)**. Interactive Session has three sub-phases: Pre-userinit, Userinit, and Shell. Hover over the Interactive Session to see a tooltip showing the following:

- subphases
- the time taken for each subphase
- the total cumulative time delay between these subphases
- link to documentation

**Note:**

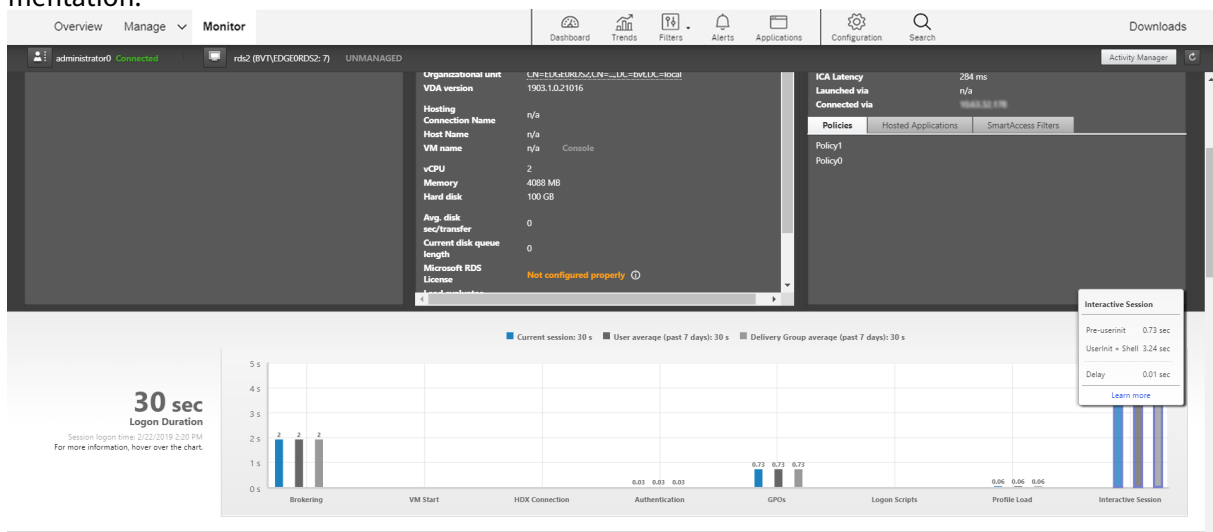
This feature is available on VDAs 1811 and later. If you have launched sessions on sites earlier

than 7.18 and then upgraded to 7.18, a 'Drilldown unavailable due to server error' message is displayed. However, if you have launched sessions after an upgrade, no error message is displayed.

To view the time duration of each subphase, enable Audit process tracking on the VM (VDA). When the Audit process tracking is disabled (default), the time duration of Pre-userinit and the combined time duration of Userinit and Shell are displayed. You can enable Audit process tracking through a Group Policy Object (GPO) as follows:

1. Create a GPO and edit it using the GPO editor.
2. Go to **Computer Configuration > Windows Settings > Security Settings > Local Policies > Audit Policy**.
3. On the right pane, double-click **Audit process tracking**.
4. Select **Success** and click OK.
5. Apply this GPO to the required VDAs or Group.

For more information about Audit process tracking and enabling or disabling it, see [https://docs.microsoft.com/en-us/previous-versions/ms813609\(v=msdn.10\)](https://docs.microsoft.com/en-us/previous-versions/ms813609(v=msdn.10)) in the Microsoft documentation.



Logon Duration panel in the User Details view.

- **Interactive Session – Pre-userinit:** This is the segment of Interactive Session which overlaps with Group Policy Objects and scripts. This subphase can be reduced by optimizing the GPOs and scripts.
- **Interactive Session – Userinit:** When a user logs on to a Windows machine, Winlogon runs userinit.exe. Userinit.exe runs logon scripts, re-establishes network connections, and then starts Explorer.exe, the Windows user interface. This subphase of Interactive Session represents the duration between the start of Userinit.exe to the start of the user interface for the virtual desktop or application.
- **Interactive Session – Shell:** In the previous phase, Userinit starts the initialization of Windows

user interface. The Shell subphase captures the duration between the initialization of the user interface to the time user receives keyboard and mouse control.

- **Delay:** This is the cumulative time delay between the **Pre-userinit and Userinit** subphases and the **Userinit and Shell** subphases.

The total logon time is not an exact sum of these phases. For example, some phases occur in parallel, and in some phases, extra processing occurs that can result in a longer logon duration than the sum. The total logon time does not include the ICA idle time that is the time between the ICA file download and the ICA file launch for an application.

To enable the automatic opening of ICA file upon application launch, configure your browser for automatic ICA file launch upon download of an ICA file. For more information, see [CTX804493](#).

**Note:**

The Logon Duration graph shows the logon phases in seconds. Any duration values below one second are displayed as subsecond values. The values above one second are rounded to the nearest 0.5 second. The graph has been designed to show the highest y-axis value as 200 seconds. Any value greater than 200 seconds is shown with the actual value displayed above the bar.

## Troubleshooting tips

To identify unusual or unexpected values in the graph, compare the amount of time taken in each phase of the current session with the average duration for this user for the last seven days, and the average duration for all users in this delivery group for the last seven days.

Escalate as needed. For example, if the VM startup is slow, the issue might be in the hypervisor, so you can escalate it to the hypervisor administrator. Or, if the brokering time is slow, you can escalate the issue to the site administrator to check the load balancing on the Delivery Controller.

Examine unusual differences, including:

- Missing (current) logon bars
- Major discrepancy between the current duration and this user's average duration. Causes include:
  - A new application was installed.
  - An operating system update occurred.
  - Configuration changes were made.
  - Profile size of the user is high. In this case, the Profile Load is high.
- Major discrepancy between the user's log on numbers (current and average duration) and the delivery group average duration.

If needed, click **Restart** to observe the user's logon process to troubleshoot issues, such as VM Start or Brokering.

## Shadow users

December 16, 2020

Use the shadow user feature to view or work directly on a user's virtual machine or session. You can shadow both Windows or and Linux VDAs. The user must be connected to the machine that you want to shadow. Verify this by checking the machine name listed in the user title bar.

Shadowing is launched in a new tab, update your browser settings to allow pop-ups from the Citrix Cloud URL.

Access the shadowing feature from the **User Details** view. Select the user session, and click **Shadow** in the Activity Manager view or the Session Details panel.

### Shadowing Linux VDAs

Shadowing is available for Linux VDAs Version 7.16 or and later running the RHEL7.3 or Ubuntu Version 16.04 Linux distributions.

#### Note:

- Monitor uses FQDN to connect to the target Linux VDA. Ensure that the Monitor client can resolve the FQDN of the Linux VDA.
- The VDA must have the python-websocketify and x11vnc packages installed.
- noVNC connection to the VDA uses the WebSocket protocol. By default, **ws://** WebSocket protocol is used. For security reasons, Citrix recommends that you use the secure **wss://** protocol. Install SSL certificates on each Monitor client and Linux VDA.

Follow the instructions in [Session Shadowing](#) to configure your VDA for shadowing.

1. After you click **Shadow**, the shadowing connection initializes and a confirmation prompt appears on the user device.
2. Instruct the user to click **Yes** to start the machine or session sharing.
3. The administrator can only view the shadowed session.

### Shadowing Windows VDAs

Windows VDA sessions are shadowed using Windows Remote Assistance. Enable User Windows Remote Assistance feature while installing the VDA. For more information, see [Enable or Disable features](#).

1. After you click **Shadow**, the shadowing connection initializes and a dialog box prompts you to open or save the .msrc incident file.
2. Open the incident file with the Remote Assistance Viewer, if not already selected by default. A confirmation prompt appears on the user device.

3. Instruct the user to click **Yes** to start the machine or session sharing.
4. For more control, ask the user to share keyboard and mouse control.

### **Streamline Microsoft Internet Explorer browsers for shadowing**

Configure your Microsoft Internet Explorer browser to automatically open the downloaded Microsoft Remote Assistance (.msra) file with the Remote Assistance client.

To do this, you must enable the Automatic prompting for file downloads setting in the Group Policy editor:

Computer Configuration > Administrative Templates > Windows Components > Internet Explorer > Internet Control Panel > Security Page > Internet Zone > Automatic prompting for file downloads.

## **Send messages to users**

April 23, 2020

From Monitor, send a message to a user who is connected to one or more machines. For example, use this feature to send immediate notices about administrative actions such as impending desktop maintenance, machine logoffs and restarts, and profile resets.

To send a message to a user, follow these steps:

1. Go to **Monitor > Filters > Machines > All Machines**.
2. Select a machine to which you want to send a message and click **Send Message**.
3. Type your message and click **Send**.

