XenCenter 7.1 LTSR

Contents

    11

Getting Started with XenCenter

September 30, 2020

With XenCenter, you can manage your XenServer environment and deploy, manage, and monitor virtual machines from your Windows desktop machine. Just click the topics below to get started.

Add a New Server	Connecting to XenServer host servers and adding them to the list of managed resources in XenCenter.
Create a Shared Storage	Creating XenServer storage repositories (SR) to provide storage that can be shared between managed servers.
Create a New Pool	Grouping managed servers together into a resource pool with shared storage using the **New Pool** wizard.
Create a Virtual Machine	Creating new virtual machines (VMs) with the **New VM** wizard.
Managing Users	Configuring access control by adding Active Directory (AD) user accounts and assigning different levels of access through the Role Based Access Control (RBAC) feature.

For information on system requirements for Citrix XenServer and XenCenter, see the *XenServer Installation Guide*.

Starting or Exiting XenCenter

June 3, 2019

Starting XenCenter

To start a XenCenter session, do one of the following:

- On the **Start** menu, choose: **Start > All Programs > Citrix > Citrix XenCenter**

- Double-click the **Citrix XenCenter** desktop shortcut.

If XenCenter was configured in an earlier session to restore your server connections on startup and a master password was set, you will be prompted to enter this password before continuing. See Store Your Server Connection State to find out more about how to set your server reconnection preferences.

Note that it is possible to run only one XenCenter session per user.

Exiting XenCenter

To exit the current XenCenter session: on the **File** menu, click **Exit**.

Any servers and VMs that are running when you exit will continue running after the XenCenter application window closes.

If there are any XenCenter tasks running, such as importing or exporting VMs, you will be warned when you try to exit. You can choose to exit anyway, in which case unfinished tasks may not complete successfully, or wait until the unfinished tasks have completed.

Uninstalling XenCenter

June 3, 2019

To uninstall XenCenter:

1. Open the Windows Control Panel.
2. On the Control Panel, under **Programs** , Click **Uninstall a program**
3. Select **Citrix XenCenter** from the list and then click **Uninstall**.

Note that XenCenter user configuration data and log files will not be removed when you uninstall the XenCenter application. The log files and user configuration data are stored in the folder:

```
1  %appdata%\Citrix\XenCenter
2  <!--NeedCopy-->
```

Exploring the XenCenter Workspace

June 3, 2019

Topics

- XenCenter Window Overview
- The Toolbar
- The Resources Pane
- The Tabs
- Keyboard Shortcuts
- Changing XenCenter Options
- Hidden Objects

XenCenter Window Overview

June 3, 2019
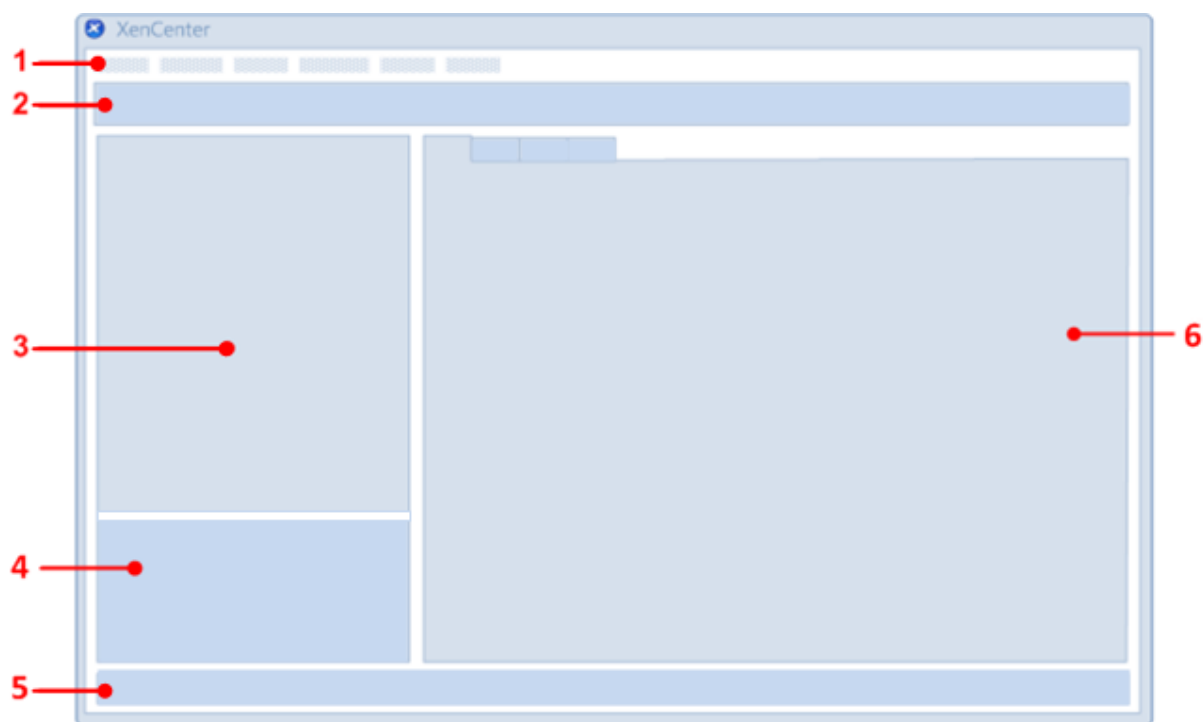
Ref #	Name	Description
1	Menu bar	Includes all of the commands you need to manage servers, pools, SRs, VMs and templates.
2	Toolbar	Provides quick access to a subset of the most frequently used menu commands. See The Toolbar.

Ref #	Name	Description
3	Resources pane	Lists all of the servers, pools, VMs, templates, and SRs currently being managed from XenCenter. See The Resources Pane
4	Navigation pane	Lists all of the navigation buttons. Click on a button to see a corresponding view of the managed resources in the resources pane.
5	Status bar	Displays progress information about the current task.
6	Properties tabs	View and set properties for the selected resource. See The Tabs.

The Toolbar

September 30, 2020

The XenCenter toolbar provides quick access to some of the most common XenCenter tasks, for example, to connect to new servers and create new VMs.
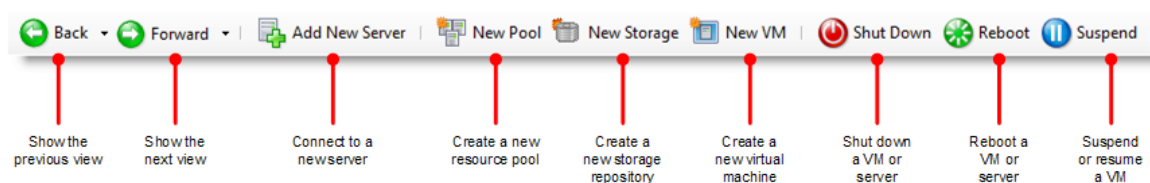
Using the Back and Forward buttons on the Toolbar

The **Back** and **Forward** buttons on the toolbar work like Back and Forward buttons on a browser and allow you to quickly move between views of your resources.

- To display your previous resource view, click **Back**.
- To display the next resource view (if you have used Back), click **Forward**.

- To display one of the resource views you have used in this session, click the ˅ arrow next to the Back or Forward buttons, and then select the view you want from the list.
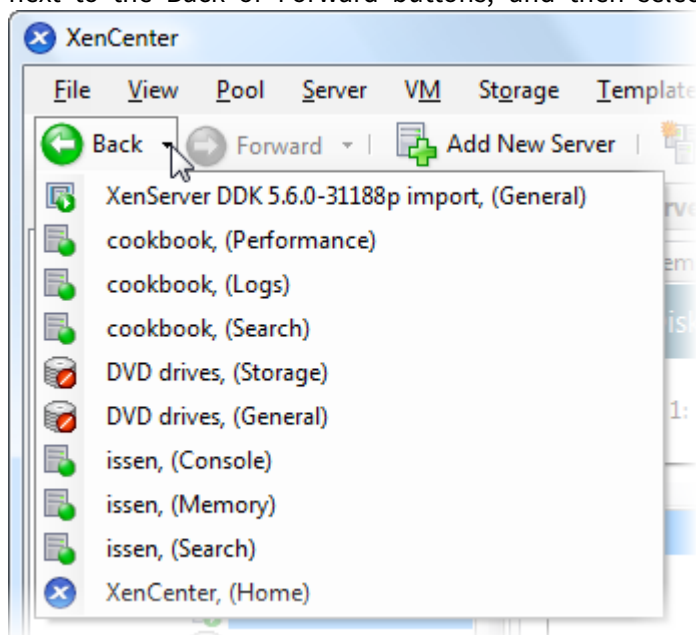
Showing and hiding the Toolbar

The XenCenter window displays the toolbar by default. However, you can hide the toolbar, for example, if you need to make more space in the XenCenter window for the console display. To hide the toolbar, do one of the following:

- Right-click anywhere on the toolbar and, on the shortcut menu, click to remove the **Show Toolbar** check mark.
- On the **View** menu, click to remove the **Toolbar** check mark.

> **Note:**
>
> Any changes you make to your XenCenter toolbar visibility are persistent and will be saved from session to session.

The Resources Pane

June 3, 2019

The **Resources pane** displays details about the managed resources - servers, pools, VMs, and storage. You can view your resources by their physical location or by properties such as folders, tags, or custom fields. The view in the Resources pane depends on the button you click in the Navigation pane. Refer to the table in the following section for information about various buttons in the Navigation pane.

You can do a simple text search on resource names by typing a word or a phrase in the Search box, located above the **Resources** pane. Matching resources are displayed as you type. To remove the query and view all your resources again, click the x button at the right of the Search box.

You can also apply a previously saved search query to the Resources pane. XenCenter includes several useful saved searches that allow you to search by network, operating system, power state, vApp membership, and XenServer Tools status. You can also create and add your own custom searches to this list at any time; see Create a Search Query and Saved searches for more information. To apply a saved search to the contents of the Resources pane, click **Saved Searches** in the Navigation pane and select a search query from the list.

The following table lists various options available in the Navigation pane.

Navigation button	Description
Infrastructure	Displays resources by their physical location, that is, by the host or pool to which they belong
Objects	Displays resources by categories such as, pools, servers, VMs, templates, etc.
Organization Views	Displays resources by folders, tags, custom fields, or by vApps
Saved Searches	Displays resources by the selected search criteria
Notifications	Displays the Notifications view which is a one-stop shop for alerts, updates, and events

For detailed information about the navigation buttons, see The Navigation Pane.

The Navigation Pane

June 3, 2019

The XenCenter Navigation pane provides various options to view and access managed resources. The navigation buttons **Infrastructure** , **Objects** , **Organizations Views** , **Saved Searches** , and **Notifications** provide a quick way to view and manage your resources.

The following sections provide an overview of the buttons in the Navigation pane:

Infrastructure
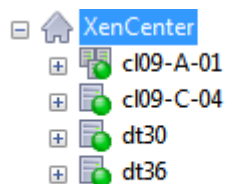
This is the default view. The Infrastructure view displays a tree view of the resources by their physical location. It provides a list of servers, VMs, templates, and storage resources by the pool or the server to which they belong.
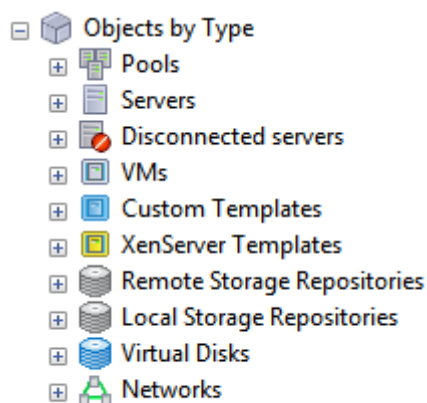
Objects

Click **Objects** to see a list of the resources by categories such as pools, servers, VMs, etc. Expand the nodes to view items in each category.

Organization Views

XenCenter allows you to group resources for ease of management. By default, XenCenter provides the following types of Organization Views:

- Objects by Folder
- Objects by Tag
- Objects by Custom Field
- vApps

Objects by Folder
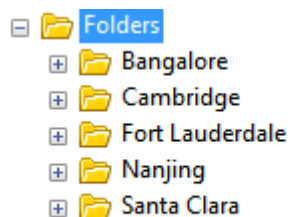
Select this option to view your resources by folders. You can create folders to group your resources by location, function, resource type, etc. You should note that organizing resources into a folder is conceptual, and not physical. The resources will not be physically moved to a folder if you choose to group them by Folders.

For detailed information about creating and managing folders to organize your resources, see Using Folders.

Objects by Tag
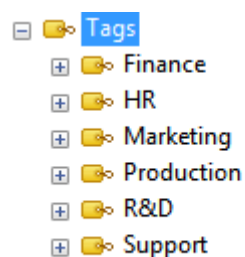
Select this option to view your resources by tags that you have previously defined. Tags are labels that you specify in order to view resources based on the criteria that you define. A single resource can contain multiple tags. For example, a server with the tag 'Production' can also be tagged as 'R&D'.

For detailed information about creating and managing tags in XenCenter, see Using Tags.

Objects by Custom Field
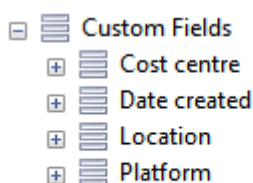
Select this option to view your resources by the customized fields you have previously defined. XenCenter enables you to add custom fields to your resources and provide a value in order to effectively

manage your resources. You simply add a custom field to a server, VM or any other resource in the pool, then give it a value. You can then use custom field values when building search queries.

For information on creating and using custom fields, see Using Custom Fields.

vApps

Select this option to view your VMs by the vApps they belong to. A vApp is a group of one or more VMs which can be managed as a single entity. For detailed information about vApps, see Managing vApps.

Saved Searches
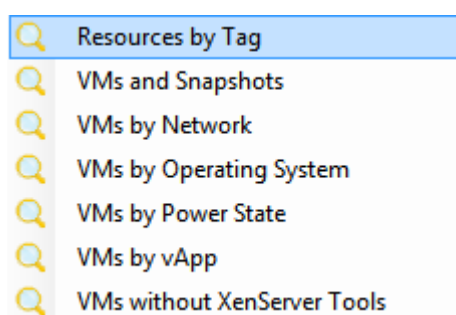
Click this button and select an option from the list to view resources that match the search criteria. By default, XenCenter includes a few saved searches that allow you to search your resources. You can create and add your own query to this list at any time.

For detailed information about the Search functionality in XenCenter, see Searching Resources.

Notifications
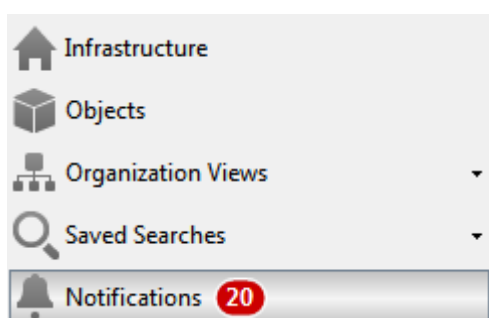
Click this button for the Notifications view. The Notifications view enables users to see all notifications in a centralized location and perform specific actions to address them. It contains Alerts, Updates, and the Events view.

⚠ Alerts (8)

⬆ Updates (12)

📅 Events

Alerts

The Alerts view displays a list of system alerts generated by XenCenter. You can filter the alerts by various options and take specific actions to address the alerts. For detailed information, see XenCenter Alerts.

Updates

Select this option to see a list of available XenServer and XenCenter updates. For more information, see Updating Managed Servers.

Events

Select this option to see a summary of all events in your current XenCenter session. For detailed information, see XenCenter Event Log.

The Tabs

September 30, 2020

The tab-based navigation in XenCenter provides quick access to your managed resources without needing to open and close dozens of windows at the same time. The tabs available at any time depend on what you have selected in the **Resources** pane; for example, most resources have a **General** tab, while the **HA** and **WLB** tabs are available only when a pool is selected, and the **Snapshots** tab is only available when a VM is selected.

Console

On this tab, you can run a console session on a VM or a managed server.

See also Run a Remote Console Session to read about the different types of remote VM console supported in XenCenter.

Switch to Remote Desktop

Switch to Default Desktop

Switches between Windows remote console types

Switch to Graphical Console

Switch to Text Console

Switches between Linux remote console types. You may need to enter your VNC password first when switching to a graphic console.

Open SSH Console

Opens an external SSH console as a pop-up window. This option is available

- On the host's Console tab to access the Control Domain (Dom0) console

- On the Console tab of a Linux VM to access the VM's console.

> **Note:**
>
> You must ensure that the Linux guest agent is installed on the VM to launch the SSH console.

Send Ctrl-Alt-Del

Sends the Ctrl+Alt+Del key sequence to the remote console.

By default, most keyboard shortcuts are transmitted to the server or virtual machine when you use a remote console. However, your local system will always intercept the **Ctrl+Alt+Del** key sequence and prevent it from being sent if you try to type it in directly at the remote console.

Undock (Alt+Shift+U)

Undocks the **Console** tab into a floating window.

To shut down or reboot a server, install XenServer Tools, shut down, reboot or suspend a virtual machine from within the floating console window, click ⏻ in the top left corner of the window and then click on a command.

To use a different keyboard shortcut for docking and undocking the console, go to the XenCenter **Options** dialog box: click **Tools > Options**.

Find Console

Opens the floating console window when it has been minimized or brings it to the front if it is hidden behind other windows.

Redock (Alt+Shift+U)

Reattach Console

Docks the floating console window back to the **Console** tab.

Scale

Scales the remote console screen to fit within the **Console** tab or window so that you can easily see everything on the remote console. Clear the check box to display the remote console screen at its normal size.

By default, the scale setting used in the **Console** tab is preserved when you undock the console or switch between console types (for example, VNC/text console), but this is configurable. To change this setting, go to the **Console** tab of the **Options** dialog box.

Fullscreen (Ctrl+Enter)

Displays the console in full-screen mode. Press **Ctrl+Alt** to exit full-screen mode; to use a different key sequence, go to the XenCenter Changing XenCenter Options dialog box.).

When you point to the top center of the screen in full-screen mode, the **Connection bar** is displayed, showing the name of the VM or server you are working on and including two controls: a Pin button to allow you to turn the **Connection bar** on permanently, and a **Restore down** button that you can click to exit full-screen mode.

You can control a number of console settings in the **Options** dialog box. For example, the text clipboard on your local machine is shared with the remote console by default, so that items you cut or copy will be placed on the clipboard and made available for pasting on either your local computer or on the remote console. You can turn clipboard sharing off and change a number of other console settings from the XenCenter **Options** dialog box; see Changing XenCenter Options.

General

View general properties of the selected container, virtual machine, server, resource pool, template or storage repository on the **General** tab; click **Properties** to set or change properties.

You can quickly copy any of the values shown on this pane to the Windows clipboard, for example, to paste into a text editor, by right-clicking on the value and then clicking **Copy** on the shortcut menu.

GPU

The GPU tab allows you to view or edit the GPU placement policy, view the available GPUs and virtual GPU types. The GPUs are grouped together based on the supported virtual GPU types. You can modify the virtual GPU types allowed on a particular GPU using the **Edit Selected GPUs** option. The horizontal bar in each group represents a physical GPU and it displays information about VMs running on the GPU.

For more information, see Configuring Virtual GPU and Change Pool Properties.

Note:

- GPU Pass-through and Graphics Virtualization is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. The GPU tab is displayed when the pool meets the license requirements and also has GPUs that support various virtual GPU types.
- There is no licensing restriction to use NVIDIA GPU pass-through for HVM Linux VMs.

HA

On the **HA** tab for a pool, you can:

- Enable HA using the **Configure HA** button.
- Change the pool's HA configuration using the **Configure HA** button.
- Disable (turn off) HA.

When HA has been enabled, you can see current HA status (failure capacity and server failure limit) and the status of the selected heartbeat SRs on the **HA** tab.

See Configure HA , Disable HA , and Change HA Settings.

Home

The **Home** tab allows you to add a server to the list of managed servers, open the XenCenter Help system, or open a browser window to find out more about XenServer and other Citrix products.

Memory

You can enable Dynamic Memory Control (DMC) and configure dynamic memory limits on the **Memory** tab. VMs can have a static memory allocation or can be configured to use DMC, which allows the amount of memory allocated to a VM to be adjusted on-the-fly as memory requirements on the server change without having to restart the VM. The Memory tab also lets you update the Control Domain (dom0) memory.

See Changing the Control Domain Memory , About VM Memory Configuration , Dynamic Memory Control (DMC) , and Configure DMC.

Networking

The **Networking** tab displays a list of networks configured on the pool, server, or the VM you have selected. It provides a centralized location to access or modify your network settings.

See About XenServer Networks and View and Change Network Properties.

NICs

View detailed information about the physical network interface cards (NICs) on the selected server and configure NIC bonds on the **NICs** tab.

NIC bonding (or "NIC teaming") can improve server resiliency by using two or more physical NICs as if they were one: if one NIC within the bond fails, the server's network traffic will automatically be routed over the second NIC, ensuring server management connectivity. See Configuring NICs.

> **Note:**
>
> You must use vSwitch as your network stack to bond four NICs. You can only bond two NICs when using Linux bridge.

Nutanix CVM Console

> **Note:**
>
> The **Nutanix CVM Console** tab will be displayed on XenServer hosts that are installed as part of Nutanix Hyperconverged infrastructure.

The **Nutanix CVM Console** tab allows you to access the CVM terminal. A CVM is a privileged VM similar to XenServer's control domain (dom0). The Nutanix software stack runs inside the CVM on each XenServer host. For more information, see the *XenServer-Nutanix Integration Guide* on the Citrix Product Documentation website.

Performance

View performance data for your VMs and managed servers on the **Performance** tab. Note that full performance data is only available for VMs with XenServer Tools installed.

The tab provides real-time monitoring of performance statistics across resource pools as well as graphical trending of virtual and physical machine performance. By default, graphs showing CPU, memory, network I/O and disk I/O are displayed on the tab. Click **Actions** to add more performance data and change the appearance of the graphs. For more information, see Configuring Performance Graphs.

Performance alerts can be generated when CPU, memory usage, network, storage throughput, or VM disk activity go over a specified threshold on a managed server, virtual machine, or storage repository. For more information, see Configuring Performance Alerts.

Search

Select the top-level XenCenter item, a pool or a server in the **Resources** pane and then click on the **Search** tab to perform complex searches of your managed resources. You can construct queries based on object types, folders, and attributes such as name, description, tags, HA status or restart priority, and power state.

See Create a Search Query , Filter and Group Search Results , Saved Searches , and Export and Import Searches.

Snapshots

Create, delete and export VM snapshots, revert a VM to a selected snapshot, and use existing snapshots to create new VMs and templates on the **Snapshots** tab.

See VM Snapshots.

Storage

View the storage configuration of the selected virtual machine, server, resource pool, or storage repository on the **Storage** tab. The settings shown on this tab depend on the type of resource currently selected in the **Resources** pane.

Selected resource	What's shown on the Storage tab	Learn more
VMs and templates	Information about each virtual disk on the VM is shown, including its size and location (the SR where the virtual disk is located), its data access status and disk access priority. To edit a virtual disk's settings, select it in the list and click Properties. Click Add to add a new disk or Attach to attach an existing disk.	Configuring Virtual Storage

Selected resource	What's shown on the Storage tab	Learn more
Servers and pools	A list of the available storage repositories (SRs) is shown, with summary information about their type, size, free space, and share status. To edit the name or description of an SR, select it in the list and click Properties. Click Add to add a SR or Detach to detach the selected SR.	Managing Storage Repositories (SRs)
Storage repositories	A list of the virtual disks or ISOs on the selected SR is shown. Click **Add** to add a new virtual disk.	Add Virtual Disks

Users

Configure role-based access to XenServer users and groups through AD user account provisioning and Role Based Access Control (RBAC) on the **Users** tab. Here, you can join a pool or server to an Active Directory (AD) domain, add an AD user or group to a pool, and assign roles to users and groups.

See Managing Users.

WLB

Access key Workload Balancing features, including configuration, optimization recommendations, and status on the **WLB** tab.

> **Note:**
>
> WLB is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. For more information about licensing, see About XenServer Licensing.

Resource Status Icons

September 30, 2020

The current status of managed resources - servers (hosts), virtual machines, storage, and templates - is represented using different icons in the **Resources** pane and elsewhere in XenCenter:

Servers

Icon	Description
	A server that is connected and is up and running normally.
	A server that is temporarily not connected to XenCenter, for example because it is being rebooted or suspended.
	A server that is disconnected, for example because it has been shut down.
	A server that is currently in Maintenance Mode. See Run in Maintenance Mode.
	A server on which a crash dump file has been created as a result of a system failure. Crash dump files are located in a folder named crash under the /var directory on the server. Crash dump file can provide invaluable information to your support engineer to aid in diagnosing XenServer-related problems, and can be included in server status reports generated in XenCenter using the **Get Server Status Report** utility. See Creating a Server Status Report. for more information on using this feature. When you remove the crash dump file from the /var directory on the server, the server status icon shown in XenCenter will be restored to normal.
	A server for which updates are available. See Updating Managed Servers.
	A server that is running an older version of XenServer than the pool master. See Updating Managed Servers.

Virtual machines, VM templates, and vApps

Icon	Description
	A virtual machine that is up and running normally.
	A virtual machine that is currently suspended.
	A virtual machine that is currently unavailable, for example because it is being rebooted or suspended.
	A virtual machine that is not running, for example because it has been shut down.
	A virtual machine that is currently migrating. See Migrate Virtual Machines.
	A XenServer VM template.
	A custom (user-defined) VM template.
	A XenServer vApp. See Managing vApps.

VM Snapshots

Icon	Description
	A disk-only VM snapshot.
	A scheduled disk-only VM snapshot.
	A disk and memory VM snapshot.
	A scheduled disk and memory VM snapshot.

Storage

Icon	Description
	A storage repository.
	The default storage repository for a pool.
	A storage repository that is not currently connected.

Icon	Description
	A storage repository that is currently unavailable. See Resolving SR Connectivity Problems.
	A storage repository that needs to be upgraded. See Upgrading Older SRs.
	A virtual disk.
	A virtual disk snapshot. This is a snapshot of a virtual machine's disks, and is created when a snapshot is made of the VM. See VM Snapshots for information about taking VM snapshots, and see the *XenServer Administrator's Guide* to find out more about disk snapshots.

Keyboard Shortcuts

June 3, 2019

You can use the keyboard as well as the mouse to navigate and perform tasks in XenCenter. For example, you can use the arrow keys to navigate between the items in the **Resources** pane and around the menus.

Navigating menus

To toggle Menu Mode on and off, press **F10** or **Alt**. In Menu Mode, you can use the keyboard to navigate menus.

Key	Action
Right Arrow , **Left Arrow**	Navigate across the menu bar, selecting each menu in turn.
Up Arrow , **Down Arrow**	Select each menu command in turn.
Enter	Activate the selected command.
Esc	Cancel the selected command and closes the menu.

Key	Action
Underlined letters (Access Keys)	Use the underlined letters to select specific menus and menu commands. For example, to copy a virtual machine, press **Alt** or **F10** , then **M** , then **C** to select the **VM** menu then **Copy VM**. In Windows Vista, the underlined letters are hidden by default, until you press **F10** or **Alt**. To change this setting, so that the underlined letters are always shown: On Windows Vista and Windows 7, in Control Panel, click **Ease of Access** , then under the **Ease of Access Center** heading, click **Change how your keyboard works**. Scroll down to the bottom of the page of options, then, under **Make it easier to use keyboard shortcuts** , select the **Underline keyboard shortcuts and access keys** check box.
Shortcut keys	Use shortcut key combinations to activate specific menu commands.

Using shortcut keys

You can use shortcut keys to perform tasks quickly with the keyboard rather than the mouse. For example, pressing **Ctrl+N** opens the **New VM** wizard, just like clicking N **ew VM** on the **VM** menu. Some shortcut keys are shown on menus and toolbar ToolTips. For numeric keypad keys, ensure that **Num Lock** is off.

Key	Action
F1	Display the online Help
Alt+F4	Exit XenCenter and close the XenCenter window
Ctrl+Alt	Toggle the console display between full screen mode and window mode
Ctrl+B	Start the selected VM
Ctrl+C	Copy the selected text to the Windows clipboard

Key	Action
Ctrl+E	Shut down the selected VM
Ctrl+N	Open the **New VM** wizard
Ctrl+R	Reboot the selected VM
Ctrl+V	Paste the selected text from the Windows clipboard
Ctrl+X	Cut the selected text to the Windows clipboard
Ctrl+Y	Suspend or resume the selected VM
Ctrl+Z	Undo the last text edit action

Changing XenCenter Options

July 15, 2020

You can change a number of settings that affect your XenCenter working environment. On the **Tools** menu, click **Options** , click on a tab and change the setting, and then click **OK** to save your changes.

Security settings

Option	Description	Default
Warn me when a new SSL certificate is found	Select this check box to have XenCenter display a warning whenever a new SSL (Secure Sockets Layer) security certificate is found on a managed server. Clear the check box if you do not want to see warnings about new certificates found on your managed servers when connecting to them.	Off

Option	Description	Default
Warn me when an SSL certificate changes	Clear this check box if you do not want to see warnings about modified certificates found on your managed servers when connecting to them. Select the check box to have XenCenter display a warning whenever a modified certificate is found on a managed server.	On
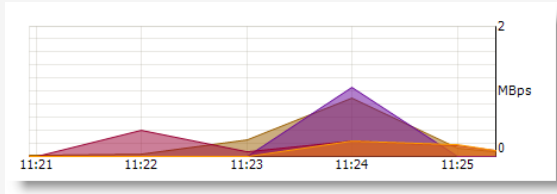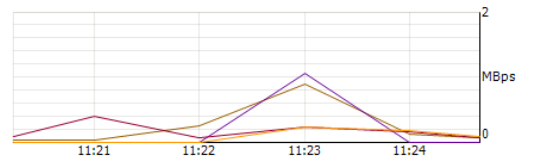
See also: Connecting and disconnecting servers.

Update settings

Option	Description	Default
Check for new versions of XenServer	Select this check box to have XenCenter periodically check and notify you when new versions of XenServer are available. Clear the check box to disable the periodic check.	On
Check for XenServer updates	Select this check box to have XenCenter periodically check and notify you when updates are available for XenServer. Clear the check box to disable the periodic check.	On
Check for new XenCenter versions	Select this check box to have XenCenter periodically check and notify you when a new version of XenCenter is available. Clear the check box to disable the periodic check.	On

See also: About software updates , Updating managed servers , Updating XenCenter.

Performance graph settings

Option	Description
Area graph	Click this radio button if you want to view the performance data shown on the **Performance** tab as area charts. For example:
Line graph	Click this radio button if you want to view the performance data shown on the **Performance** tab as line charts. For example:

See also: Monitoring system performance.