If the message is sent successfully, a confirmation message appears. If the user's machine is connected, the message appears there.

If the message is not sent successfully, an error message appears. Troubleshoot the problem according to the error message. When you have finished, type the subject and message text again and click Try again.

## **Resolve application failures**

September 14, 2020

In the **Activity Manager** view, click the **Applications** tab. You can view all the applications on all machines to which this user has access, including local and hosted applications for the currently connected machine, and the status of each.



**Note:**

If the Applications tab is grayed out, contact an administrator with the permission to enable the tab.

The list includes only those applications that were launched within the session.

For Multi-session OS machines and Single session OS machines, applications are listed for each disconnected session. If the user is not connected, no applications are displayed.

Action	Description
End the application that is not responding	Choose the application that is not responding and click <b>End Application</b> . Once the application is terminated, ask the user to launch it again.
End processes that are not responding	If you have the required permission, click the <b>Processes</b> tab. Select a process that is related to the application or using a high amount of CPU resources or memory, and click <b>End Process</b> . However, if you do not have the required permission to terminate the process, attempting to end a process fails.
Restart the user's machine	For Single session OS machines only, for the selected session, click <b>Restart</b> . Alternatively, from the Machine Details view, use the power controls to restart or shut down the machine. Instruct the user to log on again so that you can recheck the application. For Multi-session OS machines, the restart option is not available. Instead, log off from the user and let the user log on again.
Put the machine into maintenance mode	If the machine's image needs maintenance, such as a patch or other updates, put the machine into maintenance mode. From the Machine Details view, click <b>Details</b> and turn on the maintenance mode option. Escalate to the appropriate administrator.

## Restore desktop connections

June 22, 2020

From Monitor, check the user's connection status for the current machine in the user title bar.

If the desktop connection failed, the error that caused failure is displayed and can help you decide how to troubleshoot.

---

Action	Description
Ensure that the machine is not in maintenance mode	On the User Details page, make sure maintenance mode is turned off.
Restart the user's machine	Select the machine and click <b>Restart</b> . Use this option if the user's machine is unresponsive or unable to connect, such as when the machine is using an unusually high amount of CPU resources, which can make the CPU unusable.

---

## Restore sessions

June 22, 2020

If a session becomes disconnected, it is still active and its applications continue to run, but the user device is no longer communicating with the server.

In the User Details view, troubleshoot session failures in the **Session Details** panel. You can view the details of the current session, indicated by the session ID.

---

Action	Description
End applications or processes that are not responding	Click the <b>Applications</b> tab. Select any application that is not responding and click <b>End Application</b> . Similarly, select any corresponding process that is not responding and click <b>End Process</b> . Also, end processes that are consuming an unusually high amount of memory or CPU resources, which can make the CPU unusable.

---

Action	Description
Disconnect the Windows session	Click <b>Session Control</b> and then select <b>Disconnect</b> . This option is available only for brokered Multi-session OS machines. For non-brokered sessions, the option is disabled.
Log off the user from the session	Click <b>Session Control</b> and then select <b>Log Off</b> .

To test the session, the user can attempt to log back on to it. You can also shadow the user to more closely monitor this session.

## Run HDX channel system reports

August 26, 2021

In the **User Details** view, check the status of the HDX channels on the user's machine in the HDX panel. This panel is available only if the user machine is connected using HDX.

If a message appears indicating that the information is not currently available, wait for one minute for the page to refresh, or select the **Refresh** button. HDX data takes a little longer to update than other data.

Click an error or warning icon for more information.

Tip:

You can view information about other channels in the same dialog box by clicking the left and right arrows in the left corner of the title bar.

HDX channel system reports are used mainly by Citrix Support to troubleshoot further. To do this, in the HDX panel, click **Download System Report**.

## Reset a user profile

September 14, 2020

**Caution:**

When a profile is reset, although the user's folders and files are saved and copied to the new profile, most user profile data is deleted (for example, the registry is reset and application settings

might be deleted).

1. From Monitor, search for the user whose profile you want to reset and select this user's session.
2. Click **Reset Profile**.
3. Instruct the user to log off from all sessions.
4. Instruct the user to log back on. The folders and files that were saved from the user's profile are copied to the new profile.

**Important:**

If the user has profiles on multiple platforms (such as Windows 8 and Windows 7), instruct the user to log back on first to the same desktop or app that the user reported as a problem. This ensures that the correct profile is reset. For a Citrix user profile, the profile is already reset by the time the user's desktop appears. For a Microsoft roaming profile, the folder restoration might still be in progress for a brief time. The user must stay logged on until the restoration is complete.

The preceding steps assume you are using Citrix Virtual Desktops (Desktop VDA). If you are using Citrix Virtual Desktops (Server VDA) you need to be logged on to perform the profile reset. The user then needs to log off, and log back on to complete the profile reset.

If the profile is not successfully reset (for example, the user cannot successfully log back on to the machine or some of the files are missing), you must manually restore the original profile.

The folders (and their files) from the user's profile are saved and copied to the new profile. They are copied in the listed order:

- Desktop
- Cookies
- Favorites
- Documents
- Pictures
- Music
- Videos

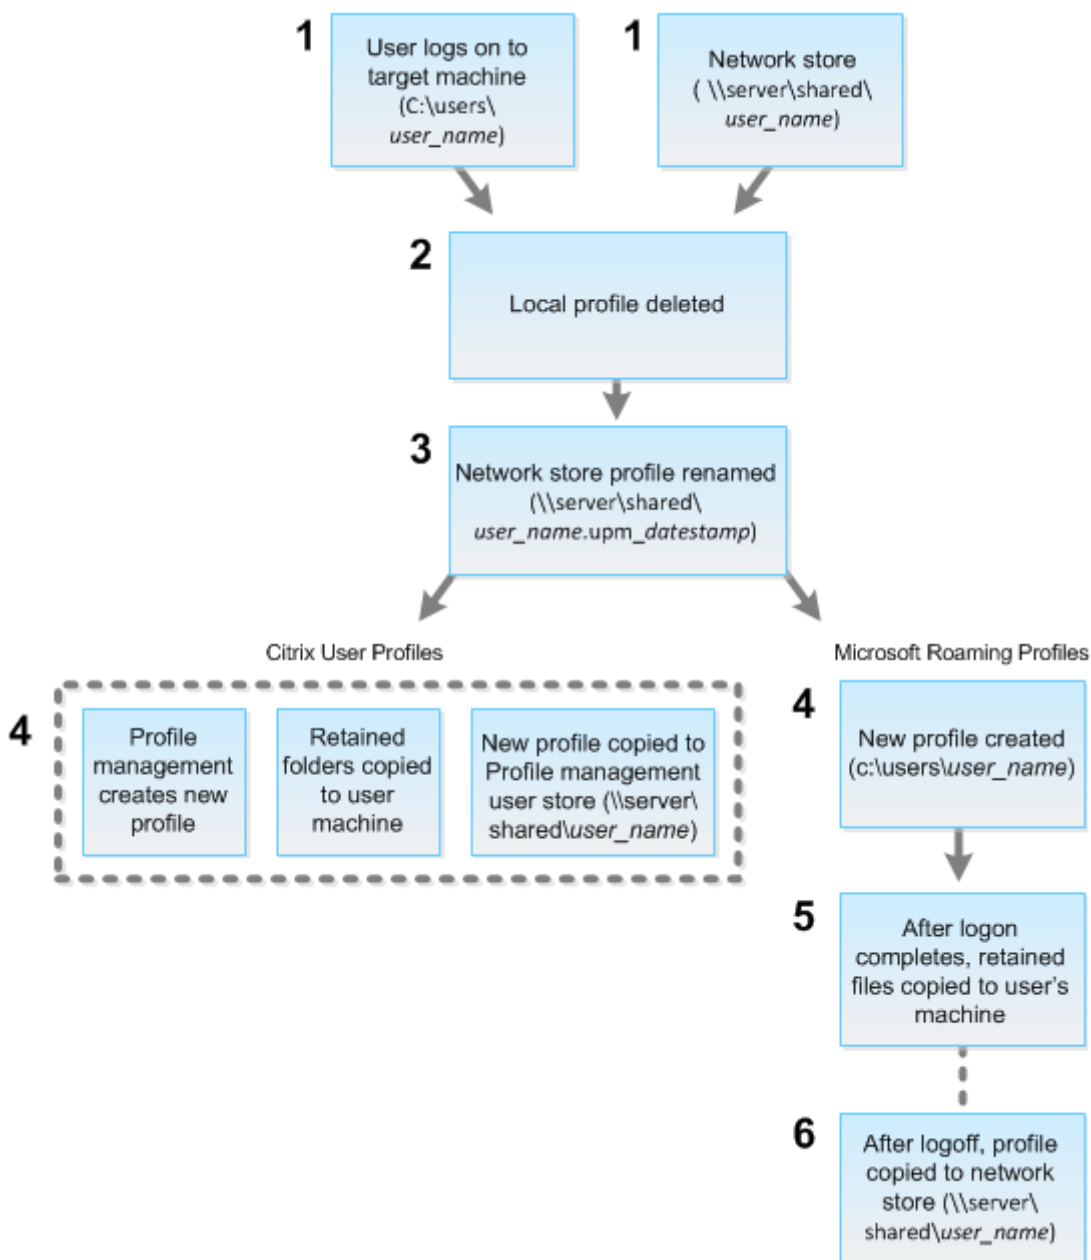
**Note:**

In Windows 8 and later, cookies are not copied when profiles are reset.

## How reset profiles are processed

Any Citrix user profile or Microsoft roaming profile can be reset. After the user logs off and you select the reset command (either in Monitor or using the PowerShell SDK), Monitor first identifies the user profile in use and issues an appropriate reset command. Monitor receives the information through Profile Management, including information about the profile size, type, and logon timings.

This diagram illustrates the process following the user log on, when a user profile is reset.



The reset command issued by Monitor specifies the profile type. The Profile Management service then attempts to reset a profile of that type and looks for the appropriate network share (user store). If Profile Management processes the user, but receives a roaming profile command, it is rejected (or the opposite way).

1. If a local profile is present, it is deleted.
2. The network profile is renamed.
3. The next action depends on whether the profile being reset is a Citrix user profile or a Microsoft roaming profile.

For Citrix user profiles, the new profile is created using the Profile Management import rules, and the folders are copied back to the network profile, and the user can log on normally. If a roaming profile is used for the reset, any registry settings in the roaming profile are preserved in the reset profile. You can configure Profile management so that a template profile overrides the roaming profile, if necessary.

For Microsoft roaming profiles, Windows creates a new profile, and when the user logs on, the folders are copied back to the user device. When the user logs off again, the new profile is copied to the network store.

### To manually restore a profile after a failed reset

1. Instruct the user to log off from all sessions.
2. Delete the local profile if one exists.
3. Locate the archived folder on the network share that contains the date and time appended to the folder name, the folder with a .upm\_datestamp extension.
4. Delete the current profile name. That is, the one without the upm\_datestamp extension.
5. Rename the archived folder using the original profile name. That is, remove the date and time extension. You have returned the profile to its original, pre-reset state.

## Feature compatibility matrix

May 24, 2021

Citrix Monitor supports three Citrix Virtual Apps and Desktops service editions. They are **Premium**, **Citrix Virtual Apps Advanced Service**, and **Citrix Virtual Apps and Desktops Advanced Service**. Specific Citrix Monitor features, VDA versions, dependent components, and their respective license editions are listed in the following table.

Feature	Dependencies - min version required	Premium	Citrix Virtual Apps Advanced Service	Citrix Virtual Apps and Desktops Advanced Service
<a href="#">Session Auto Reconnect</a>	VDA 1906	Yes	Yes	Yes
<a href="#">Session Startup Duration</a>	VDA 1903	Yes	Yes	Yes

<b>Feature</b>	<b>Dependencies - min version required</b>	<b>Premium</b>	<b>Citrix Virtual Apps Advanced Service</b>	<b>Citrix Virtual Apps and Desktops Advanced Service</b>
Desktop Probing	Citrix Probe Agent 1903	Yes	Yes	Yes
Citrix Profile Management Duration in Profile Load	VDA 1903	Yes	Yes	Yes
Profile Drilldown	VDA 1811	Yes	Yes	Yes
Hypervisor Alerts Monitoring	None	Yes	No	No
Application Probing	Citrix Application Probe Agent 1811	Yes	Yes	Yes
Microsoft RDS License Health	VDA 7.16	Yes	Yes	Yes
Key RTOP Data Display	VDA 1808	Yes	Yes	Yes
Export of Filters Data	None	Yes	Yes	Yes
Interactive Session Drill Down	VDA 1808	Yes	Yes	Yes
GPO Drill Down	VDA 1808	Yes	Yes	Yes
Machine Historical Data Available Using OData API	None	Yes	Yes	Yes
Smart Alert Policies	None	Yes	No	No

<b>Feature</b>	<b>Dependencies - min version required</b>	<b>Premium</b>	<b>Citrix Virtual Apps Advanced Service</b>	<b>Citrix Virtual Apps and Desktops Advanced Service</b>
<a href="#">Health Assistant Link</a>	None	Yes	Yes	Yes
<a href="#">Interactive Session Drill-down</a>	None	Yes	Yes	Yes
<a href="#">Application Analytics</a>	VDA 7.15	Yes	Yes	Yes
<a href="#">OData API V.4</a>	None	Yes	Yes	Yes
<a href="#">Shadow Linux VDA Users</a>	VDA 7.16	Yes	Yes	Yes
<a href="#">Machine Console Access</a>	None	Yes	Yes	Yes
<a href="#">Application Failure Monitoring</a>	VDA 7.15	Yes	Yes	Yes
<a href="#">Application-centric Troubleshooting</a>	VDA 7.13	Yes	Yes	Yes
<a href="#">Disk Monitoring</a>	VDA 7.14	Yes	Yes	Yes
<a href="#">GPU Monitoring</a>	VDA 7.14	Yes	Yes	Yes
<a href="#">Transport Protocol on Session Details Panel</a>	VDA 7.13	Yes	Yes	Yes
<a href="#">User-friendly Connection and Machine Failure Descriptions</a>	VDA 7.x	Yes	Yes	Yes
<a href="#">Historical Data Retention</a>	VDA 7.x	Yes	No	No



<b>Feature</b>	<b>Dependencies - min version required</b>	<b>Premium</b>	<b>Citrix Virtual Apps Advanced Service</b>	<b>Citrix Virtual Apps and Desktops Advanced Service</b>
Custom Reporting	VDA 7.x	Yes	No	No
Resource Utilization Reporting	VDA 7.11	Yes	Yes	Yes
Alerting Extended for CPU, Memory and ICA RTT Conditions	VDA 7.11	Yes	No	No
Export Report Improvements	VDA 7.x	Yes	Yes	Yes
Logon Duration Breakdown	VDA 7.x	Yes	Yes	Yes
Proactive Monitoring and Alerting	VDA 7.x	Yes	No	No
Hosted Applications Usage	VDA 7.x	Yes	No	No
Single-session and Multi-session OS Usage	VDA 7.x	Yes	No	No
Support for Framehawk Virtual Channel	VDA 7.6	Yes	Yes	Yes

## Delegated administration and monitoring

September 15, 2020

Delegated administration uses three concepts: administrators, roles, and scopes. Permissions are based on an administrator's role and the scope of this role. For example, an administrator might be assigned a Help Desk administrator role where the scope involves responsibility for end-users at one site only.

Administrative permissions determine the monitoring interface presented to administrators and the tasks they can perform. Permissions determine:

- The views the administrator can access, collectively referred to as a view.
- The desktops, machines, and sessions that the administrator can view and interact with.
- The commands the administrator can perform, such as shadowing a user's session or enabling maintenance mode.

Monitoring now supports delegated administrator roles that allow you to assign custom defined or built-in roles to administrators. The role determines the available permissions and hence, how an administrator uses monitoring. You can also define the scope applicable for those roles. The scope defines the objects for which the role is applicable.


For information about creating delegated administrators, see the main [Delegated administration](#) article.

The built-in roles and permissions determine how administrators use **Monitor**:

Administrator Role	Permissions in Monitor
Full Administrator	Full access to all views and can perform all commands, including shadowing a user's session, enabling maintenance mode, and exporting trends data.
Delivery group Administrator	Full access to all views and can perform all commands, including shadowing a user's session, enabling maintenance mode, and exporting trends data.

Administrator Role	Permissions in Monitor
Read Only Administrator	Can access all views and see all objects in specified scopes in addition to global information. Can download reports from HDX channels and can export Trends data using the Export option in the Trends view. Cannot perform any other commands or change anything in the views.
Help Desk Administrator	Can access only the Help Desk and User Details views and can view only objects that the administrator is delegated to manage. Can shadow a user's session and perform commands for that user. Can perform maintenance mode operations. Can use power control options for Single session OS Machines. Cannot access the Dashboard, Trends, Alerts, or Filters views. Cannot use power control options for Multi-session OS machines.
Machine catalog Administrator	Can access only the Machine Details page (Machine-based search).
Host Administrator	No access. This administrator is not supported for Monitor and cannot view data.
Probe Agent Administrator	Read-only access to Applications page, cannot access any other view. Meant to run the Citrix Probe Agent on endpoint machines.
Monitoring Full Administrator	Has full access to all views and commands in the <b>Monitor</b> tab.
Session Administrator	Can view Delivery Groups and manage their associated sessions and machines on the <b>Filters</b> page of the <b>Monitor</b> tab.