Console settings

Option	Description	Default
Share clipboard contents with remote console	Select this check box to share your local text clipboard contents with the remote console. Items cut or copied will be placed on the clipboard and made available for pasting on either your local computer or on the remote console.	On

Option	Description	Default
Full-screen mode	The keyboard shortcut to switch the console to and from full-screen mode.	Ctrl+Enter
Dock/Undock	The keyboard shortcut to undock the **Console** tab from the XenCenter window and to redock it.	Alt+Shift+U
Release keyboard and mouse	When the operating system on a VM captures your keyboard and mouse cursor for use within the guest OS, all keystrokes, mouse moves and button clicks that you make go to the VM. To return ownership of the keyboard and mouse to your host operating system, XenCenter reserves a special key on your keyboard: this is the host key . By default, the host key is the right Ctrl key on your keyboard. You can change this default here.	Right Ctrl
Preserve current scale setting when console is undocked	Select this check box to use the same console scale setting when the console is docked and when it is undocked.	On
Preserve current scale setting when switching back to the default console	Select this check box to keep the same console scale setting when switching between console types (for example, VNC/text console).	On

Option	Description	Default
Send Windows Key combinations to the Remote Desktop console	Select this check box to have XenCenter send any Windows Key combinations entered on your keyboard to the Remote Desktop console.	On
Receive sound from the Remote Desktop console	Select this check box to have XenCenter play sounds from applications running on the Remote Desktop console on your local computer (the computer on which you are running XenCenter).	On
Automatically switch to the Remote Desktop console when it becomes available	Select this check box to have XenCenter automatically switch from using the standard graphical console to using the Remote Desktop console whenever it is available.	On
Connect directly to the server's console session	Select this check box to have XenCenter connect to the existing console session on the remote server instead of creating a new virtual console session when opening a Remote Desktop console session.	On

Option	Description	Default
Enable Remote Desktop console scanning	Select this check box to have XenCenter automatically scan for an RDP connection. Clear the check box to prevent XenCenter from automatically scanning (polling) the RDP port, for example, if you have a firewall that blocks RDP traffic. Note that when this option is enabled, XenCenter will continue to scan the RDP port even if the **Automatically switch to Remote Desktop** option is turned off, so that you have the chance to switch to RDP as soon as it becomes available.	On

Any changes you make to the Windows Remote Desktop console settings will apply when you restart XenCenter.

See also: Run a remote console session.

Connection settings

Proxy server

XenCenter can be configured to connect directly to your managed servers or to use a proxy server. You can use your Internet Explorer proxy server settings, or you can specify a proxy server.

- Select **Don't use a proxy server** to have XenCenter connect directly to managed servers without using a proxy server.
- Select **Use proxy server settings from Internet Explorer** to use the same proxy settings as Internet Explorer.
- Select **Use this proxy server** if you want XenCenter to connect to the specified proxy server and use HTTP CONNECT to establish a secure SSL tunnel to your servers. Enter the address of the proxy server and the port number to use.

To have connection requests to local resources made directly and not via the proxy server, select the **Bypass proxy server for local addresses** check box. To have all connection requests made through the proxy server, clear the check box.

A local address is one that is on the LAN or intranet and is identified syntactically by the lack of a period (.) as in the URIs `http://webserver/` and `http://localhost/`.

Default: Don't use a proxy server

Connection timeout

You can specify how long to wait when trying to establish a connection with a managed server by adjusting the number of seconds to wait for a connection timeout. You should not set this value too low if you don't want to receive a lot of false alerts due to network-related problems.

Default: 20 seconds

See also: Connecting and disconnecting servers.

Save and Restore settings

Use the settings on this tab to specify if you want your login credentials for managed servers to be stored and used to automatically reconnect to all your managed servers at the start of each XenCenter session. You can also set a master password here to protect your stored login credentials.

Option	Description	Default
Save and restore server connection state on startup	Login credentials - your username and password - for all your managed servers can be stored between XenCenter sessions and used to automatically reconnect to them at the start of each new XenCenter session. When this check box is selected, XenCenter will remember the connection state of all your managed servers at the end of each session and attempt to restore them at the start of your next session.	Off

Option	Description	Default
Require a master password	When **Save and restore server connection state on startup** is enabled, you can protect your stored login credentials with a master password to ensure they remain secure; at the start of each session, you will be prompted to enter this master password before connections to your managed servers are automatically restored.	Off
Change Master Password	Click to change the current master password; you will be prompted to enter the current password then to enter and confirm the new master password.	

Plug-ins settings

Plug-ins are optional components that you can add to XenCenter to extend its functionality. You can add custom menu items or even whole tabs to the main window using a XenCenter plug-in. For example, you might do this as an ISV to integrate your own product with XenCenter, or as an end-user to integrate with your company's existing inventory management. A menu item can run a Microsoft PowerShell script or even an arbitrary executable on the client machine. Tabs are populated with a web page, and can call out to other services on your network or to your VMs.

Plug-in components

A XenCenter plug-in consists of the following components:

- An XML configuration file.
- A resource DLL for each supported locale.
- The application and any resources it requires.

Plug-in components should be placed into a plugins subfolder in your XenCenter installation folder. The components will be loaded when XenCenter starts. For example, in a default installation of Xen-

Center, your plug-ins would be located here:

```
1  C:\Program Files\Citrix\XenCenter\plugins\<your_organization_name>\<
      your_plugin_name>
2  <!--NeedCopy-->
```

Default: Off

View available plug-ins

To see a list of plug-ins currently available in XenCenter, and to enable or disable individual plug-ins, on the **Tools** menu, click **Options**. The **Options** dialog box will be displayed. From the list of options on the left pane, click **Plugins**.

Default: On

Creating plug-ins

To learn how to create plug-ins for XenCenter, see the XenCenter plugins web page.

Hidden Objects

June 3, 2019

Networks, PIFs, and VMs can be hidden from XenCenter by adding the key HideFromXenCenter=true to the other_config parameter for the object in the XenServer API. For example, certain VMs could be hidden because they shouldn't be used directly by general users in your environment. Objects flagged with this key are hidden objects , and by default they will not appear anywhere in XenCenter.

To make hidden objects visible in XenCenter, on the **View** menu, click to select **Hidden Objects**. To hide them again, on the **View** menu, clear the **Hidden Objects** check mark.

See the *XenServer Software Development Kit Guide* to find out more about flagging objects using the HideFromXenCenter key.

Organizing Resources

June 3, 2019

XenCenter provides a number of different ways of organizing your physical and virtual resources, allowing you to use the method that works best for you.

- Using Folders
- Using Tags
- Using Custom Fields

Using Folders

June 3, 2019

A folder is a container that you can use to group any managed resources in whatever manner makes sense for your environment; for example, you might create a folder for each branch office in your organization. Folders can contain any type of resource from anywhere in your XenServer environment. Resources can be accessed independently of the folder in which they are referenced.

The organization of folders in XenCenter is conceptual, not physical. The resources are not physically located in the folder, therefore, you can place resources into folders independently of their physical location. For example, placing a VM into a folder does not place its host server in the folder as well; placing a server into a folder does not place all the VMs and storage resources on that server into the folder as well.

Folders can also be used in search queries, for example, you can search by folder, with a contained in / not contained in relation and a drop down hierarchy of folders. See Create a Search Query.

The folder in which a resource is currently located is shown on the resource's **General** tab and in its **Properties** dialog box, so you can see folder information for a resource at all times. You can also move a resource into a different folder or remove it from a folder from the **General** tab.

To create a new folder

The simplest way to create a new folder is through Resources pane. Click **Organization Views** in the Navigation pane, and then select **Objects by Folder**. In the Resources pane, click on the **Folders** group, right-click, and select **New Folder** from the shortcut menu. Type a name for the new folder, select the server where your folder metadata will be stored, and then click **Create**.

You can also create new folders from the **General** tab for any resource:

1. In the Resources pane, select a pool, server, SR, virtual machine, or template, then click on the **General** tab and click **Properties**.
2. On the **General** tab of the **Properties** dialog box, click **Change** in the **Folder** box.
3. In the **Change Folder** dialog box, click **In this folder** button and then click **New Folder**.

4. Type a name for the new folder and select the server where your folder metadata will be stored, then click **Create**.

5. Click **Move** to apply the change and click **OK** on the **Properties** dialog box.

To move resources into and out of folders

The simplest way to move a resource into a folder is by dragging it from within the Resources pane or from the **Search** tab and dropping it onto the folder. Resources can only be in one folder, so if the resource is already in another folder it will be moved when you drag and drop in onto a different folder.

You can also move a resource into a different folder or remove it from a folder from the **General** tab:

1. In the Resources pane, select the pool, server, SR, virtual machine, or template you want to place in a folder, then click on the **General** tab and click **Properties**.
2. On the **General** tab of the **Properties** dialog box, click **Change** in the **Folder** box.
3. To remove the resource from its current folder, select **Not in any folder**.
4. To move the resource into a different folder, select **In this folder** and choose a folder or a sub-folder from the list.
5. To place the resource in a new folder, click **New Folder** , type a name for the new folder and select the server where your folder metadata will be stored, and then click **Create**.
6. Click **Move** to apply the change and click **OK** on the **Properties** dialog box.

To rename a folder

1. In the Resources pane, select the folder then right-click and select **Rename Folder** on the short-cut menu.
2. Type the new name.

To delete a folder

Note that you won't actually delete the resources in the folder when you delete the folder: they will be returned to the general collection.

1. In the Resources pane, select the folder, then right-click and select **Delete Folder** on the short-cut menu.
2. Click **Yes** to confirm.

All the contents will be moved out of the folder and then the folder will be deleted.

Using Tags

June 3, 2019

Tags enable you to identify your resources in new ways. Tags are like keywords or labels, and they allow you to rearrange your view of resources within XenCenter depending on criteria that are important to you, such as application, location, cost centre, owner, or lifecycle stage.

You make up tags when you need them and you can use as many as you like, and you can build searches based on your tags, for example "all Windows 7 test machines located in Cambridge".

Click on the **General** tab for a resource to see the tags currently assigned to that resource and to add and remove tags.

To create a new tag

1. In the **Resources** pane, select a resource, click on its **General** tab and then click **Properties**.
2. On the **General** tab of the **Properties** dialog box, click **Edit tags**.
3. Type a word or phrase in the **New Tag** box and then click **Create**.
4. The new tag will be automatically assigned to the selected resource. To remove the tag, clear the check box.
5. Click **OK**.
6. On the **Properties** dialog box, click **OK** to apply your changes.

To delete a tag

1. In the Navigation pane, select **Organization Views** and then **Objects by Tag**. The Tags group will be displayed on the Resources pane.
2. Select the tag you want to delete, right-click and then click **Delete Tag**. The tag will be removed from all resources that are currently tagged with it.

To tag a resource

The simplest way to assign an existing tag to a resource is by dragging and dropping it onto the tag in the Resources pane. You can drag and drop resources from within the Resources pane on the Tags group, and also from the search results list on the **Search** tab.

You can also assign an existing tag or a new tag to a resource using the **Edit tags** dialog box:

1. In the **Resources** pane, select the pool, server, SR, virtual machine, or template you want to tag, click on the **General** tab and then click **Properties**.
2. On the **General** tab of the **Properties** dialog box, click **Edit tags**.

- To create a new tag and add it to the selected resource, type a word or phrase in the **Edit Tags** box and then click **Create**.
- To add an existing tag, click to select the tag's check box in the **Tags** list and click **OK**.
3. On the **Properties** dialog box, click **OK** to apply your changes.

To untag a resource

To untag (remove a tag from) a resource, in the Navigation pane, click **Organization Views** and then click **Objects by Tag**. Select the resource that you would like to untag, right-click and then select **Untag object**.

You can also untag a resource using the **Edit Tags** dialog box:

1. In the Resources pane, select the pool, server, SR, virtual machine, or template you want to untag, click on its **General** tab and then click **Properties**.
2. On the **General** tab of the **Properties** dialog box, click **Edit tags**.
3. Clear the check box for the tag in the **Tags** list and click **OK**.
4. On the **Properties** dialog box, click **OK** to apply your changes.

Using Custom Fields

June 3, 2019

Custom fields allow you to add structured information to your resources which you can then use to help you find and manage them more effectively.

For example, you could identify all hosts with their physical location, for instance, or label the cost centre and applications running on all of your virtual machines. You simply add a custom field to a server, VM or any other resource in the pool, then give it a value. You can then use custom field values when building search queries.

Custom fields are shared at the pool level. Setting a custom field of any resource in a pool makes it available to all other resources in the pool, on the **General** tab and in the **Custom Fields** tab of the resource's **Properties** dialog box.

In the Navigation pane, click **Organization Views** and then select **Objects by Custom Field** to see your managed resources by the custom fields.

To create a new custom field

1. In the Resources pane, select any resource, click on the **General** tab and then click **Properties**.
2. On the **Custom Fields** tab, click **Edit Custom Fields**.

3. Click **Add** , type a name for the custom field and select the field type.

4. Click **OK** to apply your changes.

To assign a value to a custom field on a resource

1. In the Resources pane, select the resource and click on the **General** tab and then click **Properties**.

2. On the **Custom Fields** tab in the **Properties** dialog box, enter a value for the custom field(s).

3. Click **OK**.

The **General** tab for the resource displays all the Custom Fields set for that resource.

To delete a custom field

1. In the Resources pane, select the resource and click on the **General** tab and then click **Properties**.

2. On the **Custom Fields** tab in the **Properties** dialog box, click **Edit Custom Fields**.

3. Select the custom field in the list and then click **Delete**.

4. Click **Yes** to confirm.

Searching Resources

June 3, 2019

XenCenter enables you to perform complex searches of your managed resources. By default, XenCenter includes several searches that allow you to search resources by tag, and VMs by network, operating system, power state, vApp, and XenServer Tools status. You can also create and add your own custom searches to this list at any time. The view on the **Search** tab depends on the option you select in the Navigation pane. Select a view from the Navigation pane and then click on the **Search** tab to start your search query. The Search tab also displays a title that highlights the selection of resources for your search query.

You can also do a simple text search on resource names by typing a word or a phrase in the Search box, located above the Resources pane. Matching resources are displayed as you type. To remove the query and view all your resources again, click the **x** button at the right of the Search box. Refer to the following topics for more information:

- Create a Search Query
- Filter and Group Search Results
- Export and Import Searches
- Saved Searches

Create a Search Query

June 3, 2019

Using the **Search** tab, you can construct queries based on object types, folders, and attributes such as name, description, tags, HA status or restart priority, and power state.

To create a new search query

1. On the **Search** tab, click **New Search**.

2. Under **Search for** , select the type of resource or combination of resources you want to search for. In addition to the resource types such as servers, VMs, and storage repositories, this list also contains some common combinations of resource types, as well as options to search all resources.

   To define your own search category, click **Custom** and select the resource types you want to search for.

   The search is applied as soon as you select an option under **Search for** , and the results are displayed immediately in the bottom half of the **Search** tab.

3. Click **Save** to save the search query.

4. Type a title for your search query in the **Name** box.

5. Click on the **Location** list to choose the server where the search query metadata will be saved, and select a server.

6. Click **Save**.

Notes:

- Double-click a search result on the **Search** tab to display the **General** tab for that resource.
- To refine the search further, you can apply filters to the results. For more information, see Filter and Group Search Results. Filters are applied as soon as you select a filter option, and the results will be updated immediately.
- To quickly place search results into folders, select **Organization Views** and then **Objects by Folder**. Perform a search query, select the search results and drag them onto folders in the Resources pane. Note that resources can only be in one folder, so if the resource is already in another folder it will be moved. See Using folders to find out more.
- To quickly tag search results, select **Organization Views** and then **Objects by Tag**. Perform a search query, select the search results and drag them onto tags in the Resources pane. For more information on using tags, see Using tags.

Filter and Group Search Results

June 3, 2019

You can apply filters to a selected search category to further refine the search. The filters that are available are appropriate to the type of resource you are searching for.

For example, when searching for servers, you can filter the results by server name, address (IP address of the server), the pool it's in, and the VMs on it. When searching for virtual disks, you can filter the results using criteria such as whether or not they are shared, and the storage repository where they are located.

To filter your search results

1. Click on the button immediately below **Filters** and choose a filter category from the list:
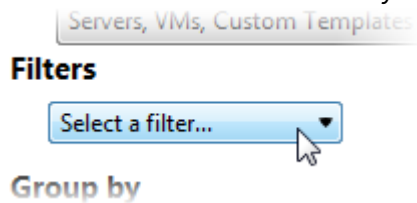
   You can only select filters that are applicable to the resources you are searching for.

2. If applicable, select an operator, such as Is , contains or Start Time and then enter a value. The filter is applied to the search result as soon as you select an option here, and the results are updated immediately.

3. To add more filters, click on the filter button immediately below **Filters** and choose a multi-filter
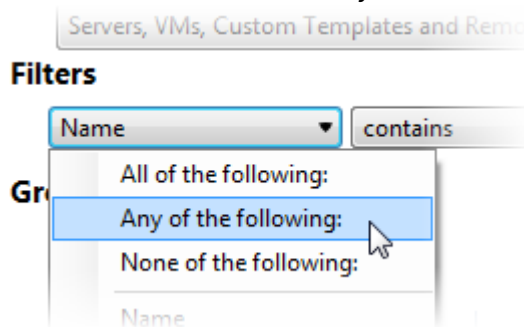
   category from the list:

4. To remove a filter, click on the **Remove** button on the right of the **Search** tab.

To group search results

1. Under **Group by** , select the grouping options you want to apply from the drop-down list. The grouping is applied as soon as you select an option here, and the results are updated immediately.

2. To add another grouping category to the search results, click **More** and then select a group category you want to add.
3. To remove a grouping, click on the group button and select **Remove Grouping**.

Saved Searches

June 3, 2019

XenCenter includes several useful saved searches that allow you to search resources by tag, and VMs by network, operating system, power state, vApp, and XenServer Tools status. You can modify these search queries at any time by clicking the **Edit Search** button on the **Search** tab.

You can also create and add your own custom searches to this list at any time; See Create a Search Query for more information. Custom searches appear at the top of the **Saved Searches** list in the XenCenter Navigation pane.

To apply a saved search

- To run a saved search in the Resources pane, select a search from the **Saved Searches** list in the XenCenter Navigation pane. Search results will be listed in the **Resources** pane.
- To run a saved search on the **Search** tab, click **Saved Searches** and then select a saved search query from the list.

To delete a saved search

On the **Search** tab, click **Saved Searches** and then **Delete** and select a saved search query from the list. Click **Yes** to confirm.

Export and Import Searches

June 3, 2019

XenCenter enables you to save search queries as XenSearch (.xensearch) files. When you export and save a search query, only the search query is saved, and not the results.

To export the current search

1. On the **Search** tab, click **Export**.
2. Enter a file name and location and then click **Save**.

To import a search

1. On the **Search** tab, click **Import**. Alternatively, on the XenCenter **File** menu, select **Import Search**.
2. Browse to locate the exported search file (file extension .xensearch) and then click **Open**. The imported search will be performed immediately, displaying results on the **Search** tab.

To save the imported search as a new custom search, click **Save** and then specify a **Name** for the search query and choose the **Location** where the search query metadata will be stored. Click **Save** to confirm.

About XenServer Licensing

June 3, 2019

XenServer uses the same licensing process as other Citrix products. XenServer requires a License Server. For details on how to install and run Citrix Licensing, see Citrix Product Documentation.

After purchasing support for XenServer you will be provided with `.LIC` license key. This license key should be installed on either:

- a Windows server running the Citrix License Server software
- a Linux-based Citrix License Server virtual appliance

XenServer Licensing is dependent on the version of XenServer you have installed on your host. Refer to the following table to know more about licensing concepts in various versions of XenServer.

Version	Topics
XenServer 7.1 CU2	Licensing Overview - XenServer 7.1 CU2 Managing XenServer 7.1 CU2 Licenses
XenServer 7.0	Licensing Overview - XenServer 7.0 Managing XenServer 7.0 Licenses

Licensing Overview - XenServer 7.1 CU2

February 3, 2021

XenServer 7.1 CU2 is available in two commercial editions:

- Standard

- Enterprise

The **Standard** edition is our entry-level commercial offering, with a range of features that will suit the needs of those customers who want a robust and high performing virtualization platform, but do not require the premium features offered by the Enterprise Edition; while still wishing to benefit from the assurance of comprehensive Citrix Support and Maintenance.

The **Enterprise** edition is our premium offering, optimized for both server, desktop and cloud workloads. In addition to the Standard edition, the Enterprise Edition offers the following features:

- Install and update I/O drivers through the Windows Update mechanism
- Automatic updating of the Management Agent
- Support for SMB storage
- Support for BitDefender hypervisor-level introspection
- Dynamic Workload Balancing
- GPU Virtualization with NVIDIA vGPU and Intel GVT-g
- VMware vSphere to XenServer Conversion utilities
- Intel Secure Measured Boot (TXT)
- Export Pool Resource Data
- In-memory Read Caching
- PVS-Accelerator
- Automated Updates using XenCenter
- XenServer Live Patching

> **Note:**
>
> Customers who have purchased Citrix Virtual Apps and Desktops continue to have an entitlement to XenServer, which includes all features listed above.

Citrix Licensing

XenServer uses the same licensing process as other Citrix products and as such requires a valid license to be installed on a License Server. You can download the License Server from Citrix Licensing. After purchasing the license for your XenServer edition, you will receive a `.LIC` license key. This license key should be installed on either:

- a Windows server running the Citrix License Server software
- a Linux-based Citrix License Server Virtual Appliance

> **Important:**
>
> Citrix XenServer 7.1 CU2 requires Citrix License Server 11.13.1.2 or higher.

When you assign a license to a XenServer host, XenServer contacts the specified Citrix License Server and requests a license for the specified server(s). If successful, a license will be checked out and the

---

License Manager displays information about the license the hosts are licensed under.

When requesting or applying licenses, information about the XenServer version and license type may be transmitted to Citrix. No other information relating to any users of the system, virtual machines running in the environment, or the environment in which XenServer is deployed, is collected or transmitted to Citrix. The limited information transmitted to Citrix during the licensing process is handled in accordance with the Citrix privacy policy, a copy of which can be reviewed at http://www.citrix.com/privacy .

Licensing XenServer Editions

XenServer does not support partial licensing, all hosts in a pool must be licensed. If you are running a XenServer pool with hosts that are licensed with different license types, the host with lowest license determines the features available to all members in the pool. Mixed pools of licensed and unlicensed hosts will behave as if all hosts were unlicensed. For information on managing licenses in XenServer 7.1 CU2, see Managing XenServer 7.1 CU2 Licenses.

Upgrades to the Enterprise Edition are available from the Standard edition. For detailed information about XenServer licensing, see the *XenServer Installation Guide*. To upgrade or to buy a XenServer license, visit the Citrix website.

License Expiry

XenCenter notifies you when your license is due to expire. You should purchase a license before it expires. When a XenServer license expires:

- XenCenter License Manager will display the status as **Unlicensed**.
- you will no longer be able to access licensed features or receive Citrix Technical Support for any host within the pool until you purchase another license.

License grace period

Citrix licensing has built-in timeout technology. After a startup license is checked out by a XenServer host, XenServer and the License Server exchange "heartbeat" messages every five minutes to indicate to each other that they are still up and running. If a XenServer host cannot contact the License Server, for example, due to problems with the License Server hardware or software or network failures, the server lapses into a 30-day licensing grace period. During the grace period, XenServer licenses itself through cached information and the hosts are allowed to continue operations as if they were still in communication with the License Server. The grace period is 30 days and when the grace period runs out, XenServer reverts to an unlicensed state. After communication is re-established between XenServer and the License Server, the grace period is reset.

Managing XenServer 7.1 CU2 Licenses

February 4, 2021

This topic contains information about managing licenses in XenServer 7.1 CU2.

All hosts in a pool must be licensed. You can manage your XenServer licenses using the **License Manager** dialog box in XenCenter. The License Manager allows you to:

- **Assign** XenServer licenses to managed servers. When you assign a license, XenServer contacts the Citrix License Server and requests the specified type of license. If a license is available, it is then checked out from the License Server.
- **Release** XenServer licenses. When you release a license, XenServer contacts the Citrix License Server and checks the license back in.

> **Important:**
>
> Citrix XenServer 7.1 CU2 requires Citrix License Server v11.13.1.2 or higher. You can download the License Server from Citrix Licensing.

To discover the license status of servers and pools

To see the license type of a server or pool, select that server or pool in the tree view. XenCenter displays the license status in the title bar for that server or pool, after the server or pool name.

You can also go to the **General** tab of the server and find the license type in the **License Details** section.

Mixed pools of licensed and unlicensed hosts behave as if all hosts were unlicensed. In the tree view XenCenter displays unlicensed pools with a warning triangle icon.

To assign a XenServer license

1. From the **Tools** menu, click **License Manager**.

2. Select the host(s) or pool(s) that you wish to assign a license and then click **Assign License**.

3. In the **Apply License** dialog box, select the license you want to request from the License Server. For more information about various XenServer licenses, see Licensing Overview - XenServer 7.1 CU2.

4. Enter the License Server details and then click **OK**.

   > **Note:**
   >
   > By default, the License Server uses port **27000** for communication with Citrix products. If you changed the default port on the License Server, enter the appropriate number in the

> **Port number** box. For more information about changing port numbers due to conflicts, refer to the Licensing topics on the Citrix Product Documentation website.

XenCenter contacts the specified Citrix License Server and requests a license for the specified server(s). If successful, a license is checked out and the information displayed in the XenCenter License Manager will be updated.

To release a XenServer license

1. On the **Tools** menu, click **License Manager**.
2. Select the host(s) or pool(s) and then click **Release License**.

Licensing Overview - XenServer 7.0

February 3, 2021

XenServer 7.0 is available in two commercial editions:

- Standard
- Enterprise

The **Standard** edition is our entry-level commercial offering, with a range of features that will suit the needs of those customers who want a robust and high performing virtualization platform, but do not require the premium features offered by the Enterprise Edition; while still wishing to benefit from the assurance of comprehensive Citrix Support and Maintenance.

The **Enterprise** edition is our premium offering, optimized for both server, desktop and cloud workloads. In addition to the Standard edition, the Enterprise Edition offers the following features:

- Install and update I/O drivers through the Windows Update mechanism
- Automatic updating of the Management Agent
- Support for SMB storage
- Support for BitDefender hypervisor-level introspection
- Dynamic Workload Balancing
- GPU Virtualization with NVIDIA vGPU and Intel GVT-g
- VMware vSphere to XenServer Conversion utilities
- Intel Secure Measured Boot (TXT)
- Export Pool Resource Data
- In-memory read caching

> **Note:**
>
> Customers who have purchased XenApp or XenDesktop continue to have an entitlement

> to XenServer, which includes all the features contained within the Enterprise Edition. Customers should note that the In-memory read caching feature is available only for XenApp and XenDesktop Platinum customers.

Citrix Licensing

XenServer uses the same licensing process as other Citrix products and as such requires a valid license to be installed on a License Server. You can download the License Server from Citrix Licensing. After purchasing the license for your XenServer edition, you will receive a `.LIC` license key. This license key should be installed on either:

- a Windows server running the Citrix License Server software
- a Linux-based Citrix License Server Virtual Appliance

Important:

Citrix XenServer 7.0 requires Citrix License Server 11.12.1 or higher.

When you assign a license to a XenServer host, XenServer contacts the specified Citrix License Server and requests a license for the specified server(s). If successful, a license will be checked out and the **License Manager** displays information about the license the hosts are licensed under.

When requesting or applying licenses, information about the XenServer version and license type may be transmitted to Citrix. No other information relating to any users of the system, virtual machines running in the environment, or the environment in which XenServer is deployed, is collected or transmitted to Citrix. The limited information transmitted to Citrix during the licensing process is handled in accordance with the Citrix privacy policy, a copy of which can be reviewed at http://www.citrix.com/privacy .

Licensing XenServer Editions

XenServer does not support partial licensing, all hosts in a pool must be licensed. If you are running a XenServer pool with hosts that are licensed with different license types, the host with lowest license determines the features available to all members in the pool. Mixed pools of licensed and unlicensed hosts will behave as if all hosts were unlicensed. For information on managing licenses in XenServer 7.0, see Managing XenServer 7.0 Licenses.

Upgrades to the Enterprise Edition are available from the Standard edition. For detailed information about XenServer licensing, see the *XenServer Installation Guide*. To upgrade or to buy a XenServer license, visit the Citrix website.

License Expiry

XenCenter notifies you when your license is due to expire. You should purchase a license before it expires. When a XenServer license expires:

- XenCenter License Manager will display the status as **Unlicensed**.
- you will no longer be able to access licensed features or receive Citrix Technical Support for any host within the pool until you purchase another license.

License grace period

Citrix licensing has built-in timeout technology. After a startup license is checked out by a XenServer host, XenServer and the License Server exchange "heartbeat" messages every five minutes to indicate to each other that they are still up and running. If a XenServer host cannot contact the License Server, for example, due to problems with the License Server hardware or software or network failures, the server lapses into a 30-day licensing grace period. During the grace period, XenServer licenses itself through cached information and the hosts are allowed to continue operations as if they were still in communication with the License Server. The grace period is 30 days and when the grace period runs out, XenServer reverts to an unlicensed state. After communication is re-established between XenServer and the License Server, the grace period is reset.

Managing XenServer 7.0 Licenses

September 30, 2020

This topic contains information about managing licenses in XenServer 7.0.

All hosts in a pool must be licensed. You can manage your XenServer licenses using the **License Manager** dialog box in XenCenter. The License Manager allows you to:

- **Assign** XenServer licenses to managed servers. When you assign a license, XenServer contacts the Citrix License Server and requests the specified type of license. If a license is available, it is then checked out from the License Server.
- **Release** XenServer licenses. When you release a license, XenServer contacts the Citrix License Server and checks the license back in.

> **Important:**
>
> Citrix XenServer 7.0 requires Citrix License Server v11.12.1 or higher. You can download the License Server from Citrix Licensing.

To assign a XenServer license

1. From the **Tools** menu, click **License Manager**.

2. Select the host(s) or pool(s) that you wish to assign a license and then click **Assign License**.

3. In the **Apply License** dialog box, select the license you want to request from the License Server. For more information about various XenServer licenses, see Licensing Overview - XenServer 7.0.

4. Enter the License Server details and then click **OK**.

   > **Note:**
   >
   > By default, the License Server uses port **27000** for communication with Citrix products. If you changed the default port on the License Server, enter the appropriate number in the **Port number** box. For more information about changing port numbers due to conflicts, refer to the Licensing topics on the Citrix Product Documentation website.

XenCenter contacts the specified Citrix License Server and requests a license for the specified server(s). If successful, a license is checked out and the information displayed in the XenCenter License Manager will be updated.

To release a XenServer license

1. On the **Tools** menu, click **License Manager**.
2. Select the host(s) or pool(s) and then click **Release License**.

Getting Help

June 3, 2019

There are a number of different places you can look to find the information you need about using XenServer:

- **XenServer product documentation**. Comprehensive reference documentation aimed at XenServer administrators and developers is available in PDF format on the XenServer product CD and online in the XenServer section of the Citrix Knowledge Center.
- **Citrix Knowledge Center**. If you are connected to the Internet, you can go online to browse or search for knowledge base articles and technical notes, as well as the most up-to-date XenServer product documentation. For more information, visit the Citrix Knowledge Center.
- **XenCenter online help**. XenCenter includes a context-sensitive help system which provides assistance appropriate to where you are in the application and what you are trying to do.

Topics

- [The Help Window](#)
- [Searching the Help](#)
- [Printing Help Topics](#)

The Help Window

June 3, 2019

XenCenter Help is displayed in an HTML Help window:

- On the left of the window is the navigation pane containing the **Contents**, **Search** and **Favorites** tabs.
- On the right of the window is the topic pane. It displays the selected Help topic, or the default Help topic.
- At the top is the toolbar.