To assign a role (built-in or custom) to a user, from the Citrix Cloud menu, go to **Identity and Access Management > Administrators**. Here, when you add or edit the access of an administrator, you can select **Custom Access** and one of the listed roles.



You can define custom roles and scopes in Manage → Citrix Studio → Configuration → Administrators.

The built-in roles and custom roles are listed for selection with custom scope.



- Cloud Administrator, All
- Delivery Group Administrator, All
- Delivery Group Administrator, rds1DGAndCatalog
- Delivery Group Administrator, vdaDGOnly
- Full Monitor Administrator, All - Access to 'Monitor' tab only
- Full Monitor Administrator, rds1DGAndCatalog - Access to 'Monitor' tab only
- Full Monitor Administrator, vdaDGOnly - Access to 'Monitor' tab only
- Help Desk Administrator, All - Access to 'Monitor' tab only
- Help Desk Administrator, rds1DGAndCatalog - Access to 'Monitor' tab only
- Help Desk Administrator, vdaDGOnly - Access to 'Monitor' tab only
- Host Administrator, All
- Host Administrator, rds1DGAndCatalog
- Host Administrator, vdaDGOnly
- Machine Catalog Administrator, All
- Machine Catalog Administrator, rds1DGAndCatalog
- Machine Catalog Administrator, vdaDGOnly
- Probe Agent Administrator, All
- Probe Agent Administrator, rds1DGAndCatalog
- Probe Agent Administrator, vdaDGOnly
- Read Only Administrator, All
- Read Only Administrator, rds1DGAndCatalog
- Read Only Administrator, vdaDGOnly
- TrendsFiltersAndUD, All
- TrendsFiltersAndUD, rds1DGAndCatalog
- TrendsFiltersAndUD, vdaDGOnly

## Data granularity and retention

February 16, 2021

### Aggregation of data values

The Monitor Service collects various data, including user session usage, user logon performance details, session load balancing details, and connection and machine failure information. Data is aggregated differently depending on its category. Understanding the aggregation of data values presented using the OData Method APIs is critical to interpreting the data. For example:

- Connected Sessions and Machine Failures occur over a period. Therefore, they are exposed as maximums over a time period.
- Logon Duration is a measure of the length of time, therefore is exposed as an average over a time period.
- Logon Count and Connection Failures are counts of occurrences over a period, therefore are exposed as sums over a time period.

### Concurrent data evaluation

Your sessions must be overlapping to be considered concurrent. However, when the time interval is 1 minute, all sessions in that minute (whether they overlap) are considered concurrent. The size of the interval is so small that the performance overhead involved in calculating the precision is not worth the value added. If the sessions occur in the same hour, but not in the same minute, they are not considered to overlap.

### Correlation of summary tables with raw data

The data model represents metrics in two different ways:

- The summary tables represent aggregate views of the metrics in per minute, hour, and day time granularities.
- The raw data represents individual events or current state tracked in the session, connection, application, and other objects.

When attempting to correlate data across API calls or within the data model itself, it is important to understand the following concepts and limitations:

- **No summary data for partial intervals.** Metrics summaries are designed to meet the needs of historical trends over long periods of time. These metrics are aggregated into the summary table for complete intervals. There is no summary data for a partial interval at the beginning (oldest available data) of the data collection nor at the end. When viewing aggregations of a

day (Interval=1440), this means that the first and most recent incomplete days have no data. Although raw data might exist for those partial intervals, it is never summarized. Pull the min and max SummaryDate from a particular summary table to determine the earliest and latest aggregate interval for a particular data granularity. The SummaryDate column represents the start of the interval. The Granularity column represents the length of the interval for the aggregate data.

- **Correlating by time.** Metrics are aggregated into the summary table for complete intervals as described in the preceding section. They can be used for historical trends, but raw events might be more current in the state than what has been summarized for trend analysis. Any time-based comparison of summary to raw data must take into account that there is no summary data for partial intervals that might occur or for the beginning and ending of the time period.
- **Missed and latent events.** Metrics that are aggregated into the summary table might be slightly inaccurate if events are missed or latent to the aggregation period. Although the Monitor Service attempts to maintain an accurate current state, it does not go back in time to recompute aggregation in the summary tables for missed or latent events.
- **Connection High Availability.** During connection HA, there are gaps in the summary data counts of current connections, but the session instances are still running in the raw data.
- **Data retention periods.** Data in the summary tables is retained on a different grooming schedule from the schedule for raw event data. Data might be missing because it has been groomed away from summary or raw tables. Retention periods might also differ for different granularities of summary data. Lower granularity data (minutes) is groomed more quickly than higher granularity data (days). If data is missing from one granularity due to grooming, it might be found in a higher granularity. Since the API calls only return the specific granularity requested, receiving no data for one granularity does not mean that the data doesn't exist for a higher granularity for the same time period.
- **Time zones.** Metrics are stored with UTC time stamps. Summary tables are aggregated on hourly time zone boundaries. For time zones that don't fall on hourly boundaries, there might be some discrepancy as to where data is aggregated.

## Granularity and retention

The granularity of aggregated data retrieved by Monitor is a function of the time (T) span requested. The rules are as follows:

- $0 < T \leq 30$  days use per-hour granularity
- $T > 31$  days use per-day granularity

Requested data that does not come from aggregated data comes from the raw Session and Connection information. This data tends to grow fast, and therefore has its own grooming setting. Grooming ensures that only relevant data is kept long term. This ensures better performance while maintaining the granularity required for reporting.

Citrix Virtual Apps and Desktops service supports historical data retention only for 90 days. Hence, one-year trends and reports in Monitor show the last 90 days of data.

	Setting name	Affected grooming	Retention days for Premium	Retention days for Advanced
1	GroomSessionsRe	Session and Connection records retention after Session termination	90	31
2	GroomFailuresReten	MinDaysFailureLog and Connection-FailureLog records	90	31
3	GroomLoadIndexe	LoadIndex records	90	31
4	GroomDeletedReten	MinDays, Catalog, DesktopGroup, and Hypervisor entities that have a LifecycleState of 'Deleted'. This also deletes any related Session, SessionDetail, Summary, Failure, or LoadIndex records.	90	31



	Setting name	Affected grooming	Retention days for Premium	Retention days for Advanced
5	GroomSummaries	DesktopGroupSun FailureLogSummary, and LoadIndexSummary records. Aggregated data - daily granularity.	90	31
6	GroomMachineHotfixes	Hotfixes applied to the VDA and Controller machines	90	31
7	GroomHourlyRetention	Aggregated data - hourly granularity	32	31
8	GroomApplicationInstanceRetention	Application Instance history	90	Not applicable
9	GroomNotification	Notification Log records	90	Not applicable
10	GroomResourceUsageRawDataRetention	Resource utilization data - raw data	3	3
11	GroomResourceUsageSummaryDataRetention	Resource utilization summary data - hour granularity	30	30
12	GroomResourceUsageDayDataRetention	Resource utilization summary data - day granularity	30	31
13	GroomProcessUsageRawDataRetention	Process utilization data - raw data	1	1

	Setting name	Affected grooming	Retention days for Premium	Retention days for Advanced
14	GroomProcessUsageHourlyDataRetentionDays	Process utilization data - hour granularity	7	7
15	GroomProcessUsageDailyDataRetentionDays	Process utilization data - day granularity	30	30
16	GroomSessionMetricsDataRetentionDays	Session metrics data	1	1
17	GroomMachineMetricsDataRetentionDays	Machine metrics data	3	3
18	GroomMachineMetricsSummaryDataRetentionDays	Machine metrics summary data	31	31
19	GroomApplicationErrorDataRetentionDays	Application error data	1	1
20	GroomApplicationFailureDataRetentionDays	Application failure data	1	1

**Caution:**

You cannot modify the values on the Monitor Service database. To change any settings in the database, contact Citrix Support.

Retaining data for long periods has the following implications on table sizes:

- Hourly data.** If hourly data is allowed to stay in the database for up to two years, a site of 1000 delivery groups can cause the database to grow as follows:
 

1000 delivery groups x 24 hours/day x 365 days/year x 2 years = 17,520,000 rows of data. The performance impact of such a large amount of data in the aggregation tables is significant. Given that the dashboard data is drawn from this table, the requirements on the database server might be large. Excessively large amounts of data can have a dramatic impact on performance.
- Session and event data.** This is the data that is collected every time a session is started and a connection/reconnection is made. For a large site (100 K users), this data grows fast. For example, two years' worth of these tables would gather more than a TB of data, requiring a high-end enterprise-level database.

## Citrix Virtual Apps and Desktops service for Citrix Service Providers

July 7, 2021

This article describes how **Citrix Service Providers (CSP)** can set up the Virtual Apps and Desktops service for tenant customers in Citrix Cloud. For an overview of the features available for Citrix Partners, see [Citrix Cloud for Partners](#).

### Requirements

- You are a [Citrix Service Provider partner](#).
- You have a Citrix Cloud account.
- You have a subscription to the Citrix Virtual Apps and Desktops service.

### Limitations and known issues

#### Limitations

- Tenant name changes take up to 24 hours to apply across all interfaces.
- When creating a new tenant, the email address must be unique.
- Filtering in **Manage > Full Configuration** by scope (similar to Monitor) is not available. To see the resources attached to a scope, select **Administrators** in the left pane. On the **Scopes** tab, select the scope and then select **Edit Scope** in the Action pane.

#### Known issues

- The customer scope name in the **Manage > Full Configuration** display shows an internal ID rather than the customer name. You can change the scope name to a friendly name. Select **Administrators** in the left pane. On the **Scopes** tab, select the scope and then select **Edit Scope** in the Action bar.
- After scopes are assigned to a resource, you cannot use the management console to remove or unassign them. Those tasks are supported only through PowerShell.
- **Manage > Full Configuration** does not enforce scopes. You are responsible for selecting the appropriate scope when creating machine catalogs, delivery groups, and application groups.
- When more than 15 scopes are created (auto-created and custom), the Citrix Cloud custom access information for an administrator (**Identity and Access Management > Administrators**) does not display correctly. Workaround: Limit scopes to 15 or fewer.
- After adding the Citrix Virtual Apps and Desktops service to a customer:
  - You cannot remove it from a customer.
  - You cannot remove the link between the customer and the CSP.

## Add a customer

1. Sign in to Citrix Cloud with your CSP credentials. Select **Customers** in the upper left menu.
2. From the Customer Dashboard, select **Invite or Add**. Provide the requested information.
3. If the customer does not have a Citrix Cloud account, adding the customer creates a customer account. Adding the customer also automatically adds you as a full access administrator of that customer's account.
4. If the customer has a Citrix Cloud account:
  - a) A Citrix Cloud URL displays, which you copy and send to the customer. For details of this process, see [Inviting a customer to connect](#).
  - b) The customer must add you as a full access administrator to their account. See [Add administrators to a Citrix Cloud account](#).

You can add more administrators later and control which customers they can see on the **Manage** and **Monitor** consoles.

## Add the Citrix Virtual Apps and Desktops service to a customer

1. Sign in to Citrix Cloud with your CSP credentials. Select **Customers** in the upper left menu.
2. From the Customer Dashboard, in the ellipsis menu for the customer, select **Add Service**.
3. In **Select a Service to Add**, select **Virtual Apps and Desktops**.
4. Select **Continue**.

After you complete this procedure, the customer is onboarded to your Citrix Virtual Apps and Desktops service subscription.

When the onboarding completes, a new customer scope is created automatically in the Citrix Virtual Apps and Desktops service. The scope is visible in the **Manage > Full Configuration** display. This scope is unique to that customer. You can [rename the scope](#), but you cannot delete it.

Use this scope to tailor access for other administrators. For example, let's say you have 10 customers and two administrators. Using the unique scope, you can restrict one administrator's access to only three of the customers. The other administrator can access one of those three customers, plus two other customers. For details, see [Control administrator access to customers](#).

## Set up a resource location

A resource location holds the machines that deliver apps and desktops for your customers, and infrastructure components such as Citrix Cloud Connectors. For details, see [Connect to Citrix Cloud](#).

## Set up catalogs and groups to deliver apps and desktops

A catalog is a group of identical virtual machines. When you create a catalog, an image is used (with other settings) as a template for creating the machines. For details, see [Create machine catalogs](#).

A delivery group is a collection of machines selected from one or more machine catalogs. The delivery group specifies which users can use those machines, plus the applications or desktops available to those users. For details, see [Create delivery groups](#).

Application groups let you manage collections of applications. You can create application groups for applications shared across different delivery groups or used by a subset of users within delivery groups. For details, see [Create application groups](#).

When configuring groups, be sure that:

- The delivery group's scope is a subset of the machine catalog's scope. For example, assume the catalog's scope is A and B. The delivery group's scope can be either A or B, or A and B.
- The application group's scope is a subset of the delivery group's scope. For example, assume the delivery groups associated with an application group have scope A and B. The application group's scope can be either A or B, or A and B.

## Federated domains

Federated domains enable customer users to use credentials from a domain attached to your resource location to sign in to their workspace. This allows you to provide dedicated workspaces to your customers that customer users can access using a custom workspace URL (for example, customer.cloud.com), while the resource location is still on your Citrix Cloud account. You can provide dedicated workspaces alongside the shared workspace that customers can access using your CSP workspace URL (for example, csppartner.cloud.com).

To enable customers to access their dedicated workspace, you add them to the appropriate domains that you manage. After configuring the workspace through [Workspace Configuration](#), customers' users can sign in to their workspace and access the apps and desktops that you've made available.

### Add a customer to a domain

1. Sign in to Citrix Cloud with your CSP credentials. Select **Customers** in the upper left menu.
2. From the Customer Dashboard, select **Identity and Access Management** in the upper left menu.
3. On the **Domains** tab, select **Manage Federated Domain** in the domain's ellipsis menu.
4. On the **Manage Federated Domain** card, in the **Available customers** column, select a customer you want to add to the domain. Select the plus sign next to the customer name. The selected customer now appears in the **Federated customers** column. Repeat to add other customers. When you're done, select **Apply**.

### Remove a customer from a domain

When you remove a customer from a domain that you manage, the customer's users can no longer access their workspaces using credentials from your domain.

1. From the Citrix Cloud menu, select **Identity and Access Management**, then select **Domains**.
2. Locate the domain you want to manage and select the ellipsis button. Select **Manage Federated Domain**.
3. From the list of federated customers, locate or search for the customers you want to remove and select the X button. Select **Remove all** to remove all the customers in the list from the domain. The selected customers move to the list of available customers.
4. Select **Apply**.
5. Review the customers you selected and select **Remove Customers**.

### Control administrator access to customers

You can control administrator access to customers by using the unique scope that was created when you added the Citrix Virtual Apps and Desktops service to the customer. You can configure access when you add an administrator or later.

To learn about restricting access using roles and scopes in the Citrix Virtual Apps and Desktops service, see [Delegated administration](#).

### Add an administrator with restricted access

1. Sign in to Citrix Cloud with your CSP credentials. Select **Customers** in the upper left menu.
2. From the Customer Dashboard, select **Identity and Access Management** in the upper left menu.
3. On the **Administrators** tab, select **Add Administrators From**, and then select **Citrix Identity**.
4. Type the email address of the person you're adding as an administrator, and then select **Invite**.
5. Configure the appropriate access permissions for the administrator. Citrix recommends selecting **Custom access**, unless you want the administrator to have management control of Citrix Cloud and all of the subscribed services.
6. After selecting **Custom access**, select one or more role and scope pairs for the Virtual Apps and Desktops service, as needed. Be sure to enable only entries that contain the unique scope that was created for the customer.
7. When you're done selecting role and scope pairs, select **Send Invite**.

When the administrator accepts the invitation, they have the access that you assigned.

### Edit delegated administration permissions for administrators

1. Sign in to Citrix Cloud with your CSP credentials. Select **Customers** in the upper left menu.

2. From the Customer dashboard, select **Identity and Access Management** in the upper left menu.
3. On the **Administrators** tab, select **Edit Access** from the ellipsis menu for the administrator.
4. Select and clear role and scope pairs for the Virtual Apps and Desktops service, as needed. Be sure to enable only entries that contain the unique scope that was created for the customer.
5. Select **Save**.

### **View customer administrators and their assigned roles and scopes**

1. Sign in to Citrix Cloud with your CSP credentials. Select **Customers** in the upper left menu.
2. From the Customer Dashboard, select **My Services > Virtual Apps and Desktops** in the upper left menu.
3. In the Citrix Virtual Apps and Desktops service, select **Manage > Full Configuration**.
4. Select **Administrators** in the left pane.