Using the Help toolbar

Icon	Description
	Hides the navigation pane, saving space on your screen by only displaying the topic pane.
	Displays the navigation pane. This button is available only if you have previously used **Hide**.
	Displays the last topic you viewed.
	Displays the next topic (if you have used **Back**).
	Prints a topic.

Using the Contents tab

Click the **Contents** tab to list the Help topics:

The folder icon starts a new subject. If the folder is closed, click the plus sign to expand it, listing its individual topics, and click the minus sign to collapse them again.

The topic icon indicates an individual help topic. Click the topic icon or its title to display it in the topic

pane.

The topic that you are currently viewing in the topic pane is highlighted in the Contents. When you move to another topic, the Contents expands as necessary to show you where you are in the Help system.

Using the Search tab

You can search the entire text of the Help system for words or phrases. You can limit the search to previous results, match similar words, or search topic titles only. See Searching the Help to find out more.

Using the Favorites tab

- To create a list of favorite Help topics, locate the Help topic you want to make a favorite topic, click the **Favorites** tab, and then click **Add**.
- To return to a favorite topic, click the **Favorites** tab, select the topic, and then click **Display**.
- To rename a favorite topic, right-click the topic in the list and click **Rename** on the shortcut menu, then type a new name for the favorite.
- To remove a favorite topic, select the topic and then click **Remove**.

Searching the Help

June 3, 2019

You can search the entire text of the Help system for words or phrases. You can limit the search to previous results, match similar words, or search topic titles only.

1. Click the **Search** tab, and then type the word or phrase you want to find.
2. To search for words in the topic titles only, select the **Search titles only** check box. All topic files are searched, including any that are not listed in the table of contents.
3. To search only the most recent group of topics found, select the **Search previous results** check box. To search through all of the files in the help system, clear this check box.
4. Click **List Topics** , and either double-click the topic you want, or select it and click **Display**.

Tips

- Searches are not case-sensitive, so you can type your search in uppercase or lowercase characters.
- You can search for any combination of letters (a-z) and numbers (0-9).

- Punctuation marks other than double quotes, such as the period, colon, semicolon, comma and hyphen, are ignored during a search.
- Use double quotes to search for a phrase. For example, "server root password" finds topics containing the phrase server root password , while server root password finds topics containing the words server and root and password somewhere but not necessarily together as one phrase.
- Use wildcard expressions to search for one or more characters: use an asterisk to search for zero or more characters and use a question mark to search for a single character.

Printing Help Topics

June 3, 2019

To print the current topic

1. Expand all the sections: select the **Show all text** check box at the top of the topic or click in the topic and then press **Alt+X**.
2. Right-click in the topic pane and click **Print** on the shortcut menu. Alternatively, click **Print** on the toolbar, click **Print the selected topic** , then click **OK**.
3. In the **Print** dialog box, set the printer options that you want, then click **Print**.

You can choose to print a group of topics shown in the Contents tab. However, note the following points:

- The topics are printed with all the sections expanded.
- The formatting of the printed topics will not look exactly as you see on screen.
- Images may print in incorrect locations, obscuring text nearby.

To print a group of topics

1. In the Contents tab, right-click the book containing the topics you want to print, then click **Print** on the shortcut menu or click **Print** on the toolbar.
2. Click **Print the selected heading and all subtopics** , then click **OK**.
3. In the Print dialog box, set the printer options that you want, then click Print.

Managing Servers

June 3, 2019

Connecting and disconnecting

- Add a new server
- Disconnect a server
- Reconnect a server
- Reboot a Server
- Shut Down a Server
- Restart Toolstack
- Configuring Host Power On
- Power on a server remotely
- Run in Maintenance Mode
- Exporting and importing a list of managed servers
- Changing Server Properties
- Changing the Control Domain Memory
- Store Your Server Connection State
- Back up and Restore a Server
- Remove a Server From XenCenter

Configuring networking

- About XenServer Networks
- Add a New Network
- Remove a Network
- View and Change Network Properties
- Configuring NICs
- Configuring IP Addresses

More information

Changing server properties

Changing the control domain memory

Exporting and importing a list of managed servers

Connecting and Disconnecting Servers

June 3, 2019

- Add a New Server
- Disconnect a Server

- Reconnect a Server
- Reboot a Server
- Shut Down a Server
- Restart Toolstack
- Configuring Host Power On
- Power on a server remotely
- Run in Maintenance Mode
- Store Your Server Connection State
- Back up and Restore a Server
- Remove a Server From XenCenter

Add a New Server

September 30, 2020

In order to monitor and manage activities on a server from XenCenter, first identify it as a managed resource. When you first connect to a server (via **Add New Server** on the toolbar or **Server** menu), it is added to the **Resources** pane on the left of the XenCenter window. The default storage repository for the server (if configured) and any physical CD or DVD drives on the server may also appear here. A managed server can subsequently be disconnected, reconnected, shut down or put into Maintenance Mode, and it remains accessible from the **Resources** pane until you remove it from XenCenter.

The first time you connect to a server using XenCenter, the **Save and Restore Connection State** dialog box appears. This enables you to set your preferences for storing your server connection information and automatically restoring server connections at the start of each XenCenter session; see Store Your Server Connection State.

To Add a Server to XenCenter

1. Click **Add New Server**. Alternatively:

    - On the **Server** menu, click **Add**.
    - In the **Resources** pane, select the top-level XenCenter entry, right-click and then click **Add** on the shortcut menu.
    - On the XenCenter Home page, click the **Add New Server** button:

2. Enter the IP address or DNS name of the server you want to add in the **Server** box. For example: 182.31.32.28 or belt.orion.com .

> **Tip:**
>
> You can add multiple servers with the same login credentials by entering the names or IP addresses separated by semicolons in the **Server** box.

3. Type the username, for example, "root", and the password set up during XenServer installation. If Active Directory (AD) authorization has been enabled in your XenServer environment, you can enter your AD credentials here. See RBAC overview for more information.

4. Click **Add**. A connection progress monitor is displayed: to cancel the connection, click **Cancel**.

Security Certificates

XenCenter may be configured to display a warning message whenever a new or modified SSL (Secure Sockets Layer) security certificate is found when connecting to a managed server. Click **View Certificate** to view the security certificate. To prevent SSL certificate warnings from being generated, use the **Security Settings** tab in the XenCenter **Options** dialog box.

Disconnect a Server

September 30, 2020

A disconnected server continues to be a managed server and remains available in the Resources pane with this status icon:

You can quickly see which of your servers are currently disconnected by switching to the **Objects** view in the Navigation Pane and then clicking on **Disconnected servers**.

- To disconnect a server, select it in the **Resources** pane and then, on the **Server** menu, click **Connect/Disconnect** and then **Disconnect**. You can reconnect to a disconnected server at any time. For more information, see Reconnect a server.
- To remove a disconnected server from the **Resources** pane, see Remove a server from XenCenter.

Reconnect a Server

June 3, 2019

Once you have added a server to XenCenter, it remains accessible in the **Resources** pane throughout the current XenCenter session regardless of the server's status - connected or disconnected, running normally or in Maintenance Mode.

To reconnect to a managed server that is disconnected, simply double-click on it in the **Resources** pane, or right-click and then click **Connect** on the shortcut menu. Server connection information is remembered for the duration of the current XenCenter session and so you will not need to enter the same login credentials more than once in the same XenCenter session if you want to reconnect using the same user account.

You can also reconnect to a connected server using different login credentials, for example, using your AD login instead of your local root account.

To reconnect to a connected server using different login credentials

1. Select the server in the **Resources** pane.
2. Do one of the following:
   - Right-click in the **Resources** pane and select **Reconnect As** on the shortcut menu.
   - On the **Server** menu, click **Connect/Disconnect** and then **Reconnect As**.
3. Enter the new username and password. If Active Directory authorization has been enabled in your XenServer environment, you can enter your AD credentials here. See RBAC overview.
4. Click **OK**.

Reboot a Server

June 20, 2019

When you reboot a server in XenCenter, any virtual machines (VMs) running on it are shut down, and then the server is disconnected and rebooted. If the server is a member of a pool, the loss of connectivity on shutdown will be handled and the pool will recover when the server returns. If you shut down another pool member (not the master), the other pool members and the master will continue to function. If you shut down the master, the pool will be out of action until the master is rebooted and back on line, at which point the other members will reconnect and synchronize with the master, or until you make one of the other members into the master (which you can do via the XenServer xe CLI).

Note that VMs with XenServer Tools installed ("paravirtualized" VMs) will be shut down gracefully when you reboot the host server, but VMs running in HVM mode (that is, VMs without XenServer Tools installed) will be shut down using a forced shutdown; to avoid this, you should Installing XenServer Tools on all HVM virtual machines.

After a server reboot, XenCenter will attempt to reconnect to the server automatically. Once the server is reconnected, you will need to restart any VMs that were running on it unless they are configured to automatically start on host server reboot (see Change VM properties).

To reboot a server

Select the server in the Resources pane and then click **Reboot** on the Toolbar.

Shut Down a Server

September 30, 2020

When you shut down a server in XenCenter, any virtual machines (VMs) running on it are shut down, and then the server is disconnected and powered off. If the server is a member of a pool, the loss of connectivity on shutdown will be handled and the pool will recover when the server returns. If you shut down another pool member (not the master), the other pool members and the master will continue to function. If you shut down the master, the pool will be out of action until the master is rebooted and back on line, at which point the other members will reconnect and synchronize with the master, or until you make one of the other members into the master (which you can do via the XenServer xe CLI).

Note that VMs with XenServer Tools installed ("paravirtualized" VMs) will be shut down gracefully, but VMs running in HVM mode (that is, VMs without XenServer Tools installed) will be shut down using a forced shutdown; to avoid this, you should install XenCenter tools on all HVM virtual machines.

After you power the server back on, you will need to connect to it again - see Reconnect a server.

To shut down a server:

Select the server in the Resources pane and then click **Shut Down** on the toolbar.

When the server has been shut down, its status in the **Resources** pane changes to **Disconnected** 📇.

Restart Toolstack

August 25, 2021

The **Restart Toolstack** option allows you to restart the XenServer management toolstack (also known as xapi). This toolstack controls VM lifecycle operations, host and VM networking, VM storage, user authentication, and allows the management of XenServer resource pools. xapi provides the publicly documented XenAPI Management Interface which is used by all tools that manage VMs and resource pools.

> **Note:**
>
> In some cases, the **Restart Toolstack** option can be used for troubleshooting XenServer issues.

> However, customers should be extremely cautious when using this option, as incorrect usage can cause unexpected results.

To restart the toolstack:
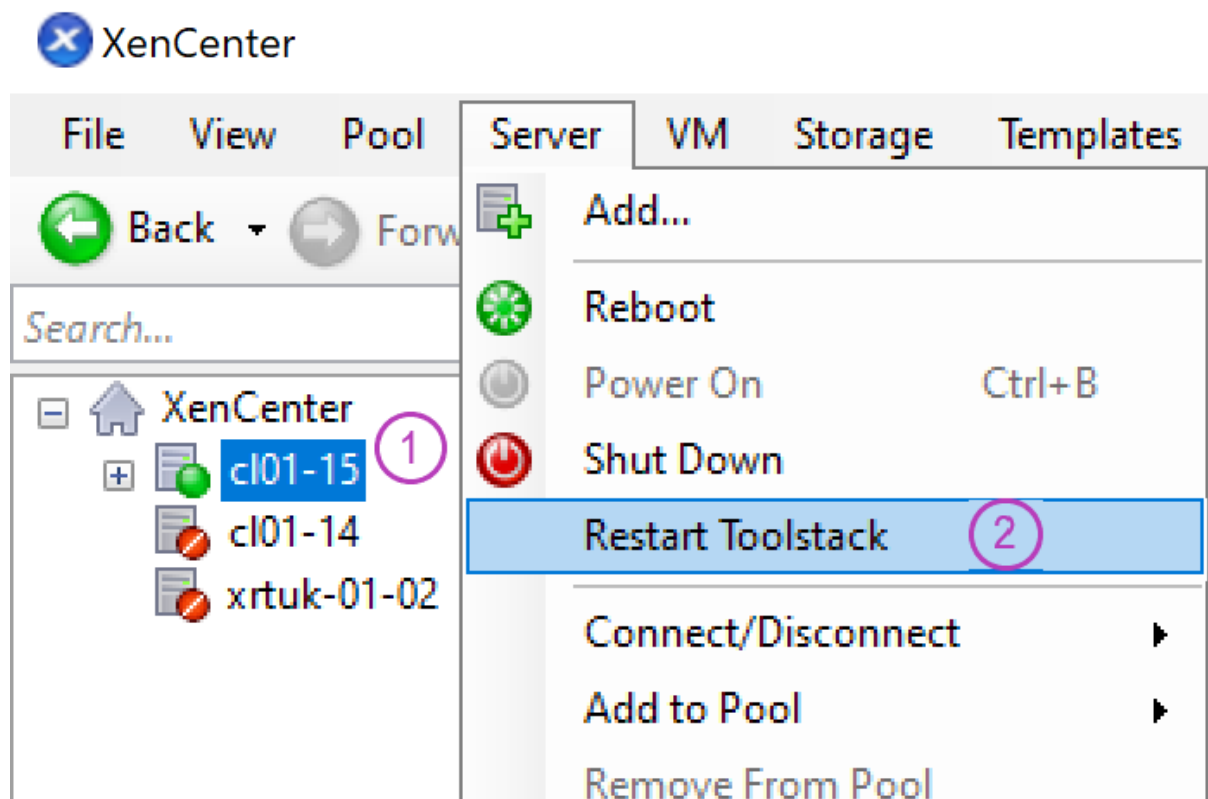
1. Select the server in the **Resources** pane.
2. On the **Server** menu, click **Restart Toolstack**.
3. Click **Yes** to confirm.

Note:

When you run the **Restart Toolstack** option on the Pool Master, XenCenter will lose connection to the pool. Wait for 30 seconds after losing connection, and then reconnect manually.

Configuring Host Power On

June 3, 2019

The XenServer Host Power On feature allows you to manually turn a remote host (server) on and off. To use this feature, you need to carry out the following steps:

1. Ensure that the host supports remote power control (that is, it has Wake-on-LAN functionality, a DRAC or iLO card, or it is using a custom script). See Prerequisites for Host Power On.

2. Enable the Host Power On functionality, as described in To enable or disable Host Power On. To perform this procedure for iLO and DRAC processors, you need the credentials for the processor, which are set in its firmware.

Once Host Power On has been configured on a server, you will be able to power the server on from XenCenter by selecting it and then, on the Server menu, clicking Power On.

If you have Workload Balancing installed and configured, you can also configure XenServer to turn hosts on and off automatically as VMs are consolidated or brought back online. This feature is known as Power Management.

Prerequisites for Host Power On

To enable the Host Power On feature, the host server must have one of the following power control solutions:

- A network card that supports **Wake On LAN (WOL)**.
- **Dell Remote Access Controller (DRAC)**. To use XenServer with DRAC, you need to follow these steps:
    1. Install the Dell supplemental pack.
    2. Install the RACADM command-line utility on the host server with the remote access controller.
    3. Enable DRAC and its interface. RACADM is often included in the DRAC management software. For more information, see Dell's DRAC documentation.
- **Hewlett-Packard Integrated Lights-Out (iLO)**. To use XenServer with iLO, iLO must be enabled on the host and be connected to the network. For more information, see HP's iLO documentation.
- A custom power-on script based on the Xen-API that enables you to turn the power on and off through XenServer. For DRAC and iLO, you can use the secrets feature (by specifying the key power_on_password_secret) to help you store your password more securely. For more information, see the *XenServer Administrator's Guide*.

To enable or disable Host Power On

You can enable Host Power On for an individual host server via the host's **Properties** window, or on multiple servers via the pool's **Properties** window.

1. Select the server or pool and open its **Properties** dialog box: on the **Server** or **Pool** menu, click **Properties**.

2. Click the **Power On** tab and under **Power On mode** , select the option you want:

    - **Disabled** Select this option to turn off the Host Power On feature.

- **Wake-on-LAN (WoL)** To use this option, the host must have a Wake on LAN-enabled network card.
- **HP Integrated Light-Out (iLO)** To use this option, iLO must be enabled on the host and be connected to the network. For more information, see HP's iLO documentation.
- **Dell Remote Access Controller (DRAC)** To use this option, the Dell supplemental pack must be installed on the host server to get DRAC support. For more information, see Dell's DRAC documentation.
- **Custom power-on script** You can use a custom Python Linux script to turn on the power on a XenServer host from a remote location. For information about creating the script, including a list of supported key/value pairs, see the XenServer Administrator's Guide.

3. If you selected either HP iLO or Dell DRAC, enter the following information:

- **IP Address** The IP address you specified configured to communicate with the power-control card. Alternatively, you can enter the domain name for the network interface where iLO or DRAC is configured.
- **User name** This is the iLO or DRAC user name that is associated with the management processor, which you may or may not have changed from its factory default settings.
- **Password** This is the password associated with that user name.

4. If you selected **Custom power-on script** , enter the filename and path to the custom script you created and then, under **Configuration options** , enter the key/value pairs you want to use to configure the script. Move between fields by clicking or tabbing.

   You do not need to specify the .py extension when you specify the filename of the custom script.

5. Click **OK** to save your configuration changes and close the **Properties** window.

After configuration, you can configure and run the Workload Balancing Automation and Host Power On features.

Power on a server remotely

June 3, 2019

The Host Power On feature allows you to remotely power on managed servers from XenCenter. To use this feature, the servers must have remote power control support (Wake-on-LAN functionality, a DRAC or iLO card, or a custom power-on script) and you must have enabled Host Power On in each server's **Properties** (this can be enabled once for multiple servers at pool-level).

See Configuring Host Power On to what you need to do to set up and enable this feature.

Once Host Power On has been configured, select the server(s) and then do one of the following:

1. On the **Server** menu, click **Power On**.

2. Right-click and select **Power On**.

Run in Maintenance Mode

April 25, 2022

You may need to take a managed server offline for a number of reasons, such as a rolling upgrade of virtualization software, adding or testing connectivity to a new network, diagnosing an underlying hardware issue or adding connectivity to a new storage system. In XenCenter, you can take a server offline temporarily by placing it into *Maintenance Mode*. If the server is in a resource pool, when you place it into Maintenance Mode, all running VMs will be automatically migrated from it to another server in the same pool. If the server is the pool master, a new master will also be selected for the pool.

When Workload Balancing is enabled, the virtual machines running on that server will be automatically migrated to their optimal servers when available, based on Workload Balancing recommendations (performance data, your placement strategy, and performance thresholds).

While a server is Maintenance Mode, you cannot create or start any VMs on it.

To place a server in Maintenance Mode

In the **Infrastructure** view of the XenCenter resources pane (left pane), select the server and then do one of the following:

- Right-click the server name and click **Enter Maintenance Mode** on the shortcut menu.
- On the **Server** menu, click **Enter Maintenance Mode**.

When all running VMs have been successfully migrated off the server, the server's status in the Resources pane is changed to .

To take a server out of Maintenance Mode

In the **Infrastructure** view of the XenCenter resources (left pane), select the server and then do one of the following:

- Right-click the server name and click **Exit Maintenance Mode** on the shortcut menu.
- On the **Server** menu, click **Exit Maintenance Mode**.

Store Your Server Connection State

September 30, 2020

Login credentials - your username and password - for all your managed servers can be stored between XenCenter sessions and used to automatically reconnect to them at the start of each new XenCenter session. When this feature is enabled, XenCenter will remember the connection state of all your managed servers at the end of each session and attempt to restore them at the start of your next session. If a server was connected at the end of your previous session, it will be reconnected automatically without you being prompting for your server login details; if a server was disconnected at the end of your previous session, it will not be reconnected automatically.

If you choose to turn off the automatic reconnection feature, you will need to reconnect to all your managed servers each time you open XenCenter, entering your user name and password for each server.

> **Note:**
>
> Your system administrator can disable the saving of server login credentials, so this feature many not be available.

You can optionally protect your stored login credentials with a master password to ensure they remain secure; at the start of each session, you will be prompted to enter this master password before connections to your managed servers are restored.

To turn automatic reconnection on or off

1. Open the XenCenter **Options** dialog box: on the **Tools** menu, click **Options**.
2. Click on the **Save and Restore** tab, then select or clear the **Save and restore server connection state on startup** check box.

Using a master password

When you choose to store server login credentials in XenCenter, you can also set a master password that must be entered before connections to your managed servers are automatically restored. You set, remove and change the master password from the **Save and Restore** tab in the XenCenter **Options** dialog box.

Note that if you lose or forget the master password, it cannot be recovered. If this happens, you will need to connect to each managed server again and then set a new master password.

To set a master password

1. Open the XenCenter **Options** dialog box: on the **Tools** menu, click **Options**.
2. Click on the **Save and Restore** tab.
3. Ensure that the **Save and restore server connection state on startup** check box is selected.
4. Under **Master password** , select the **Require a master password** check box, then enter and confirm the password, and click **OK**. Remember that passwords are case-sensitive.

To change the master password

1. Open the XenCenter **Options** dialog box: on the **Tools** menu, click **Options**.
2. Click on the **Save and Restore** tab.
3. Under **Master password** , click **Change Master Password**.
4. Enter the existing master password, then enter and confirm the new master password, and then click **OK**.

To clear the master password

1. Open the XenCenter **Options** dialog box: on the **Tools** menu, click **Options**.
2. Click on the **Save and Restore** tab.
3. Under **Master password** , clear the **Require a master password** check box.
4. When prompted, enter and confirm the current master password, then click **OK**.

Back up and Restore a Server

June 3, 2019

You can back up a managed server to a XenServer backup file (.xbk) that can then be used to restore the server in case of hardware failure. Note that this will back up just the server itself, but not any VMs that might be running on it.

We recommend that you back up your servers frequently to enable you to recover from possible server and/or software failure. When backing up servers in XenCenter, note the following points:

- You should not create the backup on the XenServer control domain (dom0). For more information about XenServer control domains, see the *XenServer Administrator's Guide*.
- XenServer backup files may be very large.

To restore a server, you can select and restore the backup file within XenCenter, and then you need to reboot the server from the XenServer installation CD to complete the restore.

To back up your server configuration and software

1. Select the server in the **Resources** pane, then, on the Server menu, click **Back Up**.
2. Browse to locate the folder where you want to create the backup file and enter the filename, then click **Save** to begin the backup.

The backup may take some time. You can click on the **Notifications** and then **Events** to view the progress.

To restore server software and configuration from backup

1. Select the server in the **Resources** pane, then, on the **Server** menu, click **Restore From Backup**.
2. Browse to locate the backup file and then click **Open** to begin the restore.
3. On the host server, reboot to the host installation CD and select **Restore from backup**.

Remove a Server From XenCenter

June 3, 2019

Removing a managed server from XenCenter stops all managing and monitoring activities for that server. It does not affect the activities running on the server itself or remove any VMs installed on it. Removing a server simply breaks the connection between XenCenter and the server and its VMs, and removes all traces of it from XenCenter.

To remove a server, select it in the **Resources** pane, and, in the **Server** menu, click **Remove from XenCenter**.

To put a server that you removed back into the list of managed resources, you need to add it again to XenCenter in the same way as the first time you connected to it; see Add a new server.

Configuring Networking

June 3, 2019

- About XenServer Networks
- Add a New Network
- Remove a Network
- View and Change Network Properties
- Configuring NICs
- Configuring IP Addresses

About XenServer Networks

June 3, 2019

Each managed server has one or more networks. A XenServer network is a virtual Ethernet switch that may be connected to an external interface (with or without a VLAN tag) or may be entirely virtual, internal to an individual server or pool.

When XenServer is installed on a physical server, a network is created for each physical NIC on the server. The network works as a bridge between a virtual network interface on a Virtual Machine (VIF) and a physical network interface (PIF) associated with a network interface card (NIC) on the host server.

When you move a managed server into a Resource Pool, these default networks are merged so that all physical NICs with the same device name are attached to the same network. Typically, you would only need to add a new network if you wished to create an internal network, to set up a new VLAN using an existing NIC, or to create a NIC bond. You can configure up to 16 networks per managed server, or up to 8 bonded network interfaces.

Jumbo frames can be used to optimize performance of storage traffic. You can set the Maximum Transmission Unit (MTU) for a new server network in the **New Network** wizard or for an existing network in its **Properties** window, allowing the use of jumbo frames. The possible MTU value range is 1500 to 9216.

Network types

There are four different physical (server) network types to choose from when creating a new network within XenCenter.

Single-Server Private network

This is an internal network that has no association with a physical network interface, and provides connectivity only between the virtual machines on a given server, with no connection to the outside world.

Cross-Server Private network

This is a pool-wide network that provides a private connection between the VMs within a pool, but which has no connection to the outside world. Cross-server private networks combine the isolation properties of a single-server private network with the ability to span a resource pool. This enables use of VM agility features such as XenMotion live migration and Workload Balancing (WLB) for VMs with connections to cross-server private networks. VLANs provide similar functionality though unlike

VLANs, cross-server private networks provide isolation without requiring configuration of the physical switch fabric, through the use of the Generic Routing Encapsulation (GRE) IP tunneling protocol. To create a cross-server private network, the following conditions must be met:

- all of the servers in the pool must be using Open vSwitch for networking;
- the pool must have a vSwitch Controller configured that handles the initialization and configuration tasks required for the vSwitch connection (this must be done outside of XenCenter).

External network

This type of network has an association with a physical network interface and provides a bridge between virtual machines and your external network, enabling VMs to connect to external resources through the server's physical network interface card.

Bonded network

This type of network is formed by bonding two or more NICs to create a single, high-performing channel that provides connectivity between VMs and your external network. Three bond modes are supported:

- **Active-active** - In this mode, traffic is balanced between the bonded NICs. If one NIC within the bond fails, all of the host's network traffic automatically routes over the second NIC. This mode provides load balancing of virtual machine traffic across the physical NICs in the bond.
- **Active-passive (active-backup)** - Only one NIC in the bond is active; the inactive NIC becomes active if and only if the active NIC fails, providing a hot-standby capability.
- **Link Aggregation Control Protocol (LACP) Bonding** - This mode provides active-active bonding, where traffic is balanced between the bonded NICs. Unlike the active-active bond in a Linux bridge environment, LACP can load balance all traffic types. Two available options in this mode are:
  - **LACP with load balancing based on source MAC address** - In this mode, the outgoing NIC is selected based on the MAC address of the VM from which the traffic originated. Use this option to balance traffic in an environment where you have several VMs on the same host. This option is not suitable if there are fewer VIFs than NICs: as load balancing is not optimal because the traffic cannot be split across NICs.
  - **LACP with load balancing based on IP and port of source and destination** - In this mode, the source IP address, the source port number, the destination IP address, and the destination port number are used to route the traffic across NICs. This option is ideal to balance traffic from VMs and the number of NICs exceeds the number of VIFs. For example, when only one virtual machine is configured to use a bond of three NICs.

Notes:

- You must configure vSwitch as the network stack to be able to view the LACP bonding options in XenCenter and to create a new LACP bond. Also, your switches must support the IEEE 802.3ad standard.
- Active-active and active-passive bond types are available for both the vSwitch and Linux bridge.
- You can bond either two, three, or four NICs when vSwitch is the network stack, whereas you can only bond two NICs when Linux bridge is the network stack.

For more information about the support for NIC bonds in XenServer, see the *XenServer Administrator's Guide*.

Add a New Network

September 30, 2020

To create a new network in a pool or on a standalone server, use the **New Network** wizard: select the server or pool in the **Resources** pane, click on the **Networking** tab and then click **Add Network**.

To add a new external network

An external network has an association with a physical network interface card (NIC) and provides a bridge between virtual machines and your external network, enabling VMs to connect to external resources through the NIC.

1. Open the **New Network** wizard.
2. On the first page of the wizard, select **External Network** and then click **Next**.
3. Enter the name and an optional description for the new network, and then click **Next**.
4. On the **Network settings** page, configure the NIC, VLAN and MTU settings for the new network:
   a) From the **NIC** list, choose a physical network interface card (NIC).
   b) In the **VLAN** box, assign a number to the new virtual network.
   c) To use jumbo frames, set the Maximum Transmission Unit (**MTU**) to a value between 1500 to 9216.
5. Select the **Automatically add this network to new virtual machines** check box to have the new network added to any new VMs created using the **New VM** wizard.
6. Click **Finish** to create the new network and close the wizard.

To add a new single-server private network

A single-server private network is an internal network that has no association with a physical network interface, and provides connectivity only between the virtual machines on a given server, with no connection to VMs on other servers in the pool or to the outside world.

1. Open the **New Network** wizard.
2. On the first page of the wizard, select **Single-Server Private Network** and then click **Next**.
3. Enter a name and an optional description for the new network, and then click **Next**.
4. On the **Network settings** page, select the **Automatically add this network to new virtual machines** check box to have the new network added to any new VMs created using the **New VM** wizard.
5. Click **Finish** to create the new network and close the wizard.

To add a new cross-server private network

A cross-server private network is a pool-wide network that provides a private connection between the VMs within a pool, but which has no connection to the outside world. To create a cross-server private network, the following conditions must be met:

- all of the servers in the pool must be using Open vSwitch for networking;
- the pool must have a vSwitch Controller configured that handles the initialization and configuration tasks required for the vSwitch connection (this must be done outside of XenCenter).

1. Open the **New Network** wizard.
2. On the first page of the wizard, select **Cross-Server Private Network** and then click **Next**.
3. Enter a name and an optional description for the new network, and then click **Next**.
4. On the **Network settings** page, select an interface for the new network to use, from the **Management interface** list.
5. Select the **Automatically add this network to new virtual machines** check box to have the new network added to any new VMs created using the **New VM** wizard.
6. Click **Finish** to create the new network and close the wizard.

To add a new bonded network

This type of network is formed by bonding two or more NICs together to create a single, high-performing channel that provides connectivity between VMs and your external network.

> **Note:**
>
> Whenever possible, create NIC bonds as part of initial resource pool creation prior to joining additional servers to the pool or creating VMs. Doing so allows the bond configuration to be automatically replicated to servers as they are joined to the pool and reduces the number of steps required.

1. Open the **New Network** wizard.
2. On the first page of the wizard, select **Bonded Network** and then click **Next**.

3. On the **Bond Members** page, select the NICs you want to bond together. To select a NIC, select its check box in the list. Up to four NICs may be selected in this list. Clear the check box to deselect a NIC.

4. Under **Bond mode** , choose the type of bond:

   - Select **Active-active** to configure an active-active bond, where traffic is balanced between the bonded NICs and if one NIC within the bond fails, the host server's network traffic automatically routes over the second NIC.
   - Select **Active-passive** to configure an active-passive bond, where traffic passes over only one of the bonded NICs. In this mode, the second NIC will only become active if the active NIC fails, for example, if it loses network connectivity.
   - Select **LACP with load balancing based on source MAC address** to configure a LACP bond, where the outgoing NIC is selected based on MAC address of the VM from which the traffic originated. Use this option to balance traffic in an environment where you have several VMs on the same host. This option is not suitable if there are fewer virtual interfaces (VIFs) than NICs: as load balancing is not optimal because the traffic cannot be split across NICs.
   - Select **LACP with load balancing based on IP and port of source and destination** to configure a LACP bond, where the source IP address, source port number, destination IP address, and destination port number are used to allocate the traffic across the NICs. Use this option to balance traffic from VMs in an environment where the number of NICs exceeds the number of VIFs.