Information is available on three tabs:

- The **Administrators** tab lists the administrators that have been created, plus their roles and scopes.
- The **Roles** tab lists all roles. To view role details, select the role in the middle pane. The lower portion of that pane lists the object types and associated permissions for the role. Select the **Administrators** tab in the lower pane to display a list of administrators who currently have this role.
- The **Scopes** tab lists all the scopes, including those generated for customers of Citrix partners.

### **Configure workspaces**

The customer has their own workspace with a unique `customer.cloud.com` URL. This is where the customer's users access their published apps and desktops.

The workspace URL is displayed in two places:

- From the Customer dashboard, select **Workspace Configuration** from the menu in the upper left menu.
- From the Citrix Virtual Apps and Desktops service **Welcome** page (the **Overview** tab), the workspace URL appears at the bottom of the page.

You can change access and authentication to a workspace. You can also customize the workspace appearance and preferences. For details, see the following articles:

- [Configure workspaces](#)
- [Secure workspaces](#)

## Monitor a customer's service

The **Monitor** dashboard in a CSP environment is essentially the same as a non-CSP environment. See [Monitor](#) for details.

By default, the **Monitor** dashboard displays information about all customers. To display information about one customer, use **Select Customer**.

Keep in mind that the ability to see Monitor displays for a customer is controlled by the administrator's configured access. The access must include a role and scope pair that includes the customer's unique scope.

If you used built-in roles to configure access: The built-in roles control whether the administrator can see the **Manage** and **Monitor** displays. If you select only role and customer-scope pairs that do not include **Monitor** tab visibility, that administrator won't see the **Monitor** tab for any selected customers. For example, if you give an administrator only **Read Only Administrator,customerABC** access, that administrator won't see the **Monitor** tab for customer ABC, because read only administrators don't have access to Monitor displays.

## Citrix Gateway service

December 9, 2020

Citrix Gateway provides users with secure access to Citrix Virtual Apps and Desktops applications.

The Citrix Gateway service enables secure, remote access to those applications, without having to deploy Citrix Gateway in the DMZ or reconfigure your firewall. The infrastructure overhead of using Citrix Gateway moves to Citrix Cloud.

For more information about the Citrix Gateway service, see the [product documentation](#). That content includes how to [enable the service](#) and [known issues](#) for the version you're using.

Citrix ADC is an application delivery controller that analyses application-specific traffic to distribute, optimize, and secure Layer 4-Layer 7 (L4-L7) network traffic intelligently for web applications. The Citrix ADC VPX virtual appliance can be hosted on various virtualization and cloud platforms. For details, see [Deploy a Citrix ADC VPX instance](#).

## SDKs and APIs

July 7, 2021



## Citrix Virtual Apps and Desktops Remote PowerShell SDK

The Remote PowerShell SDK automates complex and repetitive tasks. It provides the mechanism to set up and manage the Citrix Virtual Apps and Desktops service environment without using the Manage user interfaces.

- Cmdlet details are provided in [Citrix Virtual Apps and Desktops service SDK](#).
- The supported snap-ins are listed in Support and limitations. That section also lists the cmdlets that are disabled in this SDK.

### How this SDK differs from the SDK for customer-managed deployments

In a Citrix Virtual Apps and Desktops deployment that is installed and managed by customer administrators, those administrators run cmdlets and scripts in a site containing VDAs and Delivery Controllers within a common domain structure. In contrast, the Citrix Virtual Apps and Desktops service splits the VDAs and Controllers into a resource location and the control plane, respectively. This split means that the original Citrix Virtual Apps and Desktops PowerShell SDK does not work in a Citrix Virtual Apps and Desktops service environment. It cannot cross the secure boundary from the resource location to the control plane.

The solution is the Citrix Virtual Apps and Desktops Remote PowerShell SDK. When running in the resource location, the Remote PowerShell SDK accesses the control plane as if it is local. This provides the same functionality as a single Citrix Virtual Apps and Desktops site. There is only the lowest non-visible communication layer, enhanced to work either in a single local site or in the cloud environment. The cmdlets are the same, and most existing scripts remain unchanged.

The `Get-XdAuthentication` cmdlet provides the authorization to cross the secure resource location to control plane boundary. By default, `Get-XdAuthentication` prompts users for CAS credentials, and must be done once per PowerShell session. Alternatively, the user can define an authentication profile using an API access Secure Client, created in the Citrix Cloud console. In both cases, the security information persists for use in subsequent PowerShell SDK calls. If this cmdlet is not explicitly run, it is called by the first PowerShell SDK cmdlet.

### Install and use the Remote PowerShell SDK

Requirements:

**Note:**

Do not install the Remote PowerShell SDK on a Citrix Cloud Connector machine. It can be installed on any domain-joined machine within the same resource location.

Citrix recommends that you do not run this SDK's cmdlets on Cloud Connectors. The SDK's operation does not involve the Cloud Connectors.

If you also have an on-premises Citrix Virtual Apps and Desktops deployment (in addition to the Virtual Apps and Desktops service deployment), do not install the Remote PowerShell SDK on an on-premises Delivery Controller machine.

- Ensure that PowerShell 3.0 or later is available on the machine.
- The SDK installer downloads and installs .NET Framework 4.8 if it (or a later supported version) is not already installed.
- If the machine already has the Citrix Virtual Apps and Desktops SDK installed, remove that SDK (from Windows Programs and Features) before installing the Remote PowerShell SDK.

To install the Remote PowerShell SDK:

1. From [the download page](#), download the Virtual Apps and Desktops Remote PowerShell SDK.
2. Install and run the SDK.

Installation logs are created in %TEMP%\CitrixLogs\CitrixPoshSdk. Logs can help resolve installation issues.

Run the SDK on a domain-joined computer within that resource location:

- Open a PowerShell command prompt. You do not need to run as an administrator.
- Add the Citrix snap-ins: `asnpx citrix.*`.
- You can explicitly authenticate by using the `Get-XdAuthentication` cmdlet. Or, run your first Citrix Virtual Apps and Desktops PowerShell SDK command, which prompts you for the same authentication as `Get-XdAuthentication`.
- To bypass the authentication prompt, you can use the `Set-XdCredentials` cmdlet to create a default authentication profile, using a Secure Client created in the Citrix Cloud console.
- Continue running PowerShell SDK cmdlets or PowerShell SDK automation scripts. See an example.

To uninstall the Remote PowerShell SDK, from the Windows feature for removing or changing programs, select **Citrix Virtual Apps and Desktops Remote PowerShell SDK**. Right-click and select **Uninstall**. Follow the dialog.

### Example activities

Common activities include setting up machine catalogs, applications, and users. A sample script is shown below.

```
1     $users = "xd.local\Domain Users"
2
3     $TSVDACatalogName = "TSVDA"
4
5     $TSVDADGName = "TSVDA"
6
```

```
7   $TSVDAMachineName = "xd\ds-tsvda2"
8
9   #Create TSVDA Catalog
10
11  $brokerUsers = New-BrokerUser -Name $users
12
13  $catalog = New-BrokerCatalog -Name $TSVDACatalogName -
    AllocationType "Random" -Description $TSVDACatalogName -
    PersistUserChanges "OnLocal" -ProvisioningType "Manual" -
    SessionSupport "MultiSession" -MachinesArePhysical $true
14
15  #Add TSVDA Machine to Catalog
16
17  $BrokeredMachine = New-BrokerMachine -MachineName $TSVDAMachineName
    -CatalogUid $catalog.uid
18
19  #Create new desktops & applications delivery group
20
21  $dg = New-BrokerDesktopGroup -Name $TSVDADGName -PublishedName
    $TSVDADGName -DesktopKind "Shared" -SessionSupport "MultiSession"
    -DeliveryType DesktopsAndApps -Description $TSVDADGName
22
23  #Create notepad application
24
25  New-BrokerApplication -ApplicationType HostedOnDesktop -Name "
    Notepad" -CommandLineExecutable "notepad.exe" -DesktopGroup $dg
26
27  #Assign users to desktops and applications
28
29  New-BrokerEntitlementPolicyRule -Name $TSVDADGName -DesktopGroupUid
    $dg.Uid -IncludedUsers $brokerUsers -description $TSVDADGName
30
31  New-BrokerAccessPolicyRule -Name $TSVDADGName -
    IncludedUserFilterEnabled $true -IncludedUsers $brokerUsers -
    DesktopGroupUid $dg.Uid -AllowedProtocols @("HDX","RDP")
32
33  New-BrokerAppEntitlementPolicyRule -Name $TSVDADGName -
    DesktopGroupUid $dg.Uid -IncludedUsers $brokerUsers -description
    $TSVDADGName
34
35  #Add machine to delivery group
36
37  Add-BrokerMachine -MachineName $TSVDAMachineName -DesktopGroup $dg
38 <!--NeedCopy-->
```

## Support and limitations

The following Citrix Virtual Apps and Desktops PowerShell snap-ins are supported in this release:

- Broker
- Active Directory (AD) Identity
- Machine creation
- Configuration
- Configuration logging
- Host
- Delegated administration
- Analytics

For details about cmdlets in those snap-ins, see [Citrix Virtual Apps and Desktops SDK](#).