   Notes:

   - You must configure vSwitch as the network stack to be able to view the LACP bonding options in XenCenter and to create a new LACP bond. Also, your switches must support the IEEE 802.3ad standard.
   - Active-active and active-passive bond types are available for both the vSwitch and Linux bridge.
   - You can bond either two, three, or four NICs when vSwitch is the network stack, whereas you can only bond two NICs when Linux bridge is the network stack.

5. To use jumbo frames, set the Maximum Transmission Unit (**MTU**) to a value between 1500 to 9216.

6. Select the **Automatically add this network to new virtual machines** check box to have the new network added to any new VMs created using the **New VM** wizard.

7. Click **Finish** to create the new network and close the wizard.

For more information, see Configuring NICs.

Remove a Network

June 3, 2019

1. Select the server or pool in the **Resources** pane then click on the **Networking** tab.
2. On the **Networking** tab, select the network in the list and then click **Remove Network**.

View and Change Network Properties

June 3, 2019

To view a server's current networking configuration

Click on the **Networking** tab for a server to see all the networks currently configured on the server, with information about each one:

Name	The name of the network.
Description	(Optional) A description of the network.
NIC	The physical network interface card (NIC), NIC bond or internal virtual network used by the network.
VLAN	For external networks, this column shows the virtual LAN (VLAN) tag.
Auto	This column shows whether the network will be automatically added to any new virtual machines created using the **New VM** wizard.
Link Status	The link status of the network: connected or disconnected.
MAC	The MAC (Media Access Control) address of the network adapter (NIC). This is a unique identifier for a particular network adapter.
MTU	A Maximum Transmission Unit value between 1500 to 9216 allows the use of jumbo frames.

To change a server's networking configuration

On the XenCenter **Networking** tab, select the network and click **Properties**. In addition to the name, description, folder, tags and custom fields properties, you can also change a number of network configuration settings on the **Network Settings** tab:

Bond mode

This configuration option will appear on bonded networks only.

- Select **Active-active** to configure an active-active bond, where traffic is balanced between the bonded NICs and if one NIC within the bond fails, the host server's network traffic automatically routes over the second NIC.
- Select **Active-passive** to configure an active-passive bond, where traffic passes over only one of the bonded NICs. In this mode, the second NIC will only become active if the active NIC fails, for example, if it loses network connectivity.
- Select **LACP with load balancing based on source MAC address** to configure a LACP bond, where the outgoing NIC is selected based on MAC address of the VM from which the traffic originated. Use this option to balance traffic in an environment where you have several VMs on the same host. This option is not suitable if there are fewer virtual interfaces (VIFs) than NICs: as load balancing is not optimal because the traffic cannot be split across NICs.
- Select **LACP with load balancing based on IP and port of source and destination** to configure a LACP bond, where the source IP address, source port number, destination IP address, and destination port number are used to allocate the traffic across the NICs. Use this option to balance the traffic in an environment where the number of NICs exceeds the number of VIFs.

Notes:

- You must configure vSwitch as the network stack to be able to view the LACP bonding options in XenCenter and to create a new LACP bond. Also, your switches must support the IEEE 802.3ad standard.
- Active-active and active-passive bond types are available for both the vSwitch and Linux bridge.
- You can bond either two, three, or four NICs when vSwitch is the network stack, whereas you can only bond two NICs when Linux bridge is the network stack.

For more information, see Configuring NICs.

MTU

To use jumbo frames, set the Maximum Transmission Unit (**MTU**) to any value between 1500 to 9216.

Automatically add this network to new virtual machines

Select this check box to have the network automatically added to new VMs when they are created using the **New VM** wizard.

Configuring NICs

December 28, 2021

XenServer automatically manages NICs as needed based on the related network, virtual network interface, server network, and bond configuration. You can view the available NICs, configure NIC bonds, and dedicate NICs to a specific function from the **NICs** tab.

NIC bonding can improve server resiliency by using two or more physical NICs as if they were one. Two or more NICs may be bonded to create a single, high-performing channel that provides connectivity between VMs and your external network. Three bond modes are supported:

Active-active: This mode provides load balancing of virtual machine traffic across the physical NICs in the bond. If one NIC within the bond fails, all of the host's network traffic automatically routes over the second NIC.

Active-passive (active-backup): This mode provides failover capability. Only one NIC in the bond is active; the inactive NIC becomes active if and only if the active NIC fails.

Link Aggregation Control Protocol (LACP) Bonding: This mode provides active-active bonding, where traffic is balanced between the bonded NICs. Unlike the active-active bond in a Linux bridge environment, LACP can load balance all traffic types.

> **Note:**
>
> You must configure vSwitch as the network stack to be able to view the LACP bonding options in XenCenter and to create a new LACP bond. Also, your switches must support the IEEE 802.3ad standard. The switch must contain a separate LAG group configured for each LACP bond on the host. For more details about creating LAG groups, see *XenServer Administrator's Guide*.

When you bond separate NICs using XenCenter, a new NIC is created: this is the bond master, and the other NICs are referred to as the bonded NICs. The NIC bond can then be connected to a XenServer network to allow virtual machine traffic and server management functions to take place. You can create NIC bonds in XenCenter from the **NICs** tab or from the server's **Networking** tab (network type = Bonded Network).

Viewing available NICs

For each available NIC on a server, the following device properties are shown on the **NICs** tab:

NIC	Identifies the physical network interface card or internal virtual network.
MAC	The MAC (Media Access Control) address of the NIC.
Link Status	The connection status of the NIC: Connected or Disconnected.
Speed	The data transfer rate of the NIC.
Duplex	The duplexing mode of the NIC: full or half.
Vendor, Device	The NIC vendor and device names.
PCI Bus Path	The PCI bus path for pass-through devices.

When you add a new physical interface on a XenServer host, for example, a new Ethernet controller, it may not appear in the list on the **NICs** tab. If this happens, click **Rescan** on the **NICs** tab to force XenServer to scan for new cards.

To create a new NIC bond

1. Ensure that the NICs you want to bind together are not in use: you must shut down any VMs with virtual network interfaces using the bonded NICs prior to creating the bond. After you have created the bond, you will need to reconnect the virtual network interfaces to an appropriate network.

2. Select the server in the **Resources** pane on the left then click on the **NICs** tab and click **Create Bond**.

3. Select the NICs you want to bond together. To select a NIC, select its check box in the list. Up to four NICs may be selected in this list. Clear the check box to deselect a NIC. To maintain a flexible and secure network, You can bond either two, three, or four NICs when vSwitch is the network stack, whereas you can only bond two NICs when Linux bridge is the network stack.

4. Under **Bond mode** , choose the type of bond:

   - Select **Active-active** to configure an active-active bond, where traffic is balanced between the bonded NICs and if one NIC within the bond fails, the host server's network traffic automatically routes over the second NIC.

   - Select **Active-passive** to configure an active-passive bond, where traffic passes over only one of the bonded NICs. In this mode, the second NIC will only become active if the active NIC fails, for example, if it loses network connectivity.

   - Select **LACP with load balancing based on source MAC address** to configure a LACP bond, where the outgoing NIC is selected based on MAC address of the VM from which

© 1999–2022 Citrix Systems, Inc. All rights reserved.   80

the traffic originated. Use this option to balance traffic in an environment where you have several VMs on the same host. This option is not suitable if there are fewer virtual interfaces (VIFs) than NICs: as load balancing is not optimal because the traffic cannot be split across NICs.

- Select **LACP with load balancing based on IP and port of source and destination** to configure a LACP bond, where the source IP address, source port number, destination IP address, and destination port number are used to allocate the traffic across the NICs. Use this option to balance traffic from VMs in an environment where the number of NICs exceeds the number of VIFs.

  > **Note:**
  >
  > LACP bonding is only available for the vSwitch, whereas active-active and active-passive bonding modes are available for both the vSwitch and Linux bridge.

  For more information about the support for NIC bonds in XenServer, see the *XenServer Administrator's Guide*.

5. To use jumbo frames, set the Maximum Transmission Unit (**MTU**) to a value between 1500 to 9216.
6. To have the new bonded network automatically added to any new VMs created using the **New VM** wizard, select the check box.
7. Click **Create** to create the NIC bond and close the dialog box.

XenCenter will automatically move management and secondary interfaces from bonded NICs to the bond master when the new bond is created.

Note that a server with its management interface on a bond will not be permitted to join a pool. You will need to reconfigure the server's management interface and move it back on to a physical NIC before it can join a pool.

Deleting a NIC bond

If reverting a server to a non-bonded configuration, be aware of the following requirements:

- As when creating a bond, all virtual machines with virtual network interfaces that use the bond must be shut down prior to destroying the bond. After reverting to a non-bonded configuration, reconnect the virtual network interfaces to an appropriate network.
- Move the management interface to another NIC using the **Management interfaces** dialog box before you delete the bond, otherwise connections to the server (including XenCenter) will be dropped.

To delete a bond:

1. Select the server in the **Resources** pane on the left, then click on the **NICs** tab.
2. Click **Delete Bond**.

Dedicating a NIC to a specific function

You can assign IP addresses to NICs to dedicate a NIC to a specific function, such as storage or other types of network traffic. For more information, see Configuring IP Addresses.

Configuring IP Addresses

September 30, 2020

The NIC used as the management interface on a managed server is initially specified during XenServer installation. In XenCenter, the XenServer xe CLI, and any other management software running on separate machines all connect to the server via the IP address of the management interface.

If a server has two or more NICs, you can select a different NIC or NIC bond to use as its management interface. You can assign IP addresses to NICs (also known as configuring secondary interface) and dedicate NICs to a specific function, such as storage or other types of network traffic. Note that putting the management interface on a VLAN network is not supported.

Note also that when a new server joins a pool, the pool master's networking configuration, including network and bond information, is automatically inherited by the joining server. However, the joining server's management interface will not be changed to match the master, so you will need to reconfigure it after joining in order to use the same bond as the pool master server. Note also that a server with its management interface on a bond will not be permitted to join a pool, and so you will need to reconfigure the server's management interface and move it back on to a physical NIC before it can join a pool.

The following sections provide instructions to assign an IP address to a NIC and change the management interface for a server or pool in XenCenter, using the **Configure IP Addresses** dialog box.

To assign an IP address to a NIC

You can use XenCenter to configure a NIC an IP address to carry out a specific function, such as storage traffic. When you configure a NIC with an IP address, you will be essentially creating a secondary interface.

To maintain a flexible and secure network, you can segment network traffic by creating secondary interfaces that use a dedicated NIC, for example, to establish separate networks for server management, application production traffic, and storage traffic. In the default XenServer networking configuration, all network traffic to IP-based storage devices occurs over the NIC used for the management interface. Also, it is important to note that the secondary interfaces inherit the DNS server settings from the management interface.

Assigning an IP address to a NIC, to carry out a specific function, requires that the appropriate network configuration be in place in order to ensure the NIC is used for the desired traffic. For example, to dedicate a NIC to storage traffic, you must assign the newly created interface an IP address that (a) is on the same subnet as the storage controller, if applicable, (b) is on a different subnet than the management interface, and (c) is not on the same subnet as any other secondary interfaces. Also, the NIC, storage target, switch, and/or VLAN must be configured such that the target is only accessible over the assigned NIC. This allows use of standard IP routing to control how traffic is routed between multiple NICs within a managed server.

Perform the following tasks to assign an IP address to a NIC and create a secondary interface:

1. On the **Networking** tab for a server or pool, under **IP Address Configuration** , click **Configure**.

2. Click **Add IP address**.

3. Enter a name for the new secondary interface.

4. Choose a XenServer network from the **Network** list.

5. Configure the networking settings for the new interface:

   - To use automated Dynamic Host Configuration Protocol (DHCP) to automatically assign networking settings including the IP address, subnet mask and gateway, select **Automatically obtain network settings using DHCP**.
   - To configure networking settings manually, click **Use these settings** and enter the required values. You must enter an IP address and a subnet mask, but the gateway settings are optional.

6. To configure additional interfaces, click **Add IP address** again and repeat the configuration steps above.

7. When you have finished, click **OK** to save your configuration choices.

   > **Note:**
   >
   > If you choose to configure the network settings manually, you will be prompted to confirm your settings. Click **Reconfigure anyway** to confirm.

To remove a secondary interface

1. On the **Networking** tab for a server or pool, under **IP Address Configuration** , click **Configure**.
2. In the list of configured interfaces on the left of the dialog box, select the one you want to remove and then click **Remove this Interface**.
3. Click **OK** to save your configuration choices.

To change the management interface

1. On the **Networking** tab for a server or pool, under **IP Address Configuration** , click **Configure**.

2. On the **Primary** tab, choose a XenServer network from the **Network** list.

3. Configure the networking settings for the management interface:

   - To use automated Dynamic Host Configuration Protocol (DHCP) to automatically assign network settings including the IP address, subnet mask, gateway and DNS server, select **Automatically obtain network settings using DHCP**.
   - To configure network settings manually, click **Use these settings** and enter the required values. You must enter an IP address and a subnet mask, but the gateway and DNS server settings are optional.

4. When you have finished, click **OK** to save your configuration choices.

   > **Note:**
   >
   > If you choose to configure the network settings manually, you will be prompted to confirm your settings. Click **Reconfigure anyway** to confirm.

Changing Server Properties

September 30, 2020

Select any connected server in the Resources pane and click on the **General** tab to see its properties and current status. Click **Properties** to change the properties of a server.

General properties - Name, Description, Folder and Tags

You can change the name, description, folder, and tags for a server on the General Properties tab of the server's Properties dialog box.

- To change the server's name, enter a new name in the **Name** box.
- To change its description, enter new text in the **Description** box.
- To place the server in a folder or to move it to a different folder, click **Change** in the **Folder** box and select a folder. See Using folders for more information about using folders.
- To tag and untag the server and to create and delete tags, see Using tags.

iSCSI IQN (General tab)

The server's iSCSI IQN is used to uniquely identify it when connecting to iSCSI storage repositories (SRs). XenServer hosts support a single iSCSI initiator which is automatically created and configured

with a random IQN during host installation. The single initiator can be used to connect to multiple iSCSI targets (SRs) concurrently. For more detailed information about XenServer support for iSCSI storage, see the *XenServer Administrator's Guide*.

> **Important:**
>
> The iSCSI target (SR) and all servers in the pool must have different IQNs set. If a non-unique IQN identifier is used, data corruption can occur and/or access to the target may be denied.

To change the iSCSI IQN value for a managed server

> **Note:**
>
> Before changing a server's iSCSI IQN value, all existing SRs must be detached. Note also that changing the server IQN may make it impossible for the server to connect to new or existing SRs unless the storage target is updated appropriately.

1. Select the server in the **Resources** pane, click on the **General** tab, and then click **Properties**.
2. On the **General** tab in the **Properties** dialog box, enter the new value in the **iSCSI IQN** box.
3. Click **OK** to save your changes and close the dialog box.

Custom Fields

Custom fields allow you to add information to managed resources to make it easier to search and organize them. See Using custom fields to find out how to assign custom fields to your managed resources.

Alerts

Use this tab to configure performance alerts for the server's CPU, memory usage, and network activity. For more information, see Configuring performance alerts.

Email Options (standalone servers)

Use this tab to configure email notification for system alerts generated on a standalone server. This feature is configured at pool level for servers in a pool. See XenCenter Alerts for information on how to receive alert notifications by email.

Multipathing

Dynamic storage multipathing support is available for Fibre Channel and iSCSI storage repositories, and can be enabled via the Multipathing tab on the server's Properties dialog.

See Storage Multipathing for information about enabling and disabling multipathing.

Power On (standalone servers)

Use this tab to configure the XenServer Host Power On feature, allowing managed servers to be powered on remotely. For more information about configuring this feature, see Configuring Host Power On. For servers in a pool, this feature is configured at pool level.

Log Destination

XenServer system log messages may be stored locally on the server itself or on a remote server.

The remote server must be running a syslogd daemon in order to receive the logs and aggregate them correctly. The syslog daemon is a standard part of all flavors of Linux and Unix, and third-party versions are available for Windows and other operating systems. The remote server should also be configured to allow remote connections from the hosts in the pool, and have its firewall configured appropriately.

To view or change the XenServer log destination

1. Select the server in the **Resources** pane, click on the **General** tab, and then click **Properties**.
2. Click on the **Log Destination** tab in the **Properties** dialog box and then:
   - To store XenServer log message locally, click **Local**.
   - To store XenServer log messages on a different server, click **Remote** and enter an IP address or the hostname of a server running the syslogd daemon.
3. Click **OK** to save your changes and close the dialog box.

GPU

The GPU tab allows you to:

1. Set a GPU placement policy
2. Enable Intel GPU pass-through for Windows VMs

Placement Policy

The GPU tab allows you to set a host-wide policy to assign VMs to available GPUs to achieve either maximum density or maximum performance. Select an option based on your requirements.

The GPU tab displays **Mixed** setting only when different settings are used for different GPU groups, that is, when certain GPU groups are configured to achieve maximum density, and the rest are configured to achieve maximum performance. Customers should note that it is **not** possible to set or edit the Mixed setting using XenCenter. You should use the xe Command Line Interface if you would like to use different settings for different GPU groups.

> **Note:**
>
> GPU Virtualization is available for XenServer Enterprise Edition customers or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. The **GPU** tab is displayed when the host meets the license requirements and also has GPUs that support various virtual GPU types. For more information, see About XenServer Licensing.

Integrated GPU pass-through

When your XenServer host is connected to an Intel GPU on an Intel Server, the XenServer host's control domain (dom0) will be connected to the integrated GPU device. In such cases, the GPU will not be available for pass-through. Select **This server will not use the Integrated GPU** to disable the connection between dom0 and the GPU and reboot the host for the changes to take effect. For more information, see GPU.

> **Note:**
>
> The XenServer host's external console output (for example, VGA, HDMI, DP) will not be available after disabling the connection between dom0 and the GPU.

Changing the Control Domain Memory

September 30, 2020

The Control Domain, also known as 'Domain0' or 'dom0' is a secure, privileged Linux Virtual Machine (VM) that runs the XenServer management toolstack (xapi). Besides providing XenServer management functions, the Control Domain also runs the driver stack that provides user-created VMs access to physical devices. The amount of memory allocated to the Control Domain is set automatically during XenServer installation and is based on the amount of physical memory on the server. For more information, see the *XenServer 7.1 CU2 Administrator's Guide*.

In Storage Read Caching and PVS-Accelerator scenarios, or when running more than 50 VMs per XenServer host, customers may wish to increase the memory allocated to dom0. On the other hand, on servers with a smaller amount of memory (less than 16GB), customers may wish to reduce the memory allocated to dom0. The following section provides instructions to update the dom0 memory using XenCenter. For information about changing the dom0 memory using the xe CLI, refer to the *XenServer 7.1 CU2 Administrator's Guide*.

Notes:

- Citrix recommends that you do not reduce the dom0 memory below 400MB.
- Increasing the amount of dom0 memory will result in less memory being available to VMs.

- Customers cannot use XenCenter to reduce dom0 memory below the value that was initially set during XenServer installation.

To update the dom0 memory

> **Note:**
>
> You must place the server in Maintenance mode before updating the dom0 memory. For more information, see Run in Maintenance Mode.

1. Select the server in the Resources pane and click **Memory**. The Memory tab displays information about the memory currently used by the server, available memory, Control domain memory, total memory and the percentage of the total memory used by the server.
2. Click the hyperlink displayed next to **Control domain memory**. Alternatively, from the **Server** menu, select **Control Domain Memory**.
3. Update the memory allocated to dom0 on the **Control Domain Memory Settings** dialog. Note that any change to the dom0 memory will cause the server to reboot.
4. Click **OK** to confirm the changes and reboot the server.

Exporting and Importing a List of Managed Servers

September 30, 2020

You can export your list of managed servers from XenCenter to a configuration file which can then be imported into a XenCenter session running on another computer. This could be useful, for example, to copy your list of managed servers from your desktop computer to a laptop, avoiding having to manually adding a long list of servers on the new machine.

The IP address/DNS name, port and display name of each managed VM is saved in XML format in a XenCenter configuration file with a .config file extension. Your login credentials are not stored.

To export your list of managed servers

1. On the **File** menu, click **Export Server List**.
2. Specify the name and location of the export file and then click **Save**.

To import a list of servers

1. On the **File** menu, click **Import Server List**.

2. Locate the XenCenter configuration file and then click **Open**.

   The servers appear in the XenCenter **Resources** pane with a disconnected status .

---

3. Double click on each imported server in the **Resources** pane to connect to it.

Managing Pools

June 3, 2019

- About Resource Pools
- Pool Requirements
- Create a New Pool
- Add a Server to a Pool
- Remove a Server From a Pool
- Destroy a Server from a Pool
- Export Resource Data
- Migrate Virtual Machines
- Change Pool Properties
- Change the Root Password
- Delete a Pool

About Resource Pools

June 3, 2019

Citrix XenServer pools allow you to view multiple servers and their connected shared storage as a single unified resource, enabling flexible deployment of virtual machines based on their resource needs and business priorities. A pool may contain up to 16 servers running the same version of XenServer software, at the same patch level, and with broadly compatible hardware - see Pool Requirements for details of hardware and configuration prerequisites.

One server in the pool is designated as the pool master , and provides a single point of contact for all of the servers in the pool, routing communication to other members of the pool as necessary.

If the pool master is shut down, the pool will be unavailable until the master is rebooted and back on line or until you nominate one of the other members as the new pool master. Every member of a resource pool contains all the information necessary to take over the role of master if required. On an HA-enabled pool, a new pool master is automatically nominated if the master is shut down.

Pool Requirements

September 30, 2020

A resource pool is a homogeneous or heterogeneous aggregate of one or more servers, up to a maximum of 16. Before you create a pool or join a server to an existing pool, you should make sure that the requirements identified below are satisfied for all the servers in the pool.

Hardware requirements

All of the servers in a XenServer resource pool must have broadly compatible CPUs, that is:

- The CPU vendor (Intel, AMD) must be the same on all CPUs on all servers.
- To run HVM virtual machines, all CPUs must have virtualization enabled.

Other requirements

In addition to the hardware prerequisites identified above, there are a number of other configuration prerequisites for a server joining a pool:

- It must have a consistent IP address (a static IP address on the server or a static DHCP lease). This also applies to the servers providing shared NFS or iSCSI storage.
- Its system clock must be synchronized to the pool master (for example, via NTP).
- It may not be a member of an existing resource pool.
- It may not have any running or suspended VMs or any active operations in progress on its VMs, such as shutting down or exporting; all VMs must be shut down before a server can join a pool.
- It may not have any shared storage already configured.
- It may not have a bonded management interface. You will need to reconfigure the joining server's management interface and move it back on to a physical NIC before joining the pool, then reconfigure it again once the server has successfully joined the pool; see Configuring IP Addresses.
- It must be running the same version of XenServer software, at the same patch level, as servers already in the pool.
- It must be configured with the same supplemental pack(s) as the servers already in the pool. Supplemental packs are used to install add-on software into the XenServer control domain, dom0. To prevent inconsistencies in the user experience across a pool, it is necessary to have the same supplemental packs at the same revision installed on all the servers in the pool.
- It must have the same XenServer license as the servers already in the pool. For example, you cannot add a server with XenServer Standard license to an existing resource pool containing servers with XenServer Enterprise or other licenses. You can change the license of any pool members

after joining the pool. The server with the lowest license determines the features available to all members in the pool. For more information about licensing, see About XenServer Licensing.

Homogeneous pool

A homogeneous resource pool is an aggregate of servers with identical CPUs. In addition to the pool requirements stated above, CPUs on a server joining a homogeneous resource pool must be the same (in terms of vendor, model, and features) as the CPUs on servers already in the pool.

Heterogeneous pool

XenServer enables expanding deployments over time by allowing disparate host hardware to be joined into a resource pool, known as heterogeneous resource pools. Heterogeneous resource pools are made possible by leveraging technologies in Intel (FlexMigration) and AMD (Extended Migration) CPUs that provide CPU "masking" or "leveling". These features allow a CPU to be configured to appear as providing a different make, model, or functionality than it actually does. This enables you to create pools of hosts with disparate CPUs but still safely support live migrations. Servers joining heterogeneous pools should meet the following requirements:

- the CPUs of the server joining the pool must be of the same vendor (that is, AMD, Intel) as the CPUs on servers already in the pool, though the specific type, (family, model and stepping numbers) need not be.
- the CPUs of the server joining the pool must support either Intel FlexMigration or AMD Enhanced Migration.

XenServer simplifies the support for heterogeneous pools. In XenServer 6.5 and earlier versions, a new pool member with a feature set different to that of the pool had to be CPU masked by users before it was allowed to join the pool. Starting with XenServer 7.0, servers can be added to existing resource pools, irrespective of the underlying CPU type (as long as the CPU is from the same vendor family). The pool feature set is dynamically calculated every time:

- a new server joins the pool
- a pool member leaves the pool
- a pool member reconnects following a reboot

Any change in the pool feature set does not affect VMs that are currently running in the pool. A Running VM will continue to use the feature set which was applied when it was started. This feature set is fixed at boot and persists across migrate, suspend, and resume operations. In scenarios where a pool level drops when a less-capable server joins the pool, a running VM can be migrated to any server in the pool, except the newly added server. When you attempt to move or migrate a VM to a different server within or across pools, XenServer performs migration checks to compare the VM's feature set against the feature set of the destination server. If the feature sets are found to be compatible, the VM will be

allowed to migrate. This enables the VM to move freely within and across pools, regardless of the CPU features the VM is using. If you use Workload Balancing (WLB) to choose an optimal destination server to migrate your VM, a server with incompatible feature set will not be recommended as the destination server.

> **Note:**
>
> To update a running VM to use the pool's new feature set, the VM must be powered off and then started. Rebooting the VM, for example, by clicking Reboot in XenCenter, does not cause the VM to update its feature set.

Shared pool storage

Although not a strict technical requirement for creating a resource pool, the advantages of pools (for example, running a VM on the most appropriate server and VM migration between servers) are only available if the pool has one or more shared storage repositories (SRs).

We recommend that you do not attempt to create a pool until shared storage is available. Once shared storage has been added, you can quickly move any existing VMs whose disks are in local storage into shared storage by copying them.

When a server with a shared SR becomes a pool master, this SR becomes a shared SR for the pool. If the new pool master does not have any shared storage, you will have to create a new shared SR for the pool: see Creating a New SR.

Create a New Pool

June 3, 2019

Before attempting to create a new pool, make sure that the requirements identified in Pool requirements are satisfied for all the servers that will be in the new pool.

To create a new pool

1. Open the New Pool dialog box by clicking **New Pool** on the Toolbar.
2. Enter a name for the new pool and an optional description; the name will be displayed in the Resources pane.
3. Nominate the pool master by selecting a server from the **Master** list.
4. Select more servers to place in the new pool from the **Additional members** list. All available managed servers are listed. If a server not listed, you may be able to add it to the list by clicking **Add New Server**. If a managed server is not listed, it may be because it does not satisfy one or more of the pool join requirements listed in Pool requirements.

---

5. Click **Create Pool** to create the new pool and close the dialog box.

If the pool master already has a shared storage repository (SR), this repository becomes a shared SR for the pool. If the new pool master does not have any shared storage, you will have to create a new shared SR for the pool: see Creating a New SR.

Additional pool configuration steps

To configure the new pool, use the property tabs:

1. To add shared storage to the pool, see Creating a New SR.
2. To add more servers to the pool, see Add a Server to a Pool.

Add a Server to a Pool

September 30, 2020

Before you add any new servers to a resource pool, make sure that the hardware and configuration requirements identified in Pool requirements are satisfied for the joining servers.

> **Important:**
>
> You should back up any virtual machines hosted on a server before attempting to add it to a pool.

To add a server to an existing pool

1. Select the server in the **Resources** pane, then do one of the following:
    *   Drag the selected server onto the target pool in the **Resources** pane.
    *   On the **Server** menu, click **Add to Pool** and then click on the target pool.
    *   Right-click and click **Add to Pool** on the shortcut menu, and then click on the target pool.
2. Click **OK** to confirm.

Once you have placed a server in a pool, it is shown as a pool member in the Resources pane, for example:
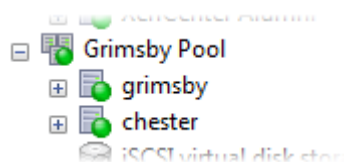
When you add a server to a pool, XenCenter will attempt to resolve any pool configuration issues if possible:

- The joining server must be licensed at the same level as the pool master. You cannot add a server to a pool whose master has a different license type. For example, in XenServer 6.5, if you add a server with Standard license edition to a pool whose master is licensed with Enterprise Edition, you will be prompted to upgrade the joining server's license to match the master's license. You cannot add the server to the pool if there are no licenses available.

  You can change the license of any pool members after joining the pool. The server with the lowest license determines the features available to all members in the pool. For more information about licensing, see About XenServer Licensing.

- If the pool master is joined to a domain, you are prompted to configure Active Directory (AD) on the server joining the pool. When you are prompted for credentials on the joining server, enter your AD credentials for the domain to which the pool is joined. These credentials must have sufficient privileges to add servers to the domain.

Note that there may be other hardware or configuration issues that will prevent a server from successfully joining a pool: see Pool requirements for details of resource pool prerequisites.

When a new server joins a pool, the pool master's networking configuration, including network and bond information, is automatically inherited by the joining server. However, the joining server's management interface will not be changed to match the master, so you will need to reconfigure it after joining in order to use the same bond as the pool master. See To change the management interface for information on how to do this.

To place a server in a new pool

You place a managed server in a new pool using the **New Pool** wizard. The server will become the master in the new pool.

1. In the **Resources** pane, select the server.
2. Right-click and, on the shortcut menu, click **Add to Pool** and then **New Pool**.
3. Create the new pool using the **New Pool** dialog box. See Create a new pool.

Remove a Server From a Pool

September 30, 2020

> **Important:**
>
> When you remove a server from a resource pool, all VM data stored on local disks (local storage) will be erased. If you have important data on local virtual disks, you must move the disks to a shared storage repository in the same resource pool before removing the server.

1. Move any data stored on local disks to a shared storage repository in the same resource pool; see Move virtual disks.
2. Shut down any VMs running on the server; see Shut Down a VM.
3. In the **Resources** pane, select the server and do one of the following:
   - Right-click and click **Remove Server from Pool** in the **Resources** pane shortcut menu.
   - In the **Pool** menu, click **Remove Server**.

Destroy a Server from a Pool

September 30, 2020

> **Important:**
>
> Destroying a server from a resource pool forgets the specified XenServer host without contacting it explicitly, and permanently removes it from the pool along with its local SRs, DVD drives and removable storage. Use this option to destroy a server that cannot be contacted or has physically failed. Also, note that the destroy server operation cannot be undone.

1. In the **Resources** pane, select the server and do one of the following:
   - Right-click and select **Destroy** in the **Resources** pane shortcut menu.
   - In the **Server** menu, click **Destroy**.
2. Click **Yes, Destroy** to confirm.

Export Resource Data

September 30, 2020

Export Resource Data allows you to generate a resource data report for your pool and export the report into a .xls or .csv file. This report provides detailed information about various resources in the pool such as, servers, networks, storage, virtual machines, VDIs, and GPUs. This feature enables administrators to track, plan, and assign resources based on various workloads such as CPU, storage, and Network.

> **Note:**
>
> Export Resource Data is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. To learn more about XenServer licensing, see About XenServer Licensing.

To export resource data:

1. In the XenCenter Navigation pane, click **Infrastructure** and then click on the pool.

2. From the XenCenter menu, click **Pool** and then select **Export Resource Data**.

3. Browse to a location where you would like to save report and then click **Save**.

Resource data

The section lists the resources and various types of resource data included in the report.

Server

- Name
- Pool Master
- UUID
- Address
- CPU Usage
- Network (avg/max KBs)
- Used Memory
- Storage
- Uptime
- Description

Networks

- Name
- Link Status
- MAC
- MTU
- VLAN
- Type
- Location

VDI

- Name
- Type
- UUID
- Size
- Storage
- Description

Storage

- Name
- Type
- UUID
- Size
- Location
- Description

VMs

- Name
- Power State
- Running on
- Address
- MAC
- NIC
- Operating System
- Storage
- Used Memory
- CPU Usage
- UUID
- Uptime
- Template
- Description

GPU

Note:

Information about GPUs will be available only if there are GPUs attached to your XenServer host.

- Name
- Servers
- PCI Bus Path
- UUID
- Power Usage
- Temperature
- Used Memory
- Computer Utilization

Change Pool Properties

September 30, 2020

Select any resource pool in the **Resources** pane and click on the **General** tab to see its properties and current status. Click **Properties** on the **General** tab to change the properties of a pool.

General properties - Name, Description, Folder, Tags

On the **General Properties** tab you can change the pool's name and description, place it in a folder, and manage its tags.

- To change the pool name, enter a new name in the **Name** box.
- To change its description, enter new text in the **Description** box.
- To place the pool in a folder or to move it to a different folder, click Change in the Folder box and select a folder. See Using folders for more information about using folders.
- To tag and untag the pool and to create and delete tags, see Using tags.

Custom Fields

Custom fields allow you to add information to managed resources to make it easier to search and organize them. See Using custom fields to find out how to assign custom fields to your managed resources.

Email Options

Use this tab to configure email notification for system alerts generated on any of the servers or VMs in the pool. See XenCenter Alerts for details about how to set up alert email notification.

Power On

The Power On feature allows you to configure power management preferences for servers that support power management, allowing them to be powered off and on automatically depending on the pool's total workload (via Workload Balancing).

- In the list of servers at the top of the tab, select the server(s) for which you want to configure power management.
- Under **Power On mode** , specify the Power On settings (Disabled, Wake-on-LAN, HP iLO, DRAC, or custom script) for the selected servers.
- Under **Configuration options** , specify either the IP address and credentials or key-value pairs for a host power-on script, depending on the **Power On mode** option you chose.

See Configuring Host Power On for more information on prerequisites for configuring the Host Power On feature and the different configuration options available.

🟩 GPU

This tab allows you to set a pool-wide policy to assign VMs to available GPUs to achieve either maximum density or maximum performance. Select an option based on your requirements.

The GPU tab displays **Mixed** setting only when different settings are used for different GPU groups. That is, when certain GPU groups within a pool are configured to achieve maximum density, and the rest of the GPU groups are configured to achieve maximum performance. Customers should note that it is **not** possible to set or edit the Mixed setting using XenCenter. You should use the xe Command Line Interface if you wish to use different settings for different GPU groups.

> **Note:**
>
> GPU Virtualization is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. The **GPU** tab is displayed when the pool meets the license requirements and also has GPUs that support various virtual GPU types. For more information, see About XenServer Licensing.

🔒 Security

The **Security** tab enables you to specify a security protocol that will be used for communication with the pool.

TLS 1.2 only : This option accepts communication from XenAPI clients and appliances (including third-party appliances) that can communicate with the XenServer pool using the TLS 1.2 protocol. **TLS 1.2 only** option uses ciphersuite TLS_RSA_WITH_AES_128_CBC_SHA256 .

> **Important:**
>
> Do not choose the **TLS 1.2 only** option before you ensure that all XenAPI clients and appliances that communicate with the XenServer pool are compatible with TLS 1.2.

Backwards compatibility mode (TLS 1.2 and earlier protocols) : Choose this option if you would like to allow both TLS and SSL protocols for pool-wide communication, for example, for backward compatibility reasons. This option uses the following ciphersuites as specified to stunnel:

- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_RC4_128_SHA
- TLS_RSA_WITH_RC4_128_MD5

- TLS_RSA_WITH_3DES_EDE_CBC_SHA

Live Patching

This tab allows you to enable or disable live patching. Live patching enables customers to install some Linux kernel and Xen hypervisor updates without having to reboot the hosts. It is enabled by default.

> **Note:**
>
> XenServer Live Patching is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. For more information about licensing, see About XenServer Licensing.

Change the Root Password

June 3, 2019

You can change the root password for a pool – that is, for all of the servers in a pool – by selecting the pool or any server in the pool in the **Resources** pane and clicking **Change Server Password** on the **Pool** menu or on the **Server** menu.

To change the root password of a standalone server: select the server in the **Resources** pane, and click **Password** and then **Change** from the **Server** menu.

If XenCenter is configured to save your server login credentials between sessions, the new password will be remembered; see Store your server connection state for details.

Delete a Pool

June 3, 2019

A resource pool containing only one managed server can be deleted, effectively turning that server into a standalone server.

To delete a pool, select the pool in the **Resources** pane and from the **Pool** menu, select **make into standalone server**.

Managing Storage

June 3, 2019

- About XenServer SRs
- Creating a New SR
- Removing an SR
- Reattaching an SR
- Storage Multipathing
- Storage Read Caching
- PVS-Accelerator
- Reclaiming Freed Space
- Live LUN Expansion
- Changing SR Properties
- Upgrading Older SRs

About XenServer SRs

June 20, 2019

A XenServer storage repository (SR) is a storage container on which virtual disks are stored. Both storage repositories and virtual disks are persistent, on-disk objects that exist independently of XenServer. SRs can be shared between servers in a resource pool and can exist on different types of physical storage device, both internal and external, including local disk devices and shared network storage. A number of different types of storage are available when you create a new storage repository using the **New Storage Repository** wizard and, depending on the type of storage selected, a number of advanced storage features can be configured in XenCenter, including:

- **Dynamic multipathing**. On Fibre Channel and iSCSI SRs, you can configure storage multipathing using round robin mode load balancing. See Storage Multipathing for more information.
- **Thin provisioning**. On NetApp and Dell EqualLogic SRs, you can choose the type of space management used. By default, allocated space is thickly provisioned and all virtual allocation guarantees are fully enforced on the filer, guaranteeing that virtual disks will never run out of space and consequently experience failed writes to disk. Thin provisioning allows the disks to be better utilized, as physical capacity is allocated only as a VM needs it - when it writes - allowing for over provisioning of the available storage and maximum utilization of your storage assets.
- **Reclaiming Freed Space**. On a thinly provisioned block-based SRs, you can free up unused space (for example, deleted VDIs in a LUN). The reclaimed space can then be reused by the storage repository. For more information, see Reclaiming Freed Space.
- **Live LUN Expansion**. Live LUN Expansion enables you to increase the size of the LUN without any VM downtime. For more information, see Live LUN Expansion.

When you configure a server or pool, you nominate a default SR which is used to store crash dump

---

data and images of suspended VMs, and which will be the default SR used for new virtual disks. At pool level, the default SR must be a shared SR. Any new virtual disks, crash dump files or suspended VM images created within the resource pool are stored in the pool's default SR, providing a mechanism to recover from physical server failure. For standalone servers, the default SR can be local or shared. When you add shared storage to a standalone server, the shared storage will automatically become the default SR for that server.

Note that it is possible to use different SRs for VMs, crash dump data and suspended VM using the XenServer xe CLI. See the *XenServer Administrator's Guide* for more information.

Creating a New SR

September 30, 2020

To create a new storage repository, click **New Storage** on the toolbar.

Alternatively, do one of the following:

- On the **Storage** tab for the selected pool or server, click **New SR**.
- On the Storage menu, click **New SR**.

You select the physical storage type on the first page of the **New Storage Repository** wizard and then follow the steps in the wizard as it takes you through the configuration process for that storage type. The set of available settings in the wizard depends on the storage system vendor/model you select on the first page. Click the links below to find out more about creating different types of SR.

SR creation steps

The **New Storage Repository** wizard takes you through process of creating a new SR:

1. On the **Type** page, you select the type of underlying storage:
   - NFS In NFS VHD SRs, VM images are stored as thin-provisioned VHD format files on a shared NFS target. Existing NFS servers that support NFS V4 and NFS V3 over TCP/IP can be used immediately as a storage repository for virtual disks. NFS SRs can be shared, allowing any VMs with their virtual disks in an NFS VHD storage repository to migrated between servers in the same resource pool.
   - iSCSI Software iSCSI SRs use a shared Logical Volume Manager (LVM) on a SAN attached LUN over iSCSI. iSCSI is supported using the open-iSCSI software iSCSI initiator or by using a supported iSCSI Host Bus Adapter (HBA).
   - Hardware HBA Hardware HBA SRs connect to a Fibre Channel (FC), Fibre Channel over Ethernet (FCoE) or shared Serial Attached SCSI (SAS) LUNs via an HBA. You need to carry out the configuration required to expose the LUN before running the **New Storage Repository**

wizard: the wizard will automatically probe for available LUNs and display a list of all the LUNs found.

- **SMB/CIFS Storage** SMB servers are a common form of Windows shared filesystem infrastructure, and can be used as a storage repository substrate for virtual disks. Virtual Machine images in SMB servers are stored as thinly provisioned VHD files on an SMB target.
- **StorageLink Technology** This option allows you to configure a StorageLink SR that provides native access to automated fabric/initiator and array configuration features in NetApp/IBM N Series or Dell EqualLogic PS Series arrays. The exact features available for a given StorageLink SR depend on the capabilities of the array.

  > **Important:**
  >
  > StorageLink technology has been removed from XenServer version 6.5 and higher. Applications, code, or usage that depend on StorageLink technology will not function in XenServer version 6.5 and higher.

- **Software FCoE** This option allows you to configure a Software FCoE SR. Software FCoE provides a standard framework to which hardware vendors can plug in their FCoE offload-capable drivers and get the same benefits of a hardware-based FCoE. This eliminates the need for using expensive HBAs.

  Before you use the New Storage Repository wizard to create a new Software FCoE storage, you should manually carry out the required configuration to expose a LUN to the host.
- **Window File Sharing (SMB/CIFS)** This option allows you to handle CD images stored as files in ISO format available as a Windows (SMB/CIFS) share. This type of SR can be useful for creating shared ISO libraries, for example, VM installation images.
- **NFS ISO** NFS ISO SRs handle CD images stored as files in ISO format available as an NFS share. This type of SR can be useful for creating shared ISO libraries, for example, VM installation images.

2. On the **Name** page, enter the name of the new SR. By default, a description of the SR will be generated automatically by the wizard, including a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description** check box and type in the **Description** box.
3. On the **Location** page, you enter the location of the underlying storage array and set configuration settings. The options available on this and subsequent wizard pages will depend on the type of storage you selected on the first page of the wizard.
4. Click **Finish** to create the new SR and close the wizard.

NFS Storage

September 30, 2020

In an NFS storage repository, VM images are stored as thin-provisioned VHD format files on a shared

NFS target. Existing NFS servers that support NFSv3 and NFSv4 over TCP/IP can be used immediately as a storage repository for virtual disks.

NFS SRs can be shared, allowing any VMs with their virtual disks in an NFS VHD storage repository to migrated between servers in the same resource pool.

Since virtual disks on NFS storage repositories (SRs) are created as sparse, you must ensure that there is enough disk space on the SR for all required virtual disks.

To configure an NFS SR

1. Open the **New Storage Repository** wizard: click **New Storage** on the toolbar.
2. Select **NFS** as the physical storage type, then click **Next**.
3. On the **Name** page, enter the name of the new SR. By default, a description of the SR will be generated automatically by the wizard, including a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description based on SR settings** check box and type in the **Description** box. Click **Next** to continue.
4. On the **Location** page, specify the NFS storage target details:
   - **Share Name**
     The IP address or DNS name of the server and the path. For example, server:/path where server is the DNS name or IP address of the server computer, and path is the directory used to contain the SR.
     The NFS server must be configured to export the specified path to all servers in the pool.
   - **Advanced Options**
     You can enter any additional configuration options here.
   - **NFS Version**
     Select the NFS version used by the SR.
     > **Note:**
     > If the underlying storage array does not support NFSv4, NFSv3 will be used to mount the share.
5. Click **Scan** to have the wizard scan for existing NFS SRs in the location you specified.
6. The New Storage Repository wizard lists any existing SRs which are not already attached. You can select an SR from the list and attach it as the new storage repository. click **Reattach an existing SR** and select the SR from the list, then click **Finish**.
7. If no existing SRs are found, simply click **Finish** to complete the new SR configuration and close the wizard.

Software iSCSI Storage

September 30, 2020

A Software iSCSI SR uses a shared Logical Volume Manager (LVM) on a SAN attached LUN over iSCSI. iSCSI is supported using the open-iSCSI software iSCSI initiator or by using a supported iSCSI Host Bus Adapter (HBA).

Note that dynamic multipathing support is available for iSCSI storage repositories. By default, multipathing uses round robin mode load balancing, so both routes will have active traffic on them during normal operation. You enable and disable storage multipathing in XenCenter via the Multipathing tab on the server's Properties dialog; see Storage Multipathing.

To create a software iSCSI SR

> **Note:**
>
> Before performing the following steps, ensure the iSCSI Initiator IQN is set appropriately for all hosts in the pool. For more information, see Changing Server Properties.

1. Open the **New Storage Repository** wizard: click **New Storage** on the toolbar. Alternatively:
   - On the **Storage** tab for the selected pool or server, click **New SR**.
   - On the **Storage** menu, click **New SR**.
   - In the **Resources** pane, select a server or pool then right-click and click **New SR** on the shortcut menu.
2. Select **Software iSCSI** as the physical storage type, then click **Next**.
3. On the **Name** page, enter the name of the new SR. By default, a description of the SR will be generated automatically by the wizard, including a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description** check box and type in the **Description** box. Click **Next** to continue.
4. On the **Location** page, specify the iSCSI target details:
   - Target Host: The IP address or DNS name of the iSCSI target.
   - Use CHAP:

     If the iSCSI target is configured to used CHAP authentication, select the Use CHAP check box and fill in the details:
       - CHAP User: the CHAP authentication username credential that should be applied when connecting to the target.
       - CHAP Secret: the CHAP authentication password credential that should be applied when connecting to the target.
   - Target IQN: To specify the iSCSI target IQN, click the Discover IQNs button and then choose an IQN from the Target IQN list.

> **Important:**
>
> The iSCSI target and all servers in the pool must not have the same IQN set. It is imperative that every iSCSI target and initiator have a unique IQN. If a non-unique IQN identifier is used, data corruption can occur and/or access to the target may be denied.

- Target LUN: To specify the LUN on which the storage repository will be created, click the Discover LUNs button and then choose a LUN from the Target LUN list.

  Each individual iSCSI storage repository must be contained entirely on a single LUN (Logical Unit Number), and may not span more than one LUN. If the LUN already contains an SR, you can either use the existing SR or you can replace the existing SR with a new one, destroying any data present on the disk.

5. Click **Finish** to complete the new SR configuration and close the wizard.

Hardware HBA Storage

September 30, 2020

Hardware HBA SRs connect to Fibre Channel (FC), Fibre Channel over Ethernet (FCoE) or shared Serial Attached SCSI (SAS) LUNs via an HBA. You need to carry out the configuration required to expose the LUN before running the **New Storage Repository** wizard: the wizard will automatically probe for available LUNs and display a list of all the LUNs found.

Dynamic multipathing support is available for Fibre Channel and iSCSI storage repositories. To enable storage multipathing, open the **Multipathing** tab on the server's **Properties** dialog; see Storage Multipathing.

To create a hardware HBA SR

1. Open the **New Storage Repository** wizard: click **New Storage** on the toolbar. Alternatively:

   - On the **Storage** tab for the selected pool or server, click **New SR**.
   - On the **Storage** menu, click **New SR**.
   - In the **Resources** pane, select a server or pool then right-click and click **New SR** on the shortcut menu.

2. Select **Hardware HBA** as the physical storage type and then click **Next**.

3. On the **Name** page, enter the name of the new SR. By default, a description of the SR will be generated automatically by the wizard, including a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate**

description check box and type in the **Description** box. Click **Next** to continue to the **Location** page.

4. The wizard scans for available LUNs and then displays a page listing all the LUNs found. Select a LUN from the list and click **Create**.

   > **Note:**
   >
   > A warning message will be displayed if there are existing SRs on the LUN you have selected. Review the details and click **Reattach** to use the existing SR, or click **Format** to delete the existing SR and to create a new SR. If you prefer to select a different LUN, click **Cancel** and select a LUN from the list.

5. The **Summary** page displays information about the new SR. Read the information and then click **Finish** to complete the SR creation process.

SMB Storage

September 30, 2020

SMB servers are a common form of Windows shared filesystem infrastructure, and can be used as a storage repository substrate for virtual disks. Virtual Machine images in SMB servers are stored as thinly provisioned VHD files on an SMB target. As SMB servers are shared storage repositories, VMs with their virtual disks in SMB servers can be started on any server in a resource pool and can be readily migrated between the servers.

> **Note:**
>
> SMB storage is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. To learn more about XenServer licensing, see About XenServer Licensing.

To configure an SMB SR

1. Open the **New Storage Repository** wizard: click **New Storage** on the toolbar.
2. Select **SMB/CIFS** as the physical storage type, then click **Next**.
3. On the **Name** page, enter the name of the new SR. By default, a description of the SR will be generated automatically by the wizard, including a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description based on SR settings** check box and type in the **Description** box. Click **Next** to continue.
4. On the **Location** page, specify the details of the storage target:

- **Share Name** The IP address or DNS name of the server and the path. For example, \\server\path where server is the DNS name or IP address of the server computer, and path is a folder or file name.
  The SMB server should be configured to export the specified path to all servers in the pool.
- **User name and Password** (Optional) If you would like to connect to an SMB server using a different user name, enter your login user name and password.

5. Click **Scan** to have the wizard scan for existing SMB SRs in the location you specified.

6. The New Storage Repository wizard lists any existing SRs which are not already attached. You can select an SR from the list and attach it as the new storage repository. Click **Reattach an existing SR** and select the SR from the list, then click **Finish**.

7. If no existing SRs are found, simply click **Finish** to complete the new SR configuration and close the wizard.

StorageLink Storage

September 30, 2020

> **Important:**
>
> StorageLink technology has been removed from XenServer version 6.5 and higher. Applications, code, or usage that depend on StorageLink technology will not function in XenServer version 6.5 and higher.

A StorageLink SR can provide direct access to high-performance storage, allowing the VMs in your XenServer environment to benefit from array-side enterprise storage services such as replication, de-duplication, thin provisioning, snapshots and cloning, data protection, and performance optimization.

StorageLink SRs use storage adapters to access different storage arrays on all common storage technologies including both NAS and SAN over either Fibre Channel or iSCSI. The features available on a given StorageLink SR depend on the capabilities of the underlying array. StorageLink SRs can co-exist with other SR types on the same storage array hardware, and multiple StorageLink SRs can be defined within the same resource pool.

Depending on your environment and goals, StorageLink provides three advantages:

- StorageLink SRs use the storage management software to access the storage.
- StorageLink simplifies the storage configuration process. For example, when you create a StorageLink SR from XenCenter, StorageLink automatically creates the LUN, volume, and initiator group for you. This means that you do not have to go in to the storage management software to create these objects. However, you do need to have created an aggregate before you run the XenCenter SR wizard.

- Every VDI created gets its own LUN. All StorageLink SRs use a LUN-per-VDI model where a new LUN is provisioned for each virtual disk (VDI).

By default, the XenServer block device-based storage inserts a Logical Volume manager on a disk, either a locally attached device (LVM type SR) or a SAN-attached LUN over either Fibre Channel (LVMoHBA type SR), iSCSI (LVMoISCSI type SR) or SAS (LVMoHBA type SR). When you use StorageLink to configure an SR, LUNs are directly mapped to virtual machines as VDIs. As a result, the array storage abstraction matches the VDI storage abstraction, which is helpful in environments that manage storage provisioning at an array level.

> **Note:**
>
> StorageLink SRs are not supported with Citrix Virtual Desktops.

Methods of Creating StorageLink SRs

There are two methods of creating StorageLink SRs: allowing StorageLink to create the volume (default) or creating the volume in advance. The default method of creating SRs with StorageLink may not be optimal in all environments.

StorageLink-created Volumes (Default): The standard (default) method of creating SRs with StorageLink is to select an aggregate in the SR creation wizard and then allow StorageLink to create the volume and LUN for you automatically. When you use this method, each VDI gets a LUN in a unique volume. However, in environments that want to create many VDIs, this may not be ideal since you can end up with numerous volumes. Some storage vendors have a maximum recommended limit of number of volumes per storage array. Likewise, when you create an SR with StorageLink, the volume names are auto-generated so they are not easy-to-read or meaningful, which can complicate management. This method also limits your ability to use deduplication to reduce space consumed by common data (such as operating systems or applications) in the volume.

Pre-created Volumes: The best practice method of creating SRs with StorageLink is to specify a previously created volume by selecting the Show All button in the XenCenter SR wizard. This method enables each virtual disk to get a LUN; however, you do not end up with numerous volumes with non-meaningful names.

Specifying a previously created volume when using the StorageLink SR wizard to create an SR simplifies volume management and administration. In this situation, each VDI gets a LUN in the selected volume. However, you do not end up with excessive, potentially unmanageable numbers of volumes (that is, one for each VDI). Using this method enables you to create a high number of VDIs. You can also create a more user-friendly meaningful name for the volumes, rather than the auto-generated names. This method also provides more efficient storage because it enables you to use deduplication since virtual disk images that share common data, such as the same operating system or application, may be stored in the same volume.

Creating StorageLink SRs for NetApp Arrays

When StorageLink creates an SR (and the associated LUN and volume), StorageLink automatically creates the initiator group with the Linux operating system. Manually configuring initiator groups with other operating systems is not recommended.

To create a Fibre Channel SR using StorageLink, you should have an iSCSI license on the storage during the initial discovery phase of the NetApp storage controller or be running ONTAP 8.1. If neither one of these apply, create the SR by using the standard SR creation process (without StorageLink) as described in Creating a New SR.

> **Note:**
>
> To use StorageLink to connect to a NetApp storage array, you must enable httpd on the storage array before using the XenCenter New SR wizard.

To create a StorageLink SR

> **Note:**
>
> The XenCenter procedure for configuring Storagelink (Dell EqualLogic and NetApp) SRs shown below applies to servers running XenServer version 6.0 or higher.

1. Open the **New Storage Repository** wizard: click **New Storage** on the toolbar.

   Alternatively:

   - On the **Storage** tab for the selected pool or server, click **New SR**.
   - On the **Storage** menu, click **New SR**.
   - In the **Resources** pane, select a server or pool then right-click and click **New SR** on the shortcut menu.

2. Select **StorageLink technology** as the physical storage type and then click **Next**.

3. On the **Name** page, enter the name of the new SR. By default, a description of the SR will be generated automatically by the wizard, including a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description** check box and type in the **Description** box.

   Click **Next** to continue.

4. On the **Storage Adapter** page, select one of the available storage system adapters and then click **Next** to continue.

5. On the **Storage System** page, select the storage system that will host the storage for this SR:

   a) In the **Array target** box, enter the host name or IP address of the array management console.

---

b) Under **Credentials** , enter the username and password to use for connection to the array management console.

c) Click **Discover** to scan the target array for storage systems. After discovery has completed, select a storage system from the **Storage System** list, then click **Next** to continue.

6. On the **Settings** page, define the SR settings; the set of available settings for each new SR depends on the storage system vendor/model and the configuration of the storage pool you've selected.

- **Storage pool**: Identify the storage pool within the specified storage system to use for allocating storage. On some types of storage system you can use the Show all check box to list all storage pools here
- **RAID types**: Select the level of RAID to use for the SR, as supported by the storage array
- **Provisioning type**: Set the provisioning type (Default, Thick or Thin)
- **Provisioning options**: Set the provisioning options (for example, deduplication)
- **Protocol**: Set the storage protocol used to connect the managed server to the storage (iSCSI, Fibre Channel, or Auto to have the system choose)
- **Use CHAP authentication**: If the server is configured to used CHAP authentication, select this check box and fill in the username and password to be used

7. Click **Finish** to complete the new StorageLink SR configuration and close the wizard.

To create an SR and a LUN for Fibre Channel with StorageLink

1. In the XenCenter **Resources** pane, right-click the pool, and select **New SR**.
2. Select **StorageLink technology** , and click **Next**.
3. Enter a meaningful name for the SR on the **Name** page, and click **Next**.
4. Select **NetApp/IBM N Series Storage Adapter** on the **Storage Adapter** page, and click **Next**.
5. On the **Storage System** page, do the following:
    a) On the **Array target** box, enter the IP address for the storage controller. StorageLink initially communicates with the storage over IP.
    b) In the **Credentials** section, enter the username and password for the storage and click **Discover**.
6. After the NetApp controller appears in the **Storage system** box, click **Next**.
7. On the **Settings** page, specify where and how you want the LUN configured:
    a) From the **Storage pool** box, select the aggregate where you would like to create the LUN.
        > **Note:**
        >
        > NetApp does not recommend creating the LUN on the aggr0 aggregate, as this is reserved for system files.
    b) In the **RAID type** box, select the RAID level you want the storage to use to format the LUN. RAID6 is the equivalent of the NetApp RAID-DP RAID type.

      c) In the **Provisioning type** box, select either Thick or Thin provisioning. If you want storage to allocate space as it is used, select the Thin Provisioned check box.

      d) In the **Provisioning options** box, select either **None** or **Deduplication**.

      e) In the **Protocol** box, select **Fibre Channel**.

8. Click **Finish** to create the SR and the LUN. After StorageLink creates the SR and LUN, an SR appears in the **Resource** pane.

Software FCoE Storage

September 30, 2020

Software FCoE provides a standard framework to which hardware vendors can plug in their FCoE offload-capable NIC and get the same benefits of a hardware-based FCoE. This eliminates the need for using expensive HBAs Software FCoE can be used with Open vSwitch and Linux bridge as the network backend.

Before creating a new software FCoE storage, configuration required to expose a LUN to the host must be completed manually, including configuring the FCoE fabric and allocating LUNs to your SAN's public world wide name (PWWN). After completing this configuration, available LUN should be mounted to the host's CNA as a SCSI device. The SCSI device can then be used to access the LUN as if it were a locally attached SCSI device. For information about configuring the physical switch and the array to support FCoE, refer to the documentation provided by the vendor. For detailed information about Software FCoE, refer to the *XenServer Administrator's Guide*.

> **Note:**
>
> Software FCoE can be used when using Open vSwitch and Linux bridge as the network backend.

To create a Software FCoE SR

1. Open the **New Storage Repository wizard** and click New Storage on the toolbar. Alternatively:

   - On the **Storage** tab for the selected pool or server, click **New SR**.
   - On the **Storage** menu, click **New SR**.
   - In the **Resources** pane, select a server or pool then right-click and click **New SR** on the shortcut menu.

2. Select **Software FCoE** as the Storage type and click **Next**.

3. Enter a name for the new SR. By default, a description of the SR will be generated automatically by the wizard, including a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description** check box and type in the Description box. Click **Next** to continue to the **Location page**.

4. XenCenter probes for available LUNs and displays a list of LUNs currently exposed to the host. This page also displays detailed information about the LUN such as, the size, serial, ID, NIC, etc. Choose the LUN (s) which you would like to allocate to the SR and click **Next**.

> **Note:**
>
> If the host cannot find any LUNs, an error message will be displayed. Verify your hardware configuration and retry to continue with the SR creation process.

5. Review the summary and click **Finish** to complete the SR creation process.

ISO Storage

September 30, 2020

This type of SR can be useful for creating shared ISO libraries, for example, VM installation images. Two ISO SR types are provided for handling CD images stored as files in ISO format:

- The **NFS ISO** SR type handles CD images stored as files in ISO format available as an NFS share.
- The **Windows File Sharing (SMB/CIFS)** SR type handles CD images stored as files in ISO format available as a Windows (SMB/CIFS) share.

To configure a new ISO SR

1. Open the **New Storage Repository** wizard: click **New Storage** on the toolbar.

2. Under **ISO library** , select **NFS ISO** or **Windows File Sharing (SMB/CIFS)** as the storage type, then click **Next**.

3. On the **Name** page, enter the name of the new SR. By default, a description of the SR will be generated automatically by the wizard, including a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description** check box and type in the **Description** box.

   Click **Next** to continue.

4. On the **Location** page, specify the ISO storage target details:

   - Share Name: For example, `server`:`/path` (NFS) or `\\server\sharename` (SMB/CIFS) where server is the DNS name or IP address of the server computer, and sharename or path is a folder or file name.

Use different user name (SMB SRs only):If you want to connect to an SMB server using a different user name, select this check box and then enter your login user name and password.

NFS Version (NFS SRs only): Select the NFS version that will be used by the SR.

> **Note:**
>
> If the underlying storage array does not support NFSv4, NFSv3 will be used to mount the share.

Click **Finish** to complete the new SR configuration and close the wizard.

Removing an SR

September 30, 2020

Using XenCenter, a storage repository can be removed temporarily or permanently:

- **Detach**. Detaching a storage repository breaks the association between the storage device and the pool or server, and its virtual disks become inaccessible. The contents of the virtual disks and the meta-information used by virtual machines to access the virtual disks are preserved. **Detach** can be used when you want to temporarily take a storage repository offline, for example, for maintenance. A detached SR can then be reattached; see Reattaching an SR.

- **Forget**. When you Forget an SR, the contents of the virtual disks on the SR are preserved but the information used to connect the virtual machines to the virtual disks it contains is permanently deleted. The SR is removed from the **Resources** pane.

  A Forget operation cannot be undone.

- **Destroy**. Destroying an SR deletes the contents of the SR permanently and the SR is removed from the **Resources** pane.

  A Destroy operation cannot be undone. For information on how to destroy an SR, see the *XenServer Administrator's Guide*.

> **Note:**
>
> You cannot remove a storage repository if it contains virtual disks of a currently running virtual machine.

To detach a storage repository

1. Select the SR in the **Resources** pane and click the **Storage** tab.
2. Note the virtual machines that currently have attached virtual disks on this storage repository.
3. Ensure that the virtual machines that have disks on this storage repository are not running.
4. Select the SR in the **Resources** pane and then do one of the following:
   - Right-click and click **Detach** on the **Resources** pane shortcut menu.
   - On the **Storage** menu, click **Detach**.
5. Click **Yes** to confirm.

> **Note:**
>
> While a storage repository is detached, you cannot start any virtual machines that have attached virtual disks on that storage repository.

To forget a storage repository

Certain types of storage repositories, such as iSCSI, must be detached before attempting to forget the SR.

> **Important:**
>
> An SR Forget operation cannot be undone. The information used to connect VMs to the virtual disks on the SR are permanently deleted.