After authentication, remote access remains valid in the current PowerShell session for 24 hours. After that time, you must enter your credentials.

The Remote PowerShell SDK must be run on a computer within the resource location.

The following cmdlets are disabled in remote operations to maintain the integrity and security of the Citrix Cloud control plane.

### **Citrix.ADIIdentity.Admin.V2:**

- Copy-AcctIdentityPool
- Get-AcctDBConnection
- Get-AcctDBSchema
- Get-AcctDBVersionChangeScript
- Get-AcctInstalledDBVersion
- Remove-AcctServiceMetadata
- Reset-AcctServiceGroupMembership
- Set-AcctDBConnection
- Set-AcctServiceMetadata
- Test-AcctDBConnection

### **Citrix.Analytics.Admin.V1:**

- Get-AnalyticsDBConnection
- Get-AnalyticsDBSchema
- Get-AnalyticsDBVersionChangeScript
- Get-AnalyticsInstalledDBVersion
- Import-AnalyticsDataDefinition
- Remove-AnalyticsServiceMetadata
- Reset-AnalyticsServiceGroupMembership
- Set-AnalyticsDBConnection

- Set-AnalyticsServiceMetadata
- Set-AnalyticsSite
- Set-AnalyticsDBConnection

**Citrix.DelegatedAdmin.Admin.V1:**

- Add-AdminRight
- Get-AdminDBConnection
- Get-AdminDBSchema
- Get-AdminDBVersionChangeScript
- Get-AdminInstalledDBVersion
- Import-AdminRoleConfiguration
- New-AdminAdministrator
- Remove-AdminAdministrator
- Remove-AdminAdministratorMetadata
- Remove-AdminRight
- Remove-AdminServiceMetadata
- Reset-AdminServiceGroupMembership
- Set-AdminAdministrator
- Set-AdminAdministratorMetadata
- Set-AdminDBConnection
- Set-AdminServiceMetadata
- Test-AdminDBConnection

**Citrix.Broker.Admin.V2:**

- Get-BrokerDBConnection
- Get-BrokerDBSchema
- Get-BrokerDBVersionChangeScript
- Get-BrokerInstalledDBVersion
- Get-BrokerLease
- New-BrokerMachineConfiguration
- Remove-BrokerControllerMetadata
- Remove-BrokerLease
- Remove-BrokerLeaseMetadata
- Remove-BrokerMachineConfigurationMetadata
- Remove-BrokerMachineConfiguration
- Remove-BrokerSiteMetadata
- Remove-BrokerUserFromApplication
- Reset-BrokerLicensingConnection
- Reset-BrokerServiceGroupMembership
- Set-BrokerControllerMetadata

- Set-BrokerDBConnection
- Set-BrokerLeaseMetadata
- Set-BrokerMachineConfiguration
- Set-BrokerMachineConfigurationMetadata
- Set-BrokerSiteMetadata
- Test-BrokerDBConnection
- Test-BrokerLicenseServer
- Update-BrokerBrokerLocalLeaseCache

**Citrix.Configuration.Admin.V2:**

- Export-ConfigFeatureTable
- Get-ConfigDBConnection
- Get-ConfigDBSchema
- Get-ConfigDBVersionChangeScript
- Get-ConfigInstalledDBVersion
- Get-ConfigServiceGroup
- Import-ConfigFeatureTable
- Register-ConfigServiceInstance
- Remove-ConfigRegisteredServiceInstanceMetadata
- Remove-ConfigServiceGroup
- Remove-ConfigServiceGroupMetadata
- Remove-ConfigServiceMetadata
- Remove-ConfigSiteMetadata
- Reset-ConfigServiceGroupMembership
- Set-ConfigDBConnection
- Set-ConfigRegisteredServiceInstance
- Set-ConfigRegisteredServiceInstanceMetadata
- Set-ConfigServiceGroupMetadata
- Set-ConfigServiceMetadata
- Set-ConfigSite
- Set-ConfigSiteMetadata
- Test-ConfigDBConnection
- Unregister-ConfigRegisteredServiceInstance

**Citrix.Host.Admin.V2:**

- Get-HypDBConnection
- Get-HypDBSchema
- Get-HypDBVersionChangeScript
- Get-HypInstalledDBVersion
- Remove-HypServiceMetadata

- Reset-HypServiceGroupMembership
- Set-HypDBConnection
- Set-HypServiceMetadata
- Test-HypDBConnection

**Citrix.ConfigurationLogging.Admin.V1:**

- Get-LogDBConnection
- Get-LogDBSchema
- Get-LogDBVersionChangeScript
- Get-LogInstalledDBVersion
- Remove-LogOperation
- Remove-LogServiceMetadata
- Remove-LogSiteMetadata
- Reset-LogDataStore
- Reset-LogServiceGroupMembership
- Set-LogDBConnection
- Set-LogServiceMetadata
- Set-LogSite
- Set-LogSiteMetadata
- Test-LogDBConnection

**Citrix.MachineCreation.Admin.V2:**

- Get-ProvDBConnection
- Get-ProvDBSchema
- Get-ProvDBVersionChangeScript
- Get-ProvInstalledDBVersion
- Get-ProvServiceConfigurationData
- Remove-ProvServiceConfigurationData
- Remove-ProvServiceMetadata
- Reset-ProvServiceGroupMembership
- Set-ProvDBConnection
- Set-ProvServiceConfigurationData
- Set-ProvServiceMetadata
- Test-ProvDBConnection

**Citrix.EnvTest.Admin.V1:**

- Get-EnvTestDBConnection
- Get-EnvTestDBSchema
- Get-EnvTestDBVersionChangeScript
- Get-EnvTestInstalledDBVersion
- Remove-EnvTestServiceMetadata

- Reset-EnvTestServiceGroupMembership
- Set-EnvTestDBConnection
- Set-EnvTestServiceMetadata
- Test-EnvTestDBConnection

**Citrix.Monitor.Admin.V1:**

- Get-MonitorConfiguration
- Get-MonitorDBConnection
- Get-MonitorDBSchema
- Get-MonitorDBVersionChangeScript
- Get-MonitorDataStore
- Get-MonitorDataStore
- Get-MonitorInstalledDBVersion
- Remove-MonitorServiceMetadata
- Reset-MonitorDataStore
- Reset-MonitorServiceGroupMembership
- Set-MonitorConfiguration
- Set-MonitorDBConnection
- Set-MonitorServiceMetadata
- Test-MonitorDBConnection

**Citrix.Storefront.Admin.V1:**

- Build-SfCluster
- Get-SfClusters
- Get-SfDBConnection
- Get-SfDBSchema
- Get-SfDBVersionChangeScript
- Get-SfInstalledDBVersion

**Citrix Virtual Apps and Desktops service discovery module for App-V packages and servers**

The Citrix Virtual Apps and Desktops service can deliver applications contained in App-V packages to your endpoints using either of the following methods:

- Single admin management method (accessing packages from a network share)
- Dual admin management method (accessing packages from a Microsoft App-V Management Server)

The process of registering App-V packages, Microsoft App-V Management, and Publishing Servers with the Application Library using the Citrix Virtual Apps and Desktops service differs slightly from register-



ing packages using an on-premises deployment. However, the process of assigning applications to users and launching them on a user's endpoint is identical.

The service management console in Citrix Cloud cannot view files in a resource location. Also, it cannot directly discover App-V packages or Microsoft App-V servers in your infrastructure. The discovery module provides functions that discover App-V package information in your on-premises infrastructure and uploads the package information to your Virtual Apps and Desktops service. Package information includes App-V packages, Microsoft App-V servers, and the apps that the packages contain.

The discovery module uses the Virtual Apps and Desktops Remote PowerShell SDK. It can discover package information from either a network share or a Microsoft App-V Management Server. You use the discovery module on a machine in your resource location.

Prerequisites for using the discovery module:

- Verify that PowerShell 3.0 or later is available on the machine.
- Verify that the Citrix Virtual Apps and Desktops Remote PowerShell SDK is installed on the machine.
- Verify that you have access to the network share containing the App-V packages.
- Verify that you have access to the server where the Citrix Cloud Connectors are installed and the Microsoft App-V Management Server is hosted.

### Add App-V packages to the Application Library in Citrix Cloud

The following procedure is valid for adding App-V packages from network shares (single admin management) and adding all published App-V packages from the Microsoft App-V Management Server (dual admin management). With the dual admin management method, you must manage the added App-V packages just as you do when using the single admin management method.

1. Download the discovery module from the Citrix Virtual Apps and Desktops Service downloads page <https://www.citrix.com/downloads/citrix-cloud/product-software/xenapp-and-xendesktop-service.html>. Extract the zip file `Citrix.Cloud.AppLibrary.Admin.v1.psm1` to a convenient folder.

**Note:**

This file is also provided on the Citrix Virtual Apps and Desktops ISO in `Support\Tools\Scripts`. You can copy it locally or reference it directly from the CD drive.

2. Verify that the Virtual Apps and Desktops Remote PowerShell SDK is installed on your machine
3. Navigate to the folder containing the discovery module. In the PowerShell window, type the full path of the folder containing the discovery module and then press **Enter**.
4. Import the discovery module with the command `Import-Module.\Citrix.Cloud.AppLibrary.Admin.v1.psm1`.

5. Add the App-V packages to the Application Library in Citrix Cloud using either of the following methods.

- To add App-V packages from a network share, run the PowerShell cmdlet: `Import-AppVPackageToCloud`.

For example: `Import-AppVPackageToCloud -PackagePath \\AppVsrv\share\notepad++.appv`

For cmdlet help, type `Get-Help Import-AppVPackageToCloud`.

- To add App-V packages from a Microsoft App-V Management Server, run the PowerShell cmdlet: `Import-AppVPackagesFromManagementServerToCloud`

For example: `Import-AppVPackagesFromManagementServerToCloud -ManagementSrvFQDN AppVMngSrv.domain.local`

For cmdlet help, type `Get-Help Import-AppVPackagesFromManagementServerToCloud`

This command imports all published App-V packages from the Microsoft App-V Management Server to Citrix Cloud.

After adding the App-V packages to Citrix Cloud, you must manage them as you do using the single admin management method.

6. Sign in to Citrix Cloud. Select the target customer. After the script runs successfully, the App-V packages are added to the Application Library in Citrix Cloud.

### Remove an App-V package from the Application Library in Citrix Cloud

To remove an App-V package from the Application Library in Citrix Cloud, see [Remove an App-V package from the Application Library](#) in on-premises deployments.

### High-level PowerShell functions

The module contains the following high-level functions that you can call from your own PowerShell script:

- `Import-AppVPackageToCloud -PackagePath <Full UNC path to App-V package>`  
Discovers and uploads to the Citrix Virtual Apps and Desktops service all the information necessary to publish applications from a single App-V Package.
- `Import-AppVPackagesFromManagementServerToCloud -ManagementSrvFQDN <FQDN of a Microsoft App-V Management Server>`

Discovers the UNC paths of packages published by the Management Server and calls **Import-AppVPackageToCloud** for each one in turn.

Packages discovered in this way are loaded to the Citrix Virtual Apps and Desktops Service using the single admin management method. The Citrix Virtual Apps and Desktops Service cannot deliver packages using the dual admin management method.

- `Import-AppVDualAdminToCloud -ManagementSrvUrl <URL of a Microsoft App -V Management Server> -PublishingServerUrl <URL of a Microsoft App-V Publishing Server>`

Discovers Microsoft App-V Management and Publishing Servers and imports the content to the Application Library. This cmdlet imports all the packages managed using Microsoft App-V Management Server and related information. Servers can be added and removed through PowerShell.

This cmdlet adds App-V packages in dual admin mode. Only App-V packages that are published on the Microsoft App-V Management Server, and which have AD groups added, are imported. If you make changes to the Microsoft App-V Management Server, rerun this cmdlet to synchronize the Application Library with the Microsoft App-V Management Server.

- `Remove-AppVServerFromCloud -ManagementSrvUrl <URL of a Microsoft App -V Management Server> -PublishingServerUrl <URL of a Microsoft App-V Publishing Server>`

Removes the Microsoft App-V Management and Publishing Servers added to Application Library.

This cmdlet removes the specified Microsoft App-V Management and Publishing Servers, plus all the associated App-V packages.

Run the discovery module for App-V packages and servers on a domain-joined computer within that resource location. Follow the guidance in *Install and use the Remote PowerShell SDK* to get started. Continue running PowerShell cmdlets or scripts. See the following examples.

### Example activities

Import the Virtual Apps and Desktops service App-V package discovery module.

```
1 import-module "D:\Support\Tools\Scripts\Citrix.Cloud.AppLibrary.Admin.v1.psm1"
2 <!--NeedCopy-->
```

Loop through the App-V Package store directory and upload each package.

```
1 Get-ChildItem -Path "\\FileServer.domain.net\App-V Packages" -Filter *.
  appv |
2 Foreach-Object{
3
```

```
4     Import-AppVPackageToCloud -PackagePath $_.FullName
5 }
6
7 <!--NeedCopy-->
```

Discover and upload packages registered with a Microsoft App-V management server.

```
1 Import-AppVPackagesFromManagementServerToCloud -ManagementSrvFQDN
   AppVManagementServer.domain.net
2 <!--NeedCopy-->
```

Discover Microsoft App-V Management and Publishing Servers and add the configuration to the Application Library. This also imports all the packages managed by the Microsoft App-V Management Server in dual admin mode.

```
1 Import-AppVDualAdminCloud -ManagementSrvUrl http://AppVManagementServer
   .domain.net - PublishingServerUrl http://AppVManagementServer.domain
   .net:8001
2 <!--NeedCopy-->
```

Read PowerShell help documentation included in the module.

```
1 Get-Help Import-AppVPackageToCloud
2 <!--NeedCopy-->
```

### Limitations

- You cannot discover App-V packages on your resource location infrastructure directly from the service management console in Citrix Cloud.
- The service management console in Citrix Cloud does not have a live connection to the Microsoft App-V Management server. Changes to Packages and other configuration in the Microsoft App-V Management server are not reflected in the service management console until `Import-AppVDualAdminCloud` is rerun. This differs from the [on-premises package discovery behavior](#)).

### Monitor Service OData API

In addition to using the Monitor functions to display historical data, you can query data using the Monitor Service's API. Use the API to:

- Analyze historical trends for planning
- Perform detailed troubleshooting of connection and machine failures

- Extract information for feeding into other tools and processes; for example, using Microsoft Excel's PowerPivot tables to display the data in different ways
- Build a custom user interface on top of the data that the API provides

For details, see [Monitor Service OData API](#). To access the Monitor Service API, see [Access Monitor Service data using the OData v4 endpoint in Citrix Cloud](#).

### **Citrix Virtual Apps and Desktops service APIs**

The Virtual Apps and Desktops service APIs are available at <https://developer.cloud.com/citrixworkspace/virtual-apps-and-desktops>.

### **Disclaimer**

This software / sample code is provided to you "AS IS" with no representations, warranties, or conditions of any kind. You may use, modify, and distribute it at your own risk. CITRIX DISCLAIMS ALL WARRANTIES WHATSOEVER, EXPRESS, IMPLIED, WRITTEN, ORAL OR STATUTORY, INCLUDING WITHOUT LIMITATION WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NONINFRINGEMENT. Without limiting the generality of the foregoing, you acknowledge and agree that (a) the software / sample code may exhibit errors, design flaws or other problems, possibly resulting in loss of data or damage to property; (b) it may not be possible to make the software / sample code fully functional; and (c) Citrix may, without notice or liability to you, cease to make available the current version and/or any future versions of the software / sample code. In no event should the software / code be used to support of ultra-hazardous activities, including but not limited to life support or blasting activities. NEITHER CITRIX NOR ITS AFFILIATES OR AGENTS WILL BE LIABLE, UNDER BREACH OF CONTRACT OR ANY OTHER THEORY OF LIABILITY, FOR ANY DAMAGES WHATSOEVER ARISING FROM USE OF THE SOFTWARE / SAMPLE CODE, INCLUDING WITHOUT LIMITATION DIRECT, SPECIAL, INCIDENTAL, PUNITIVE, CONSEQUENTIAL OR OTHER DAMAGES, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. You agree to indemnify and defend Citrix against any claims arising from your use, modification, or distribution of the code.

**Locations**

Corporate Headquarters | 851 Cypress Creek Road Fort Lauderdale, FL 33309, United States  
Silicon Valley | 4988 Great America Parkway Santa Clara, CA 95054, United States

© 2021 Citrix Systems, Inc. All rights reserved. Citrix, the Citrix logo, and other marks appearing herein are property of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered with the U.S. Patent and Trademark Office and in other countries. All other marks are the property of their respective owner(s).