Perform the following steps to forget an SR:

1. Select the SR in the Resources pane and then do one of the following:
   - Right-click and click **Forget** on the **Resources** pane shortcut menu.
   - On the **Storage** menu, click **Forget**.
2. Click **Yes, Forget** to confirm.

Reattaching an SR

June 3, 2019

A detached storage device has no association with any pool or server, but the data stored on it is preserved. When you reattach a storage repository to a managed server, you need to provide the storage configuration information in the same way as when you add a new storage repository.

To reattach a detached SR

1. Select the detached SR in the **Resources** pane and then do one of the following:
   - Right-click and click **Reattach Storage Repository** on the **Resources** pane shortcut menu.
   - On the **Storage** menu, click **Reattach Storage Repository**.
2. Enter the required configuration information in the same way as when you add a new storage repository. See:
   - NFS VHD Storage
   - Software iSCSI Storage
   - Hardware HBA Storage
   - StorageLink Storage
   - ISO Storage

3. Click **Finish** to complete the SR configuration and close the wizard.

Storage Multipathing

September 30, 2020

Dynamic multipathing support is available for Fibre Channel and iSCSI storage repositories. By default, multipathing uses round robin mode load balancing, so both routes will have active traffic on them during normal operation. You enable and disable storage multipathing in XenCenter via the **Multipathing** tab on the server's **Properties** dialog.

Before you enable multipathing:

- Verify that multiple targets are available on your storage server.
- The server must be placed in Maintenance Mode; this ensures that any running virtual machines with virtual disks in the affected storage repository are migrated before the changes are made.
- Multipathing must be configured on each host in the pool. All cabling and, in the case of iSCSI, subnet configurations must match for the corresponding NICs on each host. (For example, all NIC 3s must be configured to use the same subnet. See Configuring IP Addresses for more information.)

For more in-depth multipathing information, see the *Configuring iSCSI Multipathing Support for XenServer* guide and the *XenServer Administrator's Guide*.

To enable multipathing

1. In the **Resources** pane, select the server and then put it into **Maintenance Mode**. There is a short delay while XenCenter migrates any active virtual machines and unplugs the existing storage; if the server is a pool master, it will be disconnected and may disappear from the Resources pane temporarily while a new pool master is assigned. When the server reappears in the **Resources** pane with the icon, continue to the next step.
2. On the **General** tab, click **Properties** and then click on the **Multipathing** tab.
3. To enable multipathing, select the **Enable multipathing on this server** check box. To disable multipathing, clear the check box.
4. Click **OK** to apply the new setting and close the dialog box. There is a short delay while XenCenter saves the new storage configuration.
5. Take the server back out of Maintenance Mode: select it in the **Resources** pane, right-click, and click **Exit Maintenance Mode**.

Storage Read Caching

September 30, 2020

> **Note:**
>
> Storage Read Caching is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops.

Read Caching improves a VM's disk performance as, after the initial read from external disk, data is cached within the host's free memory. It greatly improves performance in situations where many VMs are cloned off a single base VM for example, in Citrix Virtual Desktops Machine Creation Service (MCS) environments, as it will drastically reduce the number of blocks read from disk.

This performance improvement can be seen whenever data needs to be read from disk more than once, as it gets cached in memory. This is most noticeable in the degradation of service that occurs during heavy I/O situations. For example, when a significant number of end users boot up within a very narrow time frame (boot storm), or when a significant number of VMs are scheduled to run malware scans at the same time (anti-virus storms).

> **Note:**
>
> The amount of memory assigned to the XenServer control domain (dom0) may need to be increased for the most efficient use of read caching. For instructions on increasing dom0 memory, see Changing the Control Domain Memory.

XenCenter displays the status of Read Caching on the VM's **General** tab. Note that you cannot enable or disable Read Caching using XenCenter. For instructions on using Read Caching, refer to the *XenServer Administrator's Guide*.

Read Caching is enabled by default, provided:

- The XenServer host is licensed with a XenServer Enterprise Edition or a Citrix Virtual Apps and Desktops license.
- The VM is attached to a VDI on a file-based SR such as NFS or EXT3. Note that Read Caching cannot be used with other SR types.
- The VM is created from a fast clone or a snapshot, or the VM is attached to a read-only VDI.

For detailed information about Read Caching, refer to the *XenServer Administrator's Guide*.

PVS-Accelerator

June 3, 2019

The XenServer PVS-Accelerator feature offers additional capabilities for customers using XenServer and Citrix Provisioning Services (PVS). PVS is a popular choice for image management and hosting for Citrix Virtual Apps and Desktops. With this feature, PVS read requests can now be cached on each XenServer host. In order to benefit from the PVS-Accelerator feature, customers should use XenServer with Provisioning Services 7.12 or higher. For detailed information about PVS-Accelerator, see the *XenServer 7.1 Administrator's Guide* on the Citrix Product Documentation website.

Enabling PVS-Accelerator involves a simple three-step process:

1. Install the PVS-Accelerator Supplemental Pack on XenServer
2. Configure PVS-Accelerator in XenServer
3. Complete the cache configuration in PVS

Enabling PVS-Accelerator

Customers should complete the following configuration settings in XenServer and in PVS to enable the PVS-Accelerator feature:

1. Install the PVS-Accelerator Supplemental Pack on each XenServer host in the pool. The supplemental pack is available to download from the XenServer Product Downloads page. For instructions on how to install the supplemental pack, see Installing Supplemental Packs.

2. Configure PVS-Accelerator in XenServer. This configuration can be done using XenCenter or the xe CLI.

    After installing the PVS-Accelerator Supplemental Pack, customers should add the PVS-Accelerator configuration details in XenServer. This entails adding a new PVS site and specifying the PVS Cache storage.

    The following section contains XenCenter instructions. For information about configuring PVS-Accelerator using the xe CLI, see the *XenServer 7.1 Administrator's Guide* on the Citrix Product Documentation website.

    The **PVS** tab appears at the pool-level (or host-level if there is no pool) in XenCenter after installing the *PVS-Accelerator Supplemental Pack* , and assigning a license with entitlement. The **PVS** tab displays a summary of the Read caching status for all the VMs running inside the pool.

    To configure PVS-Accelerator:

    a) Click on the pool or the standalone host and then select the **PVS** tab.
    b) Click **Configure PVS-Accelerator**.
    c) On the **PVS-Accelerator configuration** dialog, click **Add cache configuration** to add a PVS site.
        - Enter a name for the PVS site in the **Site name** field.
        - For each host in the pool, specify what cache should be used:

- When you select **Memory only** , the feature will use up to the specified cache size in the Control Domain memory. This option is only available after additional memory has been assigned to the Control Domain. For information on how to assign memory to the Control Domain, see Changing the Control Domain Memory.
- When you select a Storage Repository (SR), the feature will use up to the specified cache size on the SR. It will also implicitly use available control domain memory as a best effort cache tier.

Important:

- If neither Memory only nor an SR is specified, the read cache will not be activated.
- PVS-Accelerator has been designed to utilize either memory only, or a combination of disk and memory. Irrespective of the configuration choice, customers should increase the amount of memory allocated to the Control Domain to ensure there is no system performance degradation.
- We recommend that you allocate at least 4GB of Control Domain memory per host to avoid frequent disk accesses that cause higher read-latency and consequently degrade performance. For more information, see Changing the Control Domain Memory.
- We recommend that you allocate at least 5GB of cache space per vDisk version that will be actively used.

d) Click **OK**. The new PVS site and the chosen cache storage configuration will be added in XenServer.

3. After configuring PVS-Accelerator in XenServer, customers must complete the cache configuration for the newly created site using the Provisioning Services Console or the PowerShell SnapIn CLI. For more information, refer to the *Citrix Provisioning Documentation*. When this step is complete, you can view a list of PVS Servers configured for the newly created site by clicking **View PVS Servers** on the PVS-Accelerator configuration dialog in XenCenter.

Cache Operation

After starting a VM with PVS-Accelerator, the caching status for the VM will be displayed on the **PVS** tab and on the **General** tab of the VM. The following table lists the status messages displayed on these tabs.

PVS-Accelerator status	Description
Initialized	PVS-Accelerator has been started and is ready to cache. If the cache remains in this state when the VM has been booted, it means that the PVS Server IP addresses have not been configured correctly, or the VM is not communicating with the PVS server using its primary network interface.
Caching	PVS-Accelerator is working.
Stopped	PVS-Accelerator is not running for the VM. The cache remains in this state when the VM is not running, or when the cache is not configured sufficiently.
Incompatible Write Cache Mode	There is no caching as the VM is configured to persist changes on the PVS server. Ensure the VM type is "Production" or "Test" and the vDisk is in "Standard Image" Access mode.
Incompatible Protocol Version	The PVS Server version is incorrect. Ensure that you are using Provisioning Services 7.12 or higher.

The PVS-Accelerator functionality caches:

- **Reads** from vDisks but not writes or reads from a write cache
- **Based on image versions**. Multiple VMs will share cached blocks if they use the same image version
- Devices with **any write cache type**
- vDisks with the Access mode set to **Standard Image**. Caching does not work with vDisks set to Private Image mode
- Devices that are marked as type **Production** or Test. Devices marked as type **Maintenance** are not cached

Notes:

- PVS-Accelerator is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement.
- XenCenter displays various PVS-Accelerator performance graphs on the host-level Performance tab. The performance graphs provide detailed insight into the cache operation.

- The PVS-Accelerator feature leverages capabilities of OVS and is therefore not available on hosts that use Linux Bridge as the network backend.
- PVS-Accelerator works on the first virtual network interface (VIF) of a cached VM. Therefore, the first VIF should be used for connecting the PVS storage network for the caching to work.

Reclaiming Freed Space

September 30, 2020

The **Reclaim freed space** option in XenCenter enables you to free up unused blocks (for example, deleted virtual disks in an SR) on a LUN that is thinly provisioned by the storage array. Once released, the reclaimed space can then be reused by the array. The Reclaim freed space operation is only available for LVM-based SRs that are thinly provisioned on the array (that is, iSCSI, Fibre Channel, or Local LVM). This feature is not enabled on file-based SRs such as NFS and Ext3 as these SR types do not require a manual space reclamation operation.

To reclaim freed space:

1. Select the **Infrastructure** view and then click on the host or the pool that contains the SR.

2. Click on the **Storage** tab.

3. Select the SR from the list and then click **Reclaim freed space**.

   > **Note:**
   >
   > Reclaiming freed space is an intensive operation and can affect the performance of the storage array. You should perform this operation only when space reclamation is required on the array. Citrix recommends that you schedule this work outside the peak array demand hours.

4. Click **Yes** to confirm the operation. To view the status of this operation, click **Notifications** and then **Events**.

Live LUN Expansion

June 3, 2019

In order to fulfill capacity requirements, you may need to add capacity to the storage array to increase the size of the LUN provisioned to the XenServer host. Live LUN Expansion allows you to increase the size of the LUN and use the newly gained space without detaching the storage repository (SR) or taking the hosts or VMs in the pool offline. This feature is supported on Software iSCSI and Hardware HBA SR types.

Warning: It is not possible to shrink or truncate LUNs. Reducing the LUN size on the storage array can lead to data loss.

To expand the size of the LUN:

1. Add the extra storage to the storage array.
2. Select the **Infrastructure** view and then click on the SR.
3. Click the **Storage** tab in the Properties pane.
4. Click **Rescan**. This operation rescans the SR and any extra capacity is added and made available.

Changing SR Properties

September 30, 2020

The **Properties** dialog box allows you modify the details of your SR and manage them effectively by organizing the resources using folders, tags, custom fields, etc. It also allows you to configure alerts when the storage throughput exceeds specific limits.

To view the details of a storage repository, select a server or pool in the **Resources** pane and then click on the **Storage** tab. Information about the local and shared storage repositories, including the name, description, storage type, usage, and size of the SR will be displayed. Select an SR from the list and click **Properties**.

General

The General tab allows you to change the name and description of the SR, and manage its folder and tags:

- To change the name of the SR, enter a new name in the **Name** box.
- To change its description, enter new text in the **Description** box.
- To place the SR in a folder or to move it to a different folder, click Change in the Folder box and select a folder. See Using folders for more information about using folders.
- To tag and untag the SR and to create and delete tags, see Using tags.

Custom Fields

Custom fields allow you to define or modify any additional information about the SR. This tab helps you to search and effectively organize your storage repositories. See Using custom fields for information about assigning custom fields to your managed resources.

Alerts

The Alerts tab allows you to configure alerts when the total read and write storage throughput activity on a Physical Block Device (PBD) exceeds the specified limit. Select the **Generate storage throughput alerts** check box and set the storage throughput and time threshold that triggers the alert.

> **Note:**
>
> Physical Block Devices (PBD) represent the interface between a specific XenServer host and an attached SR. When the total read/write SR throughput activity on a PBD exceeds the threshold you have specified, alerts will be generated on the host connected to the PBD. Unlike other host alerts, this must be configured on the relevant SR.

Upgrading Older SRs

June 20, 2019

Since XenServer 5.5, LVM storage repositories have used Microsoft's VHD technology, providing support for fast cloning and snapshots.

LVM storage repositories created with older versions of XenServer can be upgraded quickly and simply to LVHD format in XenCenter and using the XenServer xe CLI. SRs that can be upgraded will be flagged with the message **This SR needs upgrading** in the **General** and **Search** tabs.

Note that the storage upgrade operation cannot be undone.

To upgrade LVM storage to LVHD format

Select the SR and do one of the following:

- In the **Resources** pane, right-click and select **Upgrade Storage Repository** on the shortcut menu.
- From the **Storage** menu, select **Upgrade**.
- Click on the message, **This SR needs upgrading** on the **General** and **Search** tabs.

Creating VMs

June 3, 2019

Topics

- About VMs and templates

About VMs and Templates

September 30, 2020

A virtual machine (VM) is a software container (sometimes called a "guest") that runs on a host physical computer and that behaves as if it were a physical computer itself. VMs consist of an operating system plus CPU, memory (RAM) and networking resources, and software applications.

A template is a virtual machine encapsulated into a file, making it possible to rapidly deploy new VMs. Each template contains installation metadata—the setup information needed to create a new VM with a specific guest operating system, and with the optimum storage, CPU, memory and virtual network configuration.

You can create new VMs in XenCenter in a number of different ways:

- The New VM wizard takes you step by step through the process of creating a new VM from a template or a snapshot, allowing you to configure operating system, CPU, storage, networking and other parameters.
- You can bypass the **New VM** wizard and create an "instant VM" based on a custom VM template that specifies all of the required VM configuration parameters. You simply select your preconfigured template in XenCenter then right-click and click Instant VM from template. This mode of unattended VM installation can be useful for deploying large numbers of identical VMs.
- You can copy (or "clone") an existing VM.
- You can import a VM that has been previously exported.

XenServer Tools

VMs in a XenServer environment may be fully virtualized (HVM) or paravirtualized :

- In HVM (hardware-assisted virtualization or Hardware Virtual Machine) mode, the VM is fully virtualized and can run at near-native processor speeds on virtualization-enabled hardware, without any modification to the guest operating system.

  HVM Linux VMs can take advantage of the x86 virtual container technologies in newer processors for improved performance. Network and Storage access from these VMs will still operate in PV mode, using the drivers built into the kernels. For information about upgrading your existing Linux VMs to versions which now operate in HVM mode, see the *Updating Linux Kernels and Guest Utilities* section in *XenServer Virtual Machine User's Guide*.

- In paravirtualized (non-HVM) mode, the guest operating system is tuned and optimized to run in a virtual environment, independent of the underlying processor capabilities. The result is better performance and greater flexibility.

  For detailed information about supported guest operating systems, see the *XenServer Virtual Machine User's Guide*.

I/O drivers (also known as Paravirtualized drivers or PV drivers) are available for Windows and Linux VMs to enhance disk and network performance. These drivers should be installed on all new VMs. Starting with XenServer 7.0, the I/O drivers can be installed and updated through the Windows Update mechanism. The I/O drivers and the Management Agent are combined together and issued as **XenServer Tools** for ease of installation. For more information, see Installing XenServer Tools. XenServer features such as VM migration and historical performance data tracking are only available on VMs that have XenServer Tools installed.

Using templates

A number of different templates are supplied with XenServer, and these contain all the various configuration settings needed to install a specific guest operating system on a new VM. You can also create your own customized templates configured with the appropriate guest operating system, memory, CPU, storage and network settings, and use them to create new VMs. See *XenServer Virtual Machine User's Guide* for a list of the templates/operating systems supported at this release, and for detailed information about the different install mechanisms on Windows and Linux.

You can view the XenServer templates supplied with the product and any custom templates that you create in the **Resources** pane.

- XenServer template
- Custom template

You can control whether or not to display the XenServer and Custom templates in the Resources pane:

- In the XenCenter Navigation pane, click **Infrastructure**.

  This displays a tree view of your managed resources in the Resources pane.

- To display standard XenServer VM templates: on the **View** menu, click **XenServer Templates** ; to hide XenServer templates, click again to remove the check mark.

- To show custom VM templates: on the **View** menu, click **Custom Templates** ; to hide custom templates, click again to remove the check mark.

Creating a New VM

September 30, 2020

The **New VM** wizard takes you through the process of creating a new virtual machine (VM), step-by-step. To start the **New VM** wizard, on the toolbar, click **New VM**.

Alternatively, do one of the following:

- Press **Ctrl+N**.
- On the **VM** menu, click **New VM**.
- Select a server in the **Resources** pane, right-click and then click **New VM** on the shortcut menu.

Using the wizard, you can configure the new VM exactly the way you want it, adjusting various configuration parameters for CPU, storage and networking resources. Depending on the VM template you choose on the first page of the wizard, you will see slightly different VM configuration options presented on subsequent pages, as the installation options presented are tailored for each guest operating system. Click **Help** , or press **F1** on any wizard page for more information on what to do.

In a XenServer environment where Role-Based Access Control (RBAC) is implemented, the **New VM** wizard will perform checks when it starts to ensure that you have a role with sufficient permissions to allow you to create new VMs. If your RBAC role does not have sufficient permissions, for example, a VM Operator or Read-only role, you will not be able to continue with VM creation. See RBAC overview for more information.

Overview of VM creation steps

The New VM wizard takes you through the following steps to create a new VM:

1. Select a template.

   The first step is to choose a VM template. Templates contains the setup information needed to create a new VM with a specific guest operating system, and with the optimum storage, CPU, memory and virtual network configuration. A number of different templates are supplied, and

you can add custom templates of your own. See Template and BIOS options for more information on this step.

2. Give the new VM a name.

Next, you give the new VM a name and, optionally, a description. VM names are not checked for uniqueness within XenCenter, so it will make it easier for you to manage different VMs if you give them meaningful, memorable names. See VM name and description for more information on this step.

3. Specify the operating system installation media.

The third step in the process is to specify the type and location of the OS installation media. Windows operating systems can be installed from an ISO library, from install media in a physical DVD/CD drive (a CD or DVD-ROM) or from network boot. Linux operating systems can be installed from a network install repository. See OS installation media for more information on this step.

4. Choose a home server.

This step is optional, but you can choose a home server for the new VM. XenServer will always attempt to start the VM on the nominated home server if it can. See Home server for more information on this step.

5. Configure CPU and memory.

For Windows VMs : You can assign virtual CPUs (vCPUs) to the new VM, specify the topology, that is, specify the number of cores per socket you would like to present to the vCPUs, and allocate memory. These settings can be adjusted later, after the new VM has been created. See CPU and memory configuration for more information on this step. **For Linux VMs** : You can assign maximum number of virtual CPUs (vCPUs) to the new VM, specify the topology, that is, specify the number of cores per socket you would like to present to the vCPUs, set the initial number of vCPUs, and allocate memory. These settings can be adjusted later, after the new VM has been created. See CPU and memory configuration for more information on this step.

6. Assign a graphics processing unit (GPU).

The New VM wizard prompts you to assign a dedicated GPU or a virtual GPU to the VM. This enables the VM to use the processing power of the GPU, providing better support for high-end 3D professional graphics applications such as CAD/CAM, GIS, and Medical Imaging applications. See GPU for more information on this step.

> **Note:**
>
> GPU Virtualization is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. For more information, see About XenServer Licensing.

7. Configure storage.

   The next step is to configure virtual disks for the new VM. A minimum of one virtual disk is auto‐
   matically configured by the wizard, and the template you select may include more. See Virtual
   storage configuration for more information on this step.

8. Configure cloud-config parameters. (CoreOS VMs only)

   If you are creating a CoreOS VM, you will be prompted to specify cloud-configuration parameters
   for the VM. See Cloud-Config Parameters for more information on this step.

9. Configure networking.

   The last step in the process of provisioning a new VM is to configure networking. You can config‐
   ure up to four virtual network interfaces on each VM. See Configure virtual network interfaces
   for more information on this step.

10. Complete new VM creation.

    On the final page of the wizard, you can review all the configuration options you have chosen.
    Select the Start VM automatically check box to have the new VM start automatically as soon as
    it is created.

VM Template and BIOS Options

September 30, 2020

A number of different VM templates are supplied with XenServer, and these can be used in differ‐
ent ways to create new VMs. Each template contains installation metadata—the setup information
needed to create a new VM with a specific guest operating system, and with the optimum storage,
CPU, memory and virtual network configuration.

For a full list of guest operating systems supported at each XenServer release, see the *XenServer Virtual
Machine User's Guide*.

You can also create your own customized templates configured with the appropriate guest operating
system, memory, CPU, storage and network settings, and use them to create new VMs; see Creating
new templates.

Select a BIOS option

A XenServer VM can be BIOS-generic or BIOS-specific:

- BIOS-generic: the VM has generic XenServer BIOS strings;
- BIOS-customized: the VM has a copy of the BIOS strings of a particular server in the pool;

- Without BIOS strings: if a VM does not have BIOS strings set when it is first started, the standard XenServer BIOS strings will be inserted into it, and the VM will become BIOS-generic.

When you create a new VM using the **New VM** wizard, you can copy the BIOS strings from an OEM server in the same pool into the new VM. This will enable you to install Reseller Option Kit (BIOS-locked) OEM versions of Windows on the VM later, if needed. The OEM server from which you copy the BIOS strings will be nominated as the home server for the new VM.

BIOS-customized VMs can be migrated, imported and exported to servers with the same BIOS strings and to servers with different BIOS strings.

> **Important:**
>
> It is your responsibility to comply with any EULAs governing the use of any BIOS-locked operating systems that you install.

VM Name and Description

June 3, 2019

Enter the name of the new VM in the **Name** box. You can choose any name you like, but a descriptive name is usually best. Although it is advisable to avoid having multiple VMs with the same name, it is not a requirement, and XenCenter does not enforce any uniqueness constraints on VM names.

It will make it easier for you to manage different VMs if you give them meaningful, memorable names, for example, a name that describes the VM's operating system (Windows 7 64-bit JA), application software (Citrix XenServer Web Self Service v1.0 (Build 9057)), or role (db-server, Outlook Server, Windows 7 Test) can help you to identify it later on.

It is not necessary to use quotation marks for names that include spaces.

You can also include a longer description of the VM on this page of the wizard (optional).

OS Installation Media

June 3, 2019

The options for OS installation media available on the **Installation Media** page of the **New VM** wizard depend on the OS/template selected on the first page of the wizard.

Options

Install from ISO library or DVD drive

Templates: Windows and Linux

Select **Install from ISO library or DVD drive** and then choose an ISO image or a DVD drive from the drop-down list.

If the ISO image you want to use is not listed here, you can click **New ISO library** and create a new ISO SR using the **New Storage Repository** wizard. After creating the new ISO SR, you will be able to select it from the list of available ISO libraries here.

Note that if there are no ISO images listed here, you will need to make the ISOs available to the server by creating an external NFS or SMB/CIFS share directory, as described in the *XenServer Virtual Machine User's Guide*.

Boot from network

Templates: Windows

Select this option to use PXE/network booting for Windows and **Other install media** templates.

Selecting this option will place network at the top of the boot order for the new VM.

Install from URL

Templates: Linux

CentOS, SUSE Linux Enterprise Server and Red Hat Linux operating systems can be installed from a network install repository. Select **Install from URL** and enter a URL which must include the server IP address and the repository path in the following form:

```
1   nfs://server/path
2   ftp://server/path
3   http://server/path
4   <!--NeedCopy-->
```

For example: nfs://10.10.32.10/SLES10 , where 10.10.32.10 is the IP of the NFS server and /SLES10 is the location of the install repository.

You can also optionally provide additional operating system boot parameters, if required.

Home Server

June 3, 2019

A home server is the server which will provide the resources for a VM in a pool. When you nominate a home server for a VM, XenServer will always attempt to start up the VM on that server if it can; if this is not possible, then an alternate server within the same pool will be selected automatically.

- To nominate a home server for the VM in the **New VM** wizard, click **Place the VM on this server** and select a server from the list.
- If you do not want to nominate a home server, click **Don't assign this VM a home server** : the VM will use the resources on the most suitable available server.

If you are creating a BIOS-customized VM, the OEM server from which you copy the BIOS strings will automatically be selected as the home server for the new VM.

You can change the home server configuration for a VM from the VM's Properties dialog box; see Change VM properties.

Workload Balancing (WLB) and Virtual GPU considerations

The following section lists scenarios when the home server nomination will not take effect:

- In pools with Workload Balancing (WLB) enabled, the nominated home server will not be used for starting, restarting, resuming, or migrating the VM. Instead, WLB nominates the best server for the VM by analyzing XenServer resource pool metrics and by recommending optimizations.
- If a VM has a virtual GPU assigned to it, the home server nomination will not take effect. Instead, the server nomination will be based on the virtual GPU placement policy set by the user.

VM CPU and Memory Allocation

September 30, 2020

When you create a new VM, you can allocate virtual CPUs, specify the number of cores-per-socket you would like to present to the VM's vCPUs, and set initial memory resources to the new VM. You can change the settings at anytime after the new VM is created.

The **vCPU hotplug** feature in XenCenter enables customers to dynamically increase the number of vCPUs assigned to a running Linux VM, without having to restart the VM.

Options

Number of vCPUs

(for Windows VMs)

Enter the number of virtual CPUs (vCPUs) you would like to allocate to the new VM.

To ensure you get the best performance out of the new VM, the number of vCPUs you assign to the VM should not exceed the number of physical CPUs on the server.

> **Note:**
>
> This value can be changed later, if needed. For more information, see Change VM properties: CPU. For information about the maximum number of vCPUs supported on a VM, refer to the *XenServer Configuration Limits* document.

Maximum number of vCPUs

(for Linux VMs)

Select the maximum number of virtual CPUs (vCPUs) you would like to allocate to the new VM from the drop-down list.

To ensure you get the best performance out of the new VM, the maximum number of vCPUs you assign to the VM should not exceed the number of physical CPUs on the server.

> **Note:**
>
> This value can be changed later, if needed. For more information, see Change VM properties: CPU.

Topology

Specify the topology for the vCPU.

By default, XenServer allocates one core per socket for each vCPU. For example, allocating 4 vCPUs will appear as 4 sockets with 1 core per socket. Click the **Topology** drop-down list to change this setting and select an option from the list.

> **Note:**
>
> The cores-per-socket setting depends on the number of sockets present on the server and the operating system installed. Some operating systems have restrictions on the number of CPUs. You should comply with the operating system requirements when setting this option.

Initial number of vCPUs

(for Linux VMs)

This option displays the initial number of vCPUs allocated to the VM. By default, this number is equal to the Maximum number of vCPUs set in the previous step. You can select from the drop-down list and modify the initial number of vCPUs allocated to the VM.

Memory

Enter the amount of memory you want to allocate to the VM.

The XenServer templates provide typical VM configurations and set reasonable defaults for the memory, based on the type of guest operating system. You should also take into account the following considerations when deciding how much memory you initially allocate to a new VM:

- The kinds of applications that will run on the VM.
- Other virtual machines that will be using the same memory resource.
- Applications that will run on the server alongside the virtual machine.

You can adjust the memory allocation after the new VM is created on the VM's **Memory** tab, where you can also enable Dynamic Memory Control (DMC) to allow dynamic reallocation of memory between VMs in the same pool. See the Help topics in the Configuring VM memory folder for details.

VM Power State Scenarios

The following table lists the three VM power states and describes their various vCPU scenarios.

VM Power State	Maximum Number of vCPUs	Initial number of vCPUs	Current number of vCPUs
Running	Cannot be increased/decreased	N/A	Can only be increased.
Shutdown	Can be increased/decreased	Can be increased/decreased	N/A
Suspended	Cannot be modified	N/A	Cannot be modified

GPU

September 30, 2020

XenCenter allows you to assign a dedicated graphics processing unit (GPU) or a virtual GPU to a new

VM during VM creation. This enables a VM to use the processing power of the GPU, providing better support for high-end 3D professional graphics applications such as CAD/CAM, GIS, and Medical Imaging applications. For detailed information, refer to the *Configuring XenServer for Graphics* guide.

XenServer supports Intel's virtual GPU; a graphics acceleration solution that requires no additional hardware. It uses the Intel Iris Pro functionality embedded in some processors, and utilizes a standard Intel GPU driver installed within the VM. Note that the motherboard must be equipped with a chipset which enables GPU functionality, for example, C226 for Xeon E3 v4 CPUs or C236 for Xeon v5 CPUs. For more information, see *Configuring XenServer for Graphics* guide. For information about supported processors, refer to the XenServer Hardware Compatibility List.

The following table provides information about GPU and virtual GPU support for various types of VMs

	GPU for Windows VMs	GPU for HVM Linux VMs	Virtual GPU for Windows VMs
AMD	Yes	No	No
Intel	Yes	No	Yes
NVIDIA	Yes	Yes	Yes

The **GPU type** drop-down list displays available GPUs, supported virtual GPU types, resolution, and the maximum number of displays per virtual GPU. Select a GPU or a virtual GPU type from the list to assign a GPU or a virtual GPU to the VM.

If you are using virtual GPU, select **Pass-through whole GPU** to allow a VM to use the full processing power of the GPU. The GPU or virtual GPU selection can be modified later, if required. For more information, see Change VM Properties.

Note:

- GPU Pass-through and Graphics Virtualization is only available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. For more information about licensing, see About XenServer Licensing.
- There is no licensing restriction to use NVIDIA GPU pass-through for HVM Linux VMs.
- When you allocate a GPU to HVM Linux VMs, the **GPU type** drop-down list displays all GPU types available on the host or the pool. However, you should note that only NVIDIA GPU pass-through is supported for HVM Linux VMs.

Enabling Intel GPU pass-through

XenServer supports GPU pass-through feature for Windows 7 and Windows 8 (32-/64-bit) VMs using an Intel integrated GPU device. This is supported on Haswell (Xeon E3-12xx v3) or newer CPUs that

134

contain an Intel integrated GPU device and have a graphics-capable chipset. For more information on the supported hardware, refer to the XenServer Hardware Compatibility List.

When using Intel GPU on Intel servers, the XenServer host's control domain (dom0) will have access to the integrated GPU device. In such cases, the GPU will not be available for pass through. Customers who wish to use the Intel GPU pass-through feature on Intel servers should disable the connection between dom0 and the GPU before passing through the GPU to the VM.

To do this:

1. Select the XenServer host on the **Resources** pane.

2. On the **General** tab, click **Properties** and then click on **GPU** on the left pane.

3. In the **Integrated GPU pass-through** section, click **This server will not use the integrated GPU**.

   This disables the connection between dom0 and the Intel integrated GPU device.

4. Click **OK**.

5. Reboot the XenServer host for the changes to take effect.

   The Intel GPU will now be visible on the GPU type drop-down list during new VM creation and on the VM's **Properties** tab.

> **Note:**
>
> The XenServer host's external console output (for example, VGA, HDMI, DP) will not be available after disabling the connection between dom0 and the GPU.

Virtual Storage Configuration

June 3, 2019

Virtual machines created using the **New VM** wizard will have at least one virtual disk and the selected VM template may also include additional virtual disks. A VM can have up to seven virtual disks as well as a virtual CD-ROM.

You can add more virtual disks, remove virtual disks, and change the size and location of virtual disks from the **Storage** page in **New VM** wizard.

Options

Use these virtual disks

Select this option to use the virtual disks listed.

- To add more virtual disks, click **Add** and specify the name, size and location (SR); see Add virtual disks.
- To delete a virtual disk, click **Delete**.
- To move a virtual disk to a different SR, select it in the list and click **Properties** , then choose an SR from the **Location** list.
- To make a virtual disk bigger or smaller, select it in the list and click **Properties** , then enter a new value in the **Size** box.
- To change the name or description of a virtual disk, select it in the list and click **Properties** , then enter the new text.

Use storage-level fast disk clone

This check box appears if any of the virtual disks in the template or snapshot you are using to create the new VM are on the same SR. Select the check box to use hardware-level cloning features for copying the disks from the template/snapshot to the new VM. Using storage-level fast disk clone allows new VMs to be created very quickly.

This option is only supported for VMs using remote NFS shared storage or local VHD-based storage.

Create a diskless VM that boots from the network

If you selected the **Boot from network** option on the OS Installation media page earlier in the wizard, you can select this option to make the new VM a diskless VM.

Cloud-Config Parameters

September 30, 2020

When you are creating a CoreOS VM, the **New VM** wizard prompts you to specify cloud-config parameters for your VM. You should create a config drive if you are planning to use XenServer to manage containers running inside the VM. The config drive provides user data for the VM instance.

> **Note:**
>
> Containers may fail to run if you do not create a config drive for the VM. For more information about containers, see Managing Containers.

By default, XenCenter includes a predefined set of parameters on the **Cloud-Config Parameters** page. You can modify these parameters based on your requirements. Refer to CoreOS documentation for detailed information about supported configuration parameters.

> **Note:**
>
> You can modify the cloud-config parameters when a VM is shut down. For more information, see
> **Cloud-Config Parameters** in Change VM Properties.

Virtual Networking Configuration

June 3, 2019

You can configure up to 4 virtual network interfaces from the **Networking** page of the **New VM** wizard.
To configure more than 4, go to the VM's **Networking** tab after it has been created and add them from
there.

By default, an automatically-created random MAC address will be used for all virtual network inter-
faces. To enter a different MAC address, click **Properties** and enter a new address in the **Virtual In-
terface Properties** dialog box, using hexadecimal characters in the form aa:bb:cc:dd:ee:ff .

- To add a new virtual network interface, click **Add**.
- To remove a virtual network interface, select it in the list and then click **Delete**.
- To change the virtual disk's physical network, MAC address or quality-of-service (QoS) priority,
  select it and then click **Properties**. See Change virtual network interface properties

You can use the Networking tab to change the VM's virtual networking configuration later, if required;
see the Help topics in the Configuring VM networking folder.

Complete New VM Creation

September 30, 2020

On the last page of the **New VM** wizard, select the **Start VM automatically** check box to ensure the
new VM starts up automatically as soon as it is installed.

The process of creating the new VM may take some time, depending on the size of the template and
the speed and bandwidth of the network connection between the destination server and XenCenter.
You can view the progress on the XenCenter status bar and on the **Events** view under **Notifications**.

> **Note:**
>
> After creating a new VM, you should install the XenServer Tools to ensure optimized I/O perfor-
> mance: see Installing XenServer Tools.

Express (unattended) VM Creation

June 3, 2019

Users who want to be able to create multiple identical VMs based on a custom VM template can bypass **New VM** wizard and use the **Quick Create** feature in XenCenter:

1. Create a custom VM template that specifies all the configuration parameters you want for your new VM(s). See Creating new templates.
2. Select your custom template in the **Resources** pane and, on the **Templates** menu, point to **Create VM From Selection** and then click **Quick Create**. Alternatively, right-click in the **Resources** pane and click **Quick Create** on the shortcut menu.

The new VM is then created and provisioned using all the configuration settings specified in your template.

Creating New Templates

June 3, 2019

You can create new custom templates in several different ways in XenCenter:

- By copying an existing template; see Copying VMs and templates.
- By converting an existing VM into a new template.
- By saving a copy of a VM snapshot as a new template.
- By importing a template that has previously been exported from an existing template or VM snapshot as an XVA file.

To convert an existing VM into a template

Note that when you create a new template using this method, the VM's disks will be copied to the new template and the original VM will no longer exist. Note that a VM that is currently in a vApp cannot be converted into a template.

1. Shut down the VM as described in Shut down a VM.
2. In the **Resources** pane, select the VM, right-click and then click Convert to Template.
3. Click **Convert** to confirm. You can view the conversion progress on the status bar at the bottom of the XenCenter window and on the **Events** view under **Notifications**.

When conversion is complete, the VM disappears from the **Resources** pane and reappears as a new custom template. The new custom template can then be used to create new VMs in the same way as any other template.

To save a copy of a snapshot as a new template

1. On the **Snapshots** tab, select the snapshot, right click and then click **Create Template from Snapshot** on the shortcut menu.
2. Enter the name of the new template and then click **OK**.

Once the new template has been successfully created, it appears as a custom template in the **Resources** pane and on the **Templates** page in the **New VM** wizard.

To import a template from an XVA file

VM templates and snapshots that have been exported as XVA files can be imported into XenCenter using **Import** wizard: on the **File** menu, click **Import** , select the XVA file containing the template on the first page of the wizard and then follow the same steps as when importing a VM from XVA; see Import VMs from XVA.

The import progress is displayed on the status bar at the bottom of the XenCenter window and also on the **Events** view under **Notifications**. The import process may take some time, depending on the size of the template and the speed and bandwidth of the network connection between XenCenter and the server where you are installing it. When the newly imported template is available, it will appear in the **Resources** pane as a custom template. The new template has the same configuration properties as the original exported template. To change its configuration properties, use the template's **Properties** window.

Copying VMs and Templates

September 30, 2020

You can create new VMs and templates by copying (cloning) an existing VM or a template. XenCenter enables you to copy VMs and templates within and across pools.

XenServer has two mechanisms for copying VMs and templates, full copy or fast clone:

- **Full copy** makes a complete copy of the VM's disks.
- **Fast clone** (Copy-on-Write) writes only modified blocks to disk, using hardware-level cloning features for copying the disks from the existing VM to the new VM. This mode is only supported for file-backed VMs. Copy-on-Write is designed to save disk space and allow fast clones, but can slightly slow down normal disk performance.

Copying a VM

Important:

- Before copying a Windows VM, you should use the Windows utility sysprep to ensure the uniqueness of the Security IDs (SIDs). Copying a VM without first taking the recommended system preparation steps can lead to duplicate SIDs and other problems. See the *XenServer Virtual Machine User's Guide* instructions on cloning VMs and running sysprep .
- If the VM you want to copy is a Windows VM, run the sysprep utility as described in the *XenServer Virtual Machine User's Guide*.
- If the VM is running, you must shut it down before you can copy it.

To copy a VM within the pool

1. Select the VM in the **Resources** pane, and on the **VM** menu, click **Copy VM**.
2. On the **Destination** page, select **Within Pool**.
3. On the **Name and Storage** page, enter the name of the new VM and (optionally) a meaningful description.
4. Select the Copy Mode: **Fast clone** or **Full copy**.
5. If you choose **Full copy** as the copy mode, select the storage repository (SR) where you want to copy the VM's virtual disks. If you are moving a VM from local to shared storage, make sure that you select a shared SR here.
6. Click **Finish**.

To copy a VM to a different pool

1. Select the VM in the **Resources** pane, and on the **VM** menu, click **Copy VM**.

2. On the **Destination** page, select **Cross-pool**. and click **Next**

3. Select a standalone server or a pool from the **Destination** drop-down list.

4. Select a server from the **Home Server** list to assign a home server for the VM and click **Next**

5. On the **Storage** page, specify a storage repository where you would like to place the virtual disks of the copied VM and click **Next**.

   - The **Place all migrated virtual disks on the same SR** option is selected by default and displays the default shared SR on the destination pool.
   - Click **Place migrated virtual disks onto specified SRs** to specify an SR from the **Storage Repository** drop-down list. This option allows you to select different SR for each virtual disk on the migrated VM.

6. On the **Networking** page, map the virtual network interfaces in the selected VM to networks in the destination pool or standalone server. Specify your options using the **Target Network** drop-down list and click **Next**.

7. Select a storage network on the destination pool that will be used for the live migration of the VM's virtual disks and then click **Next**.

> **Note:**
>
> Due to performance reasons, it is recommended that you do not use management network for copying VMs.

8. Review the configuration settings and click **Finish** to start copying the VM.

Copying a Template

To copy a template within the pool

1. Select the template in the **Resources** pane, and on the **Templates** menu, click **Copy Template**.
2. On the **Destination** page, select **Within Pool**.
3. **Name and Storage** page, enter the name of the new template and (optionally) a meaningful description.
4. Select the Copy Mode: **Fast clone** or **Full copy**.
5. If you choose **Full copy** as the copy mode, select the SR where the copied virtual disks will be placed.
6. Click **Finish.**

To copy a template to a different pool

1. Select the template in the **Resources** pane, and on the **Templates** menu, click **Copy**.

2. On the **Destination** page, select **Cross-pool** and click **Next**.

3. Select a standalone server or a pool from the **Destination** drop-down list.

4. Select a server from the **Home Server** list to assign a home server for the VM and click **Next**

5. On the **Storage** page, specify a storage repository where you would like to place the virtual disks of the copied template and click **Next**.

   - The **Place all migrated virtual disks on the same SR** option is selected by default and displays the default shared SR on the destination pool.
   - Click **Place migrated virtual disks onto specified SRs** to specify an SR from the **Storage Repository** drop-down list. This option allows you to select different SR for each virtual disk on the migrated VM.

6. On the **Networking** page, map the virtual network interfaces in the selected template to networks in the destination pool or standalone server. Specify your options using the **Target Network** drop-down list and click **Next**.

7. Select a storage network on the destination pool or standalone server that will be used to copy the template's virtual disks and then click **Next**.

> **Note:**
>
> Due to performance reasons, it is recommended that you do not use management network for copying VMs.

8. Review the configuration settings and click **Finish** to start copying the template.

Configuring VMs

June 3, 2019

Topics

- [Installing XenServer Tools](#)
- VM memory configuration
    - [About VM memory configuration](#)
    - [Dynamic Memory Control (DMC)](#)
    - [Configure DMC](#)
- VM storage configuration
    - [About virtual disks](#)
    - [Add virtual disks](#)
    - [Attach virtual disks](#)
    - [Detach virtual disks](#)
    - [Move virtual disks](#)
    - [Delete virtual disks](#)
    - [Change virtual disk properties](#)
- VM networking configuration
    - [About Virtual Network Interfaces](#)
    - [Add a virtual network interface](#)
    - [Activate/deactivate a virtual network interface](#)
    - [Remove a virtual network interface](#)
    - [Change virtual network interface properties](#)
- [Configuring Virtual GPU](#)
- [Managing Containers](#)
- [Change VM properties](#)

Installing XenServer Tools

September 30, 2020

XenServer Tools provide high performance I/O services without the overhead of traditional device emulation. XenServer Tools consists of I/O drivers (also known as Paravirtualized drivers or PV drivers) and the Management Agent.

The *I/O drivers* contain storage and network drivers, and low-level management interfaces. These drivers replace the emulated devices and provide high-speed transport between Windows and the XenServer product family software. During the installation of a Windows operating system, XenServer uses traditional device emulation to present a standard IDE controller and a standard network card to the VM. This allows Windows to complete its installation using built-in drivers, but with reduced performance due to the overhead inherent in emulation of the controller drivers.

The *Management Agent* , also known as the Guest Agent, is responsible for high-level virtual machine management features and provides full functionality to XenCenter, including quiesced snapshots.

XenServer has a simple mechanism to install and update the I/O drivers and the Management Agent. The I/O drivers and the Management Agent are combined together and issued as 'XenServer Tools' for ease of installation. See the following sections for information about installing XenServer Tools and updating the I/O drivers and the Management Agent.

Customers should install XenServer Tools on each VM to be able to perform the following operations:

- Cleanly shut down, reboot, or suspend a VM
- View VM performance data in XenCenter
- Migrate a running VM (using XenMotion or Storage XenMotion)
- Create quiesced snapshots or snapshots with memory (checkpoints), or revert to snapshots
- Adjust the number of vCPUs on a running Linux VM (Windows VMs require a reboot for this to take effect)

Important:

Running a VM without installing the XenServer Tools is not a supported configuration. While a Windows VM can function without them, performance is significantly hampered when the I/O drivers (PV drivers) are not installed. Running Windows VMs without these drivers is not supported. Some features, such as live relocation across physical hosts, will only work with the I/O drivers installed and active

Finding out the virtualization state of a VM

XenCenter reports the virtualization state of a VM on the VM's **General** tab. You can find out whether or not XenServer Tools (I/O drivers and the Management Agent) are installed, and whether the VM has the capability to install and receive updates from Windows Update. The following section lists the messages displayed in XenCenter:

I/O optimized (not optimized) - displays whether or not the I/O drivers are installed on the VM. Click on the **Install I/O drivers and Management Agent** link to install the I/O drivers from the XenServer

Tools ISO.

> **Note:**
>
> I/O drivers will be automatically installed on a Windows VM that has the ability to receive updates from Windows Update. For more information, see *Updating XenServer Tools*.

Management Agent installed (not installed) - displays whether or not the Management Agent is currently installed on the VM. Click on the **Install I/O drivers and Management Agent** link to install the Management Agent from the XenServer Tools ISO.

Able to (Not able to) receive updates from Windows Update - specifies whether the VM has the capability to receive I/O drivers from Windows Update. For more information, see *Updating XenServer Tools*.

Install I/O drivers and Management Agent - This message is displayed when the VM does not have the I/O drivers or the Management Agent installed. Click the link to install XenServer Tools. For Linux VMs, clicking the status link switches to the VM's console and loads the XenServer Tools ISO. You can then mount the ISO and manually run the installation, as described in the following section.

> **Note:**
>
> If you have a large number of VMs on your server or a pool, select the server or pool on the Resources pane and click the **Search** tab. From the **Saved Searches** list, select **VMs without XenServer Tools Installed**. This displays a list of VMs that do not have XenServer Tools installed.

Installing XenServer Tools on Windows VMs

> **Important:**
>
> Note that installing XenServer Tools will cause any media in the virtual machine's CD/DVD-drive to be ejected. Do not attempt to install XenServer Tools if the virtual machine's CD/DVD-drive is in use, for example, during OS install from CD.

1. Select the VM in the **Resources** pane, right-click, and then click **Install XenServer Tools** on the shortcut menu. Alternatively, on the VM menu, click **Install XenServer Tools**.

   Or

   On the General tab of the VM, click **Install I/O drivers and Management Agent**.

   > **Note:**
   >
   > When you install XenServer Tools on your VM, you will be installing both I/O drivers (PV drivers) and the Management Agent.

2. If AutoPlay is enabled for the VM's CD/DVD drive, installation will start automatically after a few moments. The process installs the I/O drivers and the Management Agent. Restart the VM when prompted to get your VM to an optimized state.

---

3. If AutoPlay is not enabled, the XenServer Tools installer displays the installation options. Click **Install XenServer Tools** to continue with the installation. This mounts the XenServer Tools ISO (guest-tools.iso) on the VM's CD/DVD drive.

   When prompted, select to choose what happens with the XenServer Tools ISO.

   a) Click **Run Setup.exe** to begin XenServer Tools installation. This opens the **Citrix XenServer Windows Management Agent Setup** wizard. Follow the instructions on the wizard to get your VM to an optimized state and perform any actions that are required to complete the installation process.

      > **Note:**
      >
      > When you install XenServer Tools using this method, the Management Agent will be configured to get updates automatically. However, the I/O drivers will not be updated by the Management Agent update mechanism. This is the default behavior. If you prefer to change the default behavior, follow the instructions below in step b.

   b) Alternatively, click **Open folders to view files** and then run **Setup.exe** from the CD drive. This option opens the **Citrix XenServer Windows Management Agent Setup** wizard and lets you customize the XenServer Tools installation and the Management Agent update settings.

      Follow the instructions on the wizard to accept the license agreement and choose a destination folder. Customize your settings on the **Installation and Updates Settings** page. The **Citrix XenServer Windows Management Agent Setup** wizard displays the following settings by default. The wizard:

      - Installs the I/O drivers
      - Allows automatic updating of the Management Agent
      - Does not allow the Management Agent to update the I/O drivers automatically.

      If you do not want to allow the automatic updating of the Management Agent, select **Disallow automatic management agent updates**. If you prefer to update the I/O drivers automatically by the Management Agent, select **Allow automatic I/O driver updates by the management agent**. If you have chosen to receive I/O driver updates through the Windows Update mechanism, we recommend that you do not allow the Management Agent to update the I/O drivers automatically .

      Click **Install** to begin the installation process. When prompted, perform any actions that are required to complete the XenServer Tools installation process and click **Finish** to exit the setup wizard.

> **Note:**
>
> - If you prefer to install the I/O drivers and the Management Agent on a large number of Win-

---

> dows VMs, install managementagentx86.msi or managementagentx64.msi using your pre-
> ferred MSI installation tool. These files can be found on the XenServer Tools ISO.
> - I/O drivers will be automatically installed on a Windows VM that has the ability to receive
>   updates from Windows Update. However, we recommend that you install the XenServer
>   Tools package to install the Management Agent, and to maintain supported configuration.

Installing XenServer Tools on Linux VMs

1. Select the VM in the **Resources** pane, right-click, and then click **Install XenServer Tools** on the
   shortcut menu. Alternatively, on the VM menu, click Install XenServer Tools.

2. Click **Install XenServer Tools** on the message dialog to go to the VM's console.

3. As the root user, mount the image into the VM:

```
1  mount -o ro,exec /dev/disk/by-label/XenServer\\x20Tools /mnt
2  <!--NeedCopy-->
```

> **Note:**
>
> If mounting the image fails, you can locate the image by running the command

```
1  blkid -t LABEL="XenServer Tools"
2  <!--NeedCopy-->
```

4. Execute the installation script as the root user:

```
1  /mnt/Linux/install.sh
2  <!--NeedCopy-->
```

5. Unmount the image from the guest by running the command:

```
1  umount /mnt
2  <!--NeedCopy-->
```

6. If the kernel has been upgraded, or the VM was upgraded from a previous version, reboot the
   VM now.

Note that CD-ROM drives and ISOs attached to Linux VMs appear as /dev/xvdd (or /dev/sdd in Ubuntu
10.10 and later), rather than /dev/cdrom . This is because they are not true CD-ROM devices, but normal

devices. When the CD is ejected by XenCenter, it hot-unplugs the device from the VM and the device disappears. This is different from Windows VMs, where the CD remains in the VM in an empty state

Updating XenServer Tools

XenServer has a simpler mechanism to automatically update I/O drivers (PV drivers) and the Management Agent for Windows VMs. This enables customers to install updates as they become available, without having to wait for a hotfix.

The **Virtualization state** section on a VM's **General** tab specifies whether or not the VM is able to receive updates from Windows Update. The mechanism to receive I/O driver updates from Windows Update is turned on by default. If you do not want to receive I/O driver updates from Windows Update, you should disable Windows Update on your VM, or specify a group policy.

> **Important:**
>
> Ensure that all requested VM restarts are completed as part of the update. Multiple restarts might be required. If all requested restarts are not completed, this might result in unexpected behavior.

Updating the I/O drivers

If you are running newly created Windows VMs on XenServer 7.0 or higher, you will be able to get updates to I/O drivers automatically from Microsoft Windows Update, provided:

- You are running XenServer with Enterprise Edition, or have access to XenServer through Citrix Virtual Apps and Desktops entitlement
- You have created a Windows VM using XenCenter issued with XenServer 7.0 or higher
- Windows Update is enabled within the VM
- You have access to the internet, or are able to connect to a WSUS proxy server

> **Note:**
>
> Customers can also receive I/O driver updates automatically through the automatic Management Agent update mechanism. See *Updating the Management Agent* for details.

Updating the Management Agent

XenServer enables you to automatically update the Management Agent on both new and existing Windows VMs. By default, XenServer allows the automatic updating of the Management Agent. However, it does not allow the Management Agent to update the I/O drivers automatically. You can customize the Management Agent update settings during XenServer Tools installation. See *Installing XenServer Tools on Windows VMs* for details. The automatic updating of the Management Agent occurs seamlessly, and does not reboot your VM. In scenarios where a VM reboot is required, XenCenter issues notification to users about the required action.

To update the Management Agent automatically:

- You must be running XenServer 7.0 or higher with Enterprise Edition, or have access to XenServer through Citrix Virtual Apps and Desktops entitlement.
- You must have installed XenServer Tools issued with XenServer 7.0 or higher
- The Windows VM must be connected to the internet

Important:

- The ability to receive I/O drivers from Windows Update and the automatic updating of the Management Agent features are available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement.
- Updates to XenServer Tools can also be issued through the standard XenServer update (hotfix) mechanism. Such hotfixes contain updates to both I/O drivers and the Management Agent. There is no licensing restriction to update XenServer Tools issued as a hotfix.

Configuring VM Memory

June 3, 2019

Topics

- About VM Memory Configuration
- Dynamic Memory Control (DMC)
- Configure DMC

About VM Memory Configuration

June 3, 2019

When a VM is first created, it is allocated a fixed amount of memory. To improve the utilization of physical memory in your XenServer environment, you can use Dynamic Memory Control (DMC) , a memory management feature that enables dynamic reallocation of memory between VMs.

The **Memory** tab in XenCenter shows memory usage and configuration information for your VMs and servers.

- For servers, the total available memory and the current memory usage is shown, and you can see how memory is divided between hosted VMs.

- For VMs, in addition to current memory usage you can also see the VM's memory configuration information, that is, if DMC is enabled and the current dynamic minimum and maximum values, and you can edit DMC configuration settings.

VMs with the same memory configuration are grouped together on the **Memory** tab, enabling you to view and configure memory settings easily for individual VMs and for groups of VMs.

Dynamic Memory Control (DMC)

September 30, 2020

XenServer DMC (sometimes known as "dynamic memory optimization", "memory overcommit" or "memory ballooning") works by automatically adjusting the memory of running VMs, keeping the amount of memory allocated to each VM between specified minimum and maximum memory values, guaranteeing performance and permitting greater density of VMs per server.

Without DMC, when a server is full, starting further VMs will fail with "out of memory" errors. To reduce the existing VM memory allocation and make room for more VMs you must edit each VM's memory allocation and then reboot the VM. With DMC enabled, even when the server is full, XenServer will attempt to reclaim memory by automatically reducing the current memory allocation of running VMs within their defined memory ranges.

Dynamic and static memory range

For each VM, you can set a dynamic memory range which is the range within which memory can be added/removed from the VM without requiring a reboot. You can adjust the dynamic range while the VM is running, without having to reboot it. XenServer always guarantees to keep the amount of memory allocated to the VM within the dynamic range. For example, if the Dynamic Minimum Memory was set at 512 MB and the Dynamic Maximum Memory was set at 1024 MB, this would give the VM a Dynamic Memory Range (DMR) of 512 - 1024 MB, within which it would operate. With DMC, XenServer guarantees at all times to assign each VM memory within its specified DMR.

When the host server's memory is plentiful, all running VMs will receive their Dynamic Maximum Memory level; when the host's memory is scarce, all running VMs will receive their Dynamic Minimum Memory level. If new VMs are required to start on "full" servers, running VMs have their memory "squeezed" to start new ones. The required extra memory is obtained by squeezing the existing running VMs proportionally within their pre-defined dynamic ranges.

Many Operating Systems that XenServer supports do not fully support dynamically adding or removing memory. As a result, XenServer must declare the maximum amount of memory that a VM will ever be asked to consume at the time that the VM boots to allow the guest operating system to size its

page tables and other memory management structures accordingly. This introduces the concept of a static memory range within XenServer. The static memory range cannot be adjusted while the VM is running, and the dynamic range is constrained such as to be always contained within this static range until the VM is next rebooted. Note that the static minimum (the lower bound of the static range) is there to protect the administrator and is set to the lowest amount of memory that the OS can run with on XenServer.

> **Important:**
>
> Citrix advises you not to change the static minimum level as this is set at the supported level per operating system: refer to the *XenServer Administrator's Guide* for more details. By setting a static maximum level higher than a dynamic max, if you need to allocate more memory to a VM in the future, you can do so without requiring a reboot.

DMC memory constraints

XenCenter enforces the following constraints when setting DMC values:

- The minimum dynamic memory value cannot be lower than the static minimum memory value.
- The minimum dynamic memory value cannot be greater than the maximum dynamic memory value.
- The maximum dynamic memory value cannot be greater than the maximum static memory value.

You can change a VM's memory properties to any values that satisfy these constraints, subject to validation checks. In addition to these constraints, Citrix supports only certain VM memory configurations for specific operating system; these are detailed in the *XenServer Administrator's Guide*.

Configure DMC

June 3, 2019

To enable DMC

1. Select a VM or server in the **Resources** pane and click on the **Memory** tab.
2. Click the **Edit** button for the VM or group of VMs on which you want to configure DMC.
3. For multiple VMs with the same current memory configuration, select the VMs you want to configure and click **Next**.
4. Click the **Automatically allocate memory within this range** option and then set the required maximum and minimum dynamic memory range values using the slider or by typing the values directly.

---

5. Click **OK** to apply the changes and close the dialog box.

To disable DMC

1. Select the VM or host server in the **Resources** pane and click on the **Memory** tab.
2. Click the **Edit** button for the VM or group of VMs you want to configure.
3. For multiple VMs with the same current memory configuration, select the VMs you want to configure and click **Next**.
4. Click the **Set a fixed memory** option and specify the amount of memory to allocate.
5. Click **OK** to apply the changes and close the dialog box.

Configuring Virtual Storage

June 3, 2019

Select a VM in the **Resources** pane and then click on the **Storage** tab to view and change its virtual storage settings.

Topics:

- About Virtual Disks
- Add Virtual Disks
- Attach Virtual Disks
- Detach Virtual Disks
- Move Virtual Disks
- Delete Virtual Disks
- Change Virtual Disk Properties

About Virtual Disks

June 3, 2019

Storage on XenServer VMs is provided by virtual disks . A virtual disk is a persistent, on-disk object that exists independently of the VM to which it is attached. Virtual disks are stored on XenServer Storage Repositories (SRs), and can be attached, detached and re-attached to the same or different VMs when needed. New virtual disks can be created at VM creation time (from within the **New VM** wizard) and they can also be added after the VM has been created from the VM's **Storage** tab.

Virtual disks on paravirtualized VMs (that is, VMs with XenServer Tools installed) can be "hotplugged", that is, you can add, delete, attach and detach virtual disks without having to shut down the VM first.

VMs running in HVM mode (without XenServer Tools installed) must be shut down before you carry out any of these operations; to avoid this, you should install XenServer Tools on all HVM virtual machines.

On the VM's **Storage** tab in XenCenter, you can:

- Add new virtual disks.
- Configure virtual disks - change a virtual disk's size, location, read/write mode and other configuration parameters.
- Attach existing virtual disks to the VM.
- Detach virtual disks, preserving the virtual disk and all the data on it.
- Move a virtual disk to a specified storage repository.
- Delete virtual disks, permanently destroying the disk and any data stored on it.

Add Virtual Disks

September 30, 2020

To add a new virtual disk, you use the **Add Virtual Disk** dialog box.

> **Important:**
>
> If the VM is running in HVM mode (without XenServer Tools installed), you will need to shut it down before you can add any virtual disks; to avoid this, you should install XenServer Tools on all HVM virtual machines.

Procedure:

1. Open the **Add Virtual Disk** dialog box by doing any of the following:
   - Select the VM or an SR in the **Resources** pane, click on the **Storage** tab and then click **Add**.
   - On the **Storage** menu, click **Virtual Disks** and then **New Virtual Disk**.
   - On the **Storage** page of the **New VM** wizard, click **Add**.
2. Enter the name of the new virtual disk and, optionally, a description.
3. Enter the size of the new virtual disk. You should make sure that the storage repository (SR) on which the virtual disk will be stored has sufficient space for the new virtual disk.
4. Select the SR where the new virtual disk will be stored.
5. Click **Create** to add the new virtual disk and close the dialog box.

Attach Virtual Disks

September 30, 2020

You can add storage to a VM by attaching an existing virtual disk.

1. Select the VM in the **Resources** pane, click on the **Storage** tab, and then click **Attach**. Alternatively, on the **Storage** menu, click **Virtual Disks** then **Attach Virtual Disk**.

2. Select a virtual disk from the list.

3. To set access to the virtual disk to read-only, select the **Attach as read-only** check box. This can help prevent data from being overwritten or changed when the disk is accessed by more than one VM and allows the virtual disk to be attached to many VMs. To allow write access to the virtual disk, clear the check box.

4. Click **Attach**.

   Tip:

   Problems on an underlying SR can sometimes cause an attached virtual disk to become deactivated ("unplugged"). If this happens, you should be able to activate it again from the VM's **Storage** tab by selecting it and clicking **Activate**.

Detach Virtual Disks

June 3, 2019

When you detach a virtual disk from a VM, the virtual disk and all the data on it are preserved but the virtual disk is no longer available to the VM; the detached storage device can then subsequently be re-attached to the same VM, attached to a different VM, or moved to a different storage repository (SR).

You can detach a virtual disk without shutting down the VM ("hot unplug") if all of the following conditions are met:

- The VM may not be suspended.
- The VM must have XenServer Tools installed.
- The virtual disk may not be a system disk.
- The virtual disk must be deactivated in order to be able to detach it cleanly. Note that the term "deactivate" is equivalent to "unplug", which is the term used for this operation in the *XenServer Administrator's Guide* and in the CLI.

If any of these conditions are not satisfied, you will have to shut the VM down before you can detach the virtual disk.

Procedure:

1. Select the VM in the **Resources** pane and then click on the **Storage** tab.
2. Select the virtual disk in the list, click **Deactivate** and then click **Detach**.
3. Click **OK** to confirm the operation.

Move Virtual Disks

September 30, 2020

Virtual disks can be moved or migrated from one storage repository (SR) to a different SR within the same pool. This includes:

- Virtual disks that are not currently attached to any VM.
- Virtual disks attached to VMs that are not running.
- Virtual disks attached to running VMs (using Storage XenMotion)

Note that you can move a virtual disk on local storage to shared storage on a different server, but you cannot move it to a local storage on a different server.

About Storage XenMotion

Storage XenMotion allows you to move virtual disks without having to shut down the VM first, enabling administrative operations such as:

- moving a VM from cheap local storage to fast, resilient, array-backed storage.
- moving a VM from a development to production environment.
- moving between tiers of storage when a VM is limited by storage capacity.
- performing storage array upgrades.

Note the following limitations for Storage XenMotion:

- Virtual disks with more than one snapshot cannot be migrated.

- Virtual disks located on StorageLink SRs cannot be migrated.

  Important:

  StorageLink technology has been removed from XenServer version 6.5 and higher. Applications, code, or usage that depend on StorageLink technology will not function in XenServer version 6.5 and higher.

To move a virtual disk

1. In the XenCenter **Resources** pane, select the SR where the virtual disk is currently stored, and then click on the **Storage** tab. To locate a virtual disk:
   - In the XenCenter **Resources** pane, select the VM to which the virtual disk that you wish to move is attached.
   - Click on the **Storage** tab and identify the SR on which the virtual disk is currently stored.

2. From the **Virtual Disks** list, select one or more virtual disks that you would like to move, and then click **Move**. Alternatively, right-click on the selected virtual disk and select **Move Virtual Disk** from the shortcut menu.

3. In the **Move Virtual Disk** dialog box, select the target SR that you would like to move the virtual disk to. Make sure that the target SR has sufficient space for another virtual disk: the available space is shown in the list of available SRs.

4. Click **Move** to move the virtual disk.

Delete Virtual Disks

September 30, 2020

You can delete a virtual disk without shutting down the VM first if all of the following conditions are met:

- The VM may not be suspended.
- The VM must have XenServer Tools installed.
- The virtual disk may not be a system disk.
- The virtual disk must be deactivated first. Note that the term "deactivate" is equivalent to "un-plug", which is the term used for this operation in the *XenServer Administrator's Guide* and in the CLI.

If any of these conditions are not satisfied, you will have to shut the VM down before you can delete the virtual disk.

> **Important:**
>
> Deleting a virtual disk will permanently delete the disk, destroying any data stored on it.

Procedure:

1. On the VM's **Storage** tab, select the virtual disk in the list and click **Deactivate** and then **Delete**.
2. Click **OK** to confirm the deletion.

Change Virtual Disk Properties

September 30, 2020

To change the properties of a virtual disk, click on the VM's **Storage** tab, then select the virtual disk and click **Properties**.

⊞ General properties - name, description, folder, tags

Name	The virtual disk name
Description	A description of the virtual disk (optional)
Folder	The name of the resource folder where the virtual disk is located, if applicable.
Tags	A list of tags that have been applied to this virtual disk.

☰ Custom fields

On the Custom Fields tab you can assign new custom fields to a virtual disk, change the value of any existing custom fields assigned to it, and remove custom fields.

See Using custom fields for information on adding, setting, modifying and deleting custom fields.

🛢 Disk size and location

Set the size of the virtual disk on this tab and select the storage repository where the virtual disk is located.

▣ Device options

The final tab on the virtual disk **Properties** dialog box allows you to set some device options for the virtual disk.

Mode	The disk read/write permissions of a virtual disk can be changed, for example, to prevent data from being overwritten on a virtual disk that is being used for backup purposes.
Device position	The position to use for this virtual disk in the drive sequence.

Disk access priority	For some virtual disks, you may also be able to adjust the disk I/O priority (Quality of Service or QoS). This setting is only available for virtual disks on storage repositories that are LVM-based, that is local, shared iSCSI or hardware HBA. For more information about adjusting QoS settings, and about how QoS control works with different storage types, see the *XenServer Administrator's Guide*.

Configuring VM Networking

June 3, 2019

Topics:

- About Virtual Network Interfaces
- Add a Virtual Network Interface
- Activate/deactivate a Virtual Network Interface
- Remove a Virtual Network Interface
- Change Virtual Network Interface Properties

About Virtual Network Interfaces

June 3, 2019

Each virtual machine (VM) may have one or more virtual network interfaces that act as virtual NICs.

A virtual network interface has the following properties:

Network	The (physical) network location of the virtual network interface.
MAC address	The MAC address of the virtual network interface.

QoS limit	An optional I/O priority Quality of Service (QoS) setting for maximum network transfer rate. When memory resources are low, using I/O throttling in this way will slow the memory processing and help make the system more stable by preventing crashes.

You can see all the virtual network interfaces for a VM on the VM's **Networking** tab. Here, you can add new virtual network interfaces, as well as edit, activate, deactivate and remove existing ones.

Add a Virtual Network Interface

June 3, 2019

1. Open the **Add Virtual Interface** dialog box by doing one of the following:
   - Select the VM in the **Resources** pane, click on the **Networking** tab and then click **Add Interface**.
   - On the **Network** page of the **New VM** wizard, click **Add**.
2. Select a network location from the **Network** list.
3. Specify the MAC address, either using an automatically-created MAC address based on the NIC MAC address, or by entering it manually:
   - To use a generated MAC address, click **Auto-generate a MAC address**.
   - To explicitly enter a MAC address, click **Use this MAC address** and enter an address in the form XY:XX:XX:XX:XX:XX where X is any hexadecimal digit, and Y is one of 2, 6, A or E.
4. To set an optional I/O priority Quality of Service (QoS) setting for maximum network transfer rate, select the check box and enter a value in kilobytes per second (kB/s). If the pool is configured to use a vSwitch Controller, the QoS option in the **Add Virtual Interface** dialog box will be unavailable; in this case, you should configure QoS settings using the vSwitch Controller, if required.
5. Click **Add**.

Activate/deactivate a Virtual Network Interface

June 3, 2019

You can activate or deactivate a virtual network interface on a running VM when all of the following conditions are met:

- The VM may not be suspended.
- The VM must have XenServer Tools installed.

To activate ("plug") or deactivate ("unplug") a virtual network interface, on the VM's **Networking** tab, select the interface and click the button labeled **Activate** or **Deactivate**.

Remove a Virtual Network Interface

June 3, 2019

You can remove a virtual network interface from a VM without shutting down the VM ("hot unplug") if all of the following conditions are met:

- The VM may not be suspended.
- The VM must have XenServer Tools installed.

If one or both of these conditions are not satisfied, you will have to shut the VM down before you can remove the virtual network interface.

To remove a virtual network interface:

1. Select the VM in the **Resources** pane then click on the **Networking** tab.
2. On the **Networking** tab, select the virtual network interface in the list and then click **Remove**.

Change Virtual Network Interface Properties

June 3, 2019

To change properties of a virtual network interface, open the Virtual Interface Properties dialog box by doing one of the following:

- Click on the VM's **Networking** tab, then select the virtual network interface and click **Properties**.
- On the **Network** page of the **New VM** wizard, click **Properties**.

You can change the network location and MAC address of a virtual network interface, and you may also be able to set its I/O priority.

Network	The network location of the virtual network interface.
MAC address	The MAC address of the virtual network interface can be autogenerated or you can enter it manually in the form XY:XX:XX:XX:XX:XX where X is any hexadecimal digit, and Y is one of 2, 6, A or E.
Enable QoS limit	Select this option and enter a value in kilobytes per second (kB/s) to set an optional I/O priority Quality of Service (QoS) setting for maximum network transfer rate. When memory resources are low, using I/O throttling in this way will slow the memory processing and help make the system more stable by preventing crashes. If the pool is configured to use a vSwitch Controller, the QoS option here will be unavailable; in this case, you should configure QoS settings using the vSwitch Controller, if required.

Configuring Virtual GPU

June 3, 2019

GPUs are grouped together based on the types of virtual GPUs supported on a particular GPU. XenCenter enables you to modify the virtual GPU types allowed per GPU, and thereby group the GPUs based on your requirements. For more information, see GPU.

To modify the virtual GPU types allowed on a particular GPU:

1. Select the Pool in the Resources pane and click the **GPU** tab.
2. If you have selected a pool, select GPUs that you would like to modify using the check boxes located beside the GPU. Remember, each horizontal bar on the GPU tab represents a physical GPU.
3. Click **Edit Selected GPUs**. The GPU window displays a list of virtual GPU types. It contains information specific to each virtual GPU type, such as the number of virtual GPUs allowed per GPU, maximum resolution, maximum number of displays per virtual GPU, and the Video RAM.

4. Modify the selection based on your requirements and click **OK**. If you would like to pass-through the whole GPU, select **Pass-through whole GPU**.

Managing Containers

September 30, 2020

XenServer includes the following features to enhance deployments of Docker Containers on XenServer:

- Support for CoreOS Linux VMs and configuring Cloud Config Drives
- Container Management for CoreOS, Debian 8, Ubuntu 14.04, and RHEL/CentOS/Oracle Linux 7
- Container Management for Windows Server Containers on Windows Server 2016

CoreOS is a minimalist Linux distribution which has become popular for hosting Docker™ applications. CoreOS Cloud Config Drive allows the customization of various operating system configuration options. When Container Management is enabled on a VM, XenServer becomes aware of any Docker containers running in the VM.

You should install the Container Management Supplemental Pack available from the *XenServer Download* page to enable XenServer to:

- query the VMs to interact with Cloud Config Drives
- discover the application containers running on each VM and display these within XenCenter's Infrastructure view

XenCenter enables interaction with the containers and allows users to perform container lifecycle operations, access the container console, view logs, and provides other performance monitoring and diagnostic capabilities. However, you must meet the requirements documented in the *Network Requirement* section to manage containers using XenCenter.

Network Requirements

For Container Management to work, managed VMs must be reachable from the XenServer Control Domain (dom0). Depending on the guest operating system, the networking topology and firewalls must allow outbound SSH (TCP port 22) or Docker TLS (TCP port 2376) connections from dom0 (the XenServer Management Interface) to container-managed VMs (the VM network).

> **Important:**
>
> For detailed information about container management network requirements and security, refer to the *XenServer Virtual Machine User's Guide*.

To manage Docker containers using XenCenter (for CoreOS VMs)

1. Install or upgrade your host to XenServer 7.0 or higher.

2. Install the XenCenter build shipped with XenServer 7.0 or higher.

3. Install the Container Management Supplemental pack.

4. Create a CoreOS VM and include a config drive for the VM. For more information about config drive, see Cloud-Config Parameters.

5. Enable container management for the VM. You can update this setting on the VM's **Properties** tab. For more information, see Change VM Properties.

If you wish to use Ubuntu 14.04, Debian 8, RHEL/CentOS/Oracle Linux 7, Windows Server 2016 VMs to manage Docker containers, you should first enable container management using the CLI. Once the container management is enabled on these VMs, you can use XenCenter to perform lifecycle operations such as start, stop, pause, and resume the containers.

For information about configuring Ubuntu 14.04, Debian 8, RHEL/CentOS/Oracle Linux 7, Windows Server 2016 VMs for container management, refer to the *XenServer Virtual Machine User's Guide*.

To access Docker container console and logs

For Linux VMs, XenCenter enables customers to access the container console and view logs in order to manage and monitor applications running on Docker containers. To access the container console and logs:

1. Select the container on the Resources pane.

2. On the **Container General Properties** section, click **View Console** to view the container console. To see the console logs, click **View Log**.

   This opens an SSH client on the machine running XenCenter.

3. When prompted, log into the SSH client using the VM username and password.

   > **Note:**
   >
   > Customers can automate the authentication process by configuring their public/private SSH keys. See the section *Automating the authentication process* for details.

Automating the authentication process (optional)

When accessing the container console and logs, customers are required to enter the login credentials of the VM to authenticate SSH connections. However, customers can automate the authentication process to avoid entering the credentials manually. Follow the instructions below to configure the automatic authentication process:

1. Generate a public/private key pair.

2. Add the public SSH key to the user directory on the VM running the container.

   For example, for containers running on a CoreOS VM, the public key should be added to the **Cloud-Config Parameters** section on the VM's **General** tab in XenCenter. See Cloud-Config Parameters for details. For Ubuntu 14.04, RHEL/CentOS/Oracle Linux 7, and Debian 8, the public key should be manually added to ~/.ssh/authorized_keys .

3. Add the private SSH key to the %userprofile% directory on the machine running XenCenter and rename the key as ContainerManagement.ppk .

Change VM Properties

September 30, 2020

Select a virtual machine in the **Resources** pane, and on the **General** tab, click on the **Properties** button to view or change the properties of the VM.

General

On the **General Properties** tab you can change the VM's name and description, place it in a folder, and manage its tags.

- To change the VM name, enter a new name in the **Name** box.
- To change the VM description, enter new text in the **Description** box.
- To place the VM in a folder or to move it to a different folder, click **Change** in the **Folder** box and select a folder. See Using folders for more information about using folders.
- To tag and untag the VM and to create and delete tags, see Using tags.

Custom fields

Custom fields allow you to add information to managed resources to make it easier to search and organize them. See Using custom fields to find out how to assign custom fields to your managed resources.

CPU

On the **CPU** tab, you can adjust the number of virtual CPUs allocated to the VM, set cores-per-socket for the vCPU, and specify the vCPU priority. Note that you should shut down the VM before modifying these settings.

Number of vCPUs (for Windows VMs) : To modify the number of virtual CPUs allocated to the VM, change the number in the Number of vCPUs drop-down list. To ensure you get the best performance out of your VM, the number of vCPUs should not exceed the number of physical CPUs on its host server.

Maximum number of vCPUs (for Linux VMs) : To modify the maximum number of virtual CPUs allocated to the VM, change the number in the Maximum Number of vCPUs drop-down list. To ensure you get the best performance out of your VM, the maximum number of vCPUs should not exceed the number of physical CPUs on its host server.

Topology

By default, XenCenter allocates one core per socket for each vCPU. The **Topology** drop-down list displays valid cores-per-socket combinations. Select an option from the list to modify this setting.

Depending on the number of vCPUs you select, XenCenter displays a list of options where the number of vCPUs will be divisible by the number of cores per socket. For example, if you specify 8 vCPUs for your VM, the number of cores per socket can only be 1, 2, 4, or 8. If you specify 5 vCPUs, the number of cores per socket can only be 1 or 5.

Current number of vCPUs (for Linux VMs) : This drop-down displays the current number of vCPUs allocated to the VM. You can increase the number of vCPUs allocated to the VM even when the VM is running by choosing the required number of vCPUs from the drop-down list. **Note** : You should shut down the VM to decrease the number of vCPUs allocated to the VM.

vCPU priority for this virtual machine

vCPU priority is the priority given to each of the VM's vCPUs during host CPU scheduling, relative to all of the other VMs running on the same host server. To adjust the vCPU priority for the VM, move the vCPU slider.

The XenServer templates provide typical VM configurations and set reasonable defaults for the memory, based on the type of guest operating system. You should also take into account the following considerations when deciding how much memory you give to a VM:

- The kinds of applications that will run on the VM.
- Other virtual machines that will be using the same memory resource.
- Applications that will run on the server alongside the virtual machine.

Boot Options

The available boot options on this tab may vary, depending on the guest operating system. For example, on some VMs, you can change the boot order (or boot sequence), or specify additional boot parameters.

- To change the boot order, select an item in the **Boot Order** list and click **Move Up** or **Move Down**.
- To specify additional boot parameters, enter them in the **OS Boot parameters** box. For example, on a Debian VM, you can enter single to boot the VM in single-user mode.

Start Options

On this tab you can adjust the **start order**, start delay interval and **HA restart priority** for the selected VM.

Start order

Specifies the order in which individual VMs will be started up within a vApp or during an HA recovery operation, allowing certain VMs to be started before others. VMs with a start order value of 0 (zero) will be started first, then VMs with a start order value of 1, followed by VMs with a start order value of 2, and so on.

Attempt to start next VM after

This is a delay interval that specifies how long to wait after starting the VM before attempting to start the next group of VMs in the startup sequence, that is, VMs with a lower start order. This applies to VMs within a vApp and to individual VMs during an HA recovery operation.

HA restart priority

In an HA-enabled pool, this specifies which VMs will be restarted automatically in the event of an underlying hardware failure or loss of their host server.

- VMs with an HA restart priority of **Restart** are guaranteed to be restarted if sufficient resources are available within the pool. They will be restarted before VMs with a **Restart if possible** priority.
- VMs with an HA restart priority of **Restart if possible** are not considered when calculating a failure plan, but one attempt to restart them will be made if a server that is running them fails. This restart is attempted after all higher-priority VMs are restarted, and if the attempt to start them fails, then it will not be retried.
- VMs with an HA restart priority of **Do not restart** will not be restarted automatically.

See VM startup settings for more information about these settings.

⚠ Alerts

On the **Alerts** tab, you can configure performance alerts for the VM's CPU usage, network, and disk activity.

For information about configuring alerts, see Configuring performance alerts.

🏠 Home Server

On the **Home Server** tab of the VM **Properties** dialog box you can nominate a server which will provide resources for the VM. The VM will be started up on that server if possible; if this is not possible, then an alternate server within the same pool will be selected automatically. See Creating a new VM to find out more about home servers.

Note that in pools with Workload Balancing (WLB) enabled, you cannot set a home server. Instead, XenCenter nominates the best server for the VM by analyzing XenServer resource pool metrics and recommending optimizations. You can decide if you want these recommendations geared towards resource performance or hardware density, and you can fine-tune the weighting of individual resource metrics (CPU, network, memory, and disk) so that the placement recommendations and critical thresholds align with your environment's needs.

▦ GPU

On the VM's **GPU** properties tab, you can assign a dedicated graphics processing unit (GPU) or a virtual GPU to a VM, providing direct access to the graphic hardware from the VM. This allows a VM to use the processing power of the GPU, providing better support for high-end 3D professional graphics applications such as CAD/CAM, GIS and Medical Imaging applications.

The **GPU type** list displays available GPUs and virtual GPU types. Select a virtual GPU type from the list to assign a specific virtual GPU type to the VM. Alternatively, select **Pass-through whole GPU** to allow a VM to use the full processing power of the GPU.

> **Note:**
>
> GPU Virtualization is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. For more information about licensing, see About XenServer Licensing.

⚙ Advanced Options (Optimization)

On the **Advanced Options** tab, you can adjust the amount of shadow memory assigned to a hardware-assisted VM. In some specialized application workloads, such as Citrix Virtual Apps and Desktops, extra

shadow memory is required to achieve full performance. This memory is considered to be overhead, and is separate from the normal memory calculations for accounting memory to a VM.

- To optimize performance for VMs running Citrix Virtual Apps, click **Optimize for Citrix Virtual Apps**.
- To manually adjust the VM's shadow memory allocation, click **Optimize manually** and enter a number in the **Shadow memory multiplier** box.
- To restore the default settings for shadow memory, select the **Optimize for general use** option.

Container Management

The **Container Management** tab allows you to enable the container management functionality for the VM. Select the checkbox to enable container management for the VM. After you have enabled this option, details about container management will be displayed on the VM's **General** tab. For more information, see Managing Containers.

Cloud-Config Parameters

> **Note:**
>
> You must shut down the VM before making any updates to the cloud-config parameters.

On the **Cloud-Config Parameters** tab, you can review and modify the configuration parameters you have specified for the VM. For more information about cloud-config parameters, see Cloud-Config Parameters.

Managing VMs

June 3, 2019

Topics

- Start a VM
- Suspend and resume a VM
- Shut down a VM
- Reboot a VM
- Run a remote console session
- Migrate VMs
- Delete VMs

Start a VM

September 30, 2020

For VMs in a pool, you can choose where to start your VMs, subject to available resources on the selected host server. Your choice of server depends on how the VM and the pool are configured:

- In a pool with Workload Balancing (WLB) enabled, recommendations are provided to help you choose the best possible physical server for the VM's workload.
- In a pool without Workload Balancing configured, you can choose to start the VM on its Home Server or any other server in the pool (subject to available storage on that server). See Creating a new VM to find out more about home servers.

When the VM is up and running, its status indicator changes to in the Resources pane.

To start a VM on a specific server

1. Select the VM in the **Resources** pane.
2. Right-click and click **Start on Server** and then select the server you want on the shortcut menu. Alternatively, on the **VM** menu, click **Start on Server** and then select the server you want on the submenu.

To start the VM on the optimal or home server

Select the VM in the **Resources** pane and then click **Start** on the Toolbar.

Alternatively, do one of the following:

- Right-click in the **Resources** pane and select **Start** on the shortcut menu.
- On the **VM** menu, click **Start**.

In a WLB-enabled pool, this will start the VM on the optimal server.

In a pool without Workload Balancing configured, this will start the VM on its Home Server, or on the first available server, if no Home Server has been set.

Suspend and Resume a VM

October 1, 2020

When you suspend a VM, its current state is stored in a file on the default storage repository (SR). This 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.

> **Note:**
>
> It may not be possible to resume a suspended VM that was created on a different type of server. For example a VM created on a server with an Intel VT-enabled CPU may not be resumed on a server with an AMD-V CPU.

To suspend a VM

1. If the current default SR is detached, select a new default SR.

2. Select the VM in the **Resources** pane and then click ⏸ **Suspend** on the Toolbar. Alternatively:

   - Right-click and click **Suspend** on the shortcut menu.
   - On the **VM** menu, click **Suspend**.

   When a VM has been suspended, its status indicator changes to 🔲 in the **Resources** pane.

To resume a suspended VM

For VMs in a pool, you can normally choose where to resume them. Your choice of server depends on how the VM and the pool are configured:

- In a pool with Workload Balancing (WLB) enabled, recommendations are provided to help you choose the best possible physical server for the VM's workload.
- In a pool without Workload Balancing configured, you can choose to resume the VM on its home server or any other server in the pool (subject to available storage on that server). See Home server to find out more about nominating a home server for a VM.

When a suspended VM has been successfully resumed, its status indicator changes to 🔲 in the **Resources** pane.

To resume a suspended VM on specific server

1. Select the VM in the **Resources** pane.
2. Right-click and click **Resume on Server** and then select the server you want on the shortcut menu. Alternatively, on the **VM** menu, click **Resume on Server** and then select the server you want on the submenu.

To resume the VM automatically on the optimal or home server

Select the VM in the **Resources** pane and then click ▶ **Resume** on the toolbar.

Alternatively, do one of the following:

- Right-click in the **Resources** pane and select **Resume** on the shortcut menu.

- On the **VM** menu, click **Resume**.

In a WLB-enabled pool, this will start the VM on the optimal server.

In a pool without Workload Balancing configured, this will start the VM on its Home Server, or on the first available server, if no Home Server has been set or if the nominated server is unavailable.

Shut Down a VM

September 30, 2020

You may need to shut down a running VM for a number of different reasons, for example to free up its resources, or to reconfigure its virtual network hardware or virtual disk storage. You can shut down a VM via the VM's console or using XenCenter. XenCenter provides two ways to shut down a VM:

- A soft shutdown performs a graceful shutdown of the VM, and all running processes are halted individually.
- A forced shutdown performs a hard shutdown and is the equivalent of unplugging a physical server. It may not always shut down all running processes and you risk losing data if you shut down a VM in this way. A forced shutdown should only be used when a soft shutdown is not possible.

A VM running in HVM mode (that is, VMs without XenServer Tools installed) can only be shut down using a forced shutdown; to avoid this, you should install XenServer Tools on all HVM virtual machines.

To perform a soft shutdown

Select the VM in the Resources pane and then click ⏻ **Shut Down** on the toolbar.

Alternatively:

- Right-click and click **Shut Down** on the **Resources** pane shortcut menu.
- On the **VM** menu, click **Shut Down**.

To shut down a VM from within its floating console window, click 🖥️ and then click **Shut Down**.

The VM's console displays shutdown messages as running processes are stopped. When the shutdown is complete, the VM status indicator changes to 🖥️ in the **Resources** pane.

To perform a forced shutdown

Select the VM in the **Resources** pane and then click ⏻ **Force Shutdown** on the toolbar.

Alternatively:

- Right-click and click **Force Shutdown** on the **Resources** pane shortcut menu.
- On the **VM** menu, click **Force Shutdown**.

To forcibly shut down a VM from within its floating console window, click and then click **Force Shut Down**.

When the shutdown is complete, the VM status indicator changes to in the **Resources** pane.

Reboot a VM

September 30, 2020

There are two different ways of rebooting a VM in XenCenter:

- A soft reboot performs an orderly shutdown and restart of the VM.
- A forced reboot is a hard reboot which restarts the VM without first performing any shut-down procedure. This works in the same way as pulling the plug on a physical server and then plugging it back in and turning it back on.

A forced reboot should only be done as a last resort to forcibly retrieve the system from instances such as a critical error.

A VM running in HVM mode (that is, without XenServer Tools installed) can only be rebooted using a forced reboot; to avoid this, you should install XenServer Tools on all HVM virtual machines.

To reboot a VM cleanly

In the **Resources** pane, select the VM and then click **Reboot** on the toolbar.

Alternatively:

- Right-click and click **Reboot** on the **Resources** pane shortcut menu.
- On the **VM** menu, click **Reboot**.

The VM is shut down and rebooted. When this process is complete, its status indicator in the **Resources** pane changes back to .

To do a forced reboot

In the **Resources** pane, select the VM and then click **Force Reboot** on the toolbar.

Alternatively:

- Right-click and click **Force Reboot** on the **Resources** pane shortcut menu.
- On the **VM** menu, click **Force Reboot**.

The VM is immediately shut down and rebooted. When this process is complete, its status indicator in the **Resources** pane changes back to .

Run a Remote Console Session

September 30, 2020

To open a remote console session on a VM, select the VM and then click on the Console tab.

Linux VMs

You can run a console session on Linux VMs using a text console or a graphical console. The graphical console uses VNC technology and requires the VNC server and an X display manager to be installed and configured on the virtual machine. See the *XenServer Virtual Machine User's Guide* for information about configuring VNC for Linux virtual machines.

To switch between the two types of remote console, use the **Switch to Graphical Console/Switch to Text Console** button on the **Console** tab.

SSH Console

XenCenter allows you to initiate SSH connections to Linux VMs using the Open SSH Console button located on the VM's Console tab. This will launch an SSH console for the VM in an external pop-up window. The SSH console also allows you to copy/paste content to and from the VM's console. To use the SSH console feature, you must:

- Ensure that the VM and XenCenter are accessible on the same network
- Install the Linux guest agent on the VM. For more information about installing the Linux Guest agent, see *XenServer Virtual Machine User's Guide*.
- Verify that the SSH daemon is running on the VM and accepts remote connections

> **Note:**
>
> When you close the SSH console, any operations that are still running in the console will be terminated.

Windows VMs

Console sessions on Windows VMs can use either the standard graphical console or a Remote Desktop console, both of which support full keyboard and mouse interactivity. The standard graphical console uses the in-built VNC technology that XenServer developed to provide remote access to your virtual

machine console while the Remote Desktop console uses RDP (Remote Desktop Protocol) technology. You can switch between a standard graphic console and a Remote Desktop console at any time using the **Switch to Remote Desktop/Switch to Default Desktop** button on the XenCenter **Console** tab.

To use a Remote Desktop console connection, the following requirements must be met:

- Remote Desktop must be enabled on the virtual machine - see the *XenServer Virtual Machine User's Guide* for information on how to enable Remote Desktop on a Windows virtual machine.
- XenServer Tools must be installed.
- The virtual machine must have a network interface and be able to connect to XenCenter.

There are a number of different XenCenter settings that affect your Remote Desktop console environment:

- Windows Key combinations are sent to the Remote Desktop console.
- Sounds from applications running on the Remote Desktop console are played on your local computer.
- By default, when opening a Remote Desktop console session, a connection is made to the console session on the remote server instead of creating a virtual console session.
- XenCenter automatically scans for an RDP connection and can automatically switch to the Remote Desktop console when it becomes available.

You can change these and other Remote Desktop console settings via the **Console** tab in the XenCenter **Options** dialog box; see Changing XenCenter options.

> **Note:**
>
> You can enhance VNC performance by using XenCenter on the local machine rather than using RDP to connect to XenCenter.

Migrate Virtual Machines

September 30, 2020

This topic contains information about migrating and moving virtual machines within and across pools and standalone servers. Migrating a VM refers to the operation of moving a running or a suspended VM to a different server or a pool. Whereas, moving a VM refers to the act of moving a VM that is shut down.

About VM Migration with XenMotion and Storage XenMotion

XenMotion

XenMotion is available in all versions of XenServer and allows you to move a running or a suspended VM between XenServer hosts, when the VM's disks are located on storage shared by both hosts. This allows for pool maintenance features such as Workload Balancing (WLB), High Availability (HA), and Rolling Pool Upgrade (RPU) to automatically move VMs. These features allow for workload leveling, infrastructure resilience, and the upgrade of server software, without any VM downtime. Storage can only be shared between hosts in the same pool, as a result, VMs can only be moved within the same pool.

Storage XenMotion

Storage XenMotion additionally allows VMs to be moved from one host to another, where the VMs are not located on storage shared between the two hosts. As a result, VMs stored on local storage can be migrated without downtime and VMs can be moved from one pool to another with virtually no service interruption. The choice of destination server depends on how the VM and the pool are configured. The choice of destination server depends on how the VM and the pool are configured: in a pool with Workload Balancing (WLB) enabled, for example, recommendations are provided to help select the best possible physical server for the VM's workload. For more information, see Choosing an Optimal Server for VM Initial Placement, Migrate, and Resume.

Storage XenMotion enables system administrators to:

- rebalance VMs between XenServer pools (for example from a development environment to a production environment)
- upgrade and update standalone XenServer hosts without any VM downtime
- upgrade XenServer host hardware

Note:

Moving a VM from one host to another preserves the VM state. The state information includes information that defines and identifies the VM as well as the historical performance metrics, such as CPU and network usage.

Storage XenMotion also allows you to move virtual disks from one Storage Repository (SR) to a different SR within the same pool. For more information, see Move Virtual Disks.

Compatibility requirements

When migrating a VM with XenMotion or Storage XenMotion, the new VM and server must meet the following compatibility requirements:

- XenServer Tools must be installed on each VM that you wish to migrate.

- The destination server must have the same or a more recent version of XenServer installed as the source.
- For Storage XenMotion, if the CPUs on the source and destination server are different, the destination server must provide at least the entire feature set as the source server's CPU. Consequently, it is unlikely to be possible to move a VM between, for example, AMD and Intel processors.
- For Storage XenMotion, VMs with more than six attached VDIs cannot be migrated.
- The target server must have sufficient spare memory capacity or be able to free sufficient capacity using Dynamic Memory Control. If there is not enough memory, the migration will fail to complete.
- For Storage XenMotion, the target storage must have enough free disk space (for the VM and its snapshot) available for the incoming VMs. If there is not enough space, the migration will fail to complete.

XenMotion and Storage XenMotion limitations

XenMotion and Storage XenMotion are subject to the following limitations:

- Storage XenMotion must not be used in Citrix Virtual Desktops deployments.
- VMs using PCI pass-thru cannot be migrated.
- VM performance will be reduced during migration.
- For Storage XenMotion, pools protected by High Availability (HA) should have HA disabled before attempting VM migration.
- Time to completion of VM migration will depend on the memory footprint of the VM, and its activity, in addition, VMs being migrated with Storage XenMotion will be affected by the size of the VDI and its storage activity.

For step-by-step instructions on using XenMotion or Storage XenMotion to migrate your VMs, see the section *To Migrate or Move a VM*.

Move VMs

XenCenter allows you move VMs that are shut down to a new storage in the same pool using the Move VM wizard. For step-by-step instructions, see the section To migrate or move a VM.

To Migrate or Move a VM

1. In the Resources pane, select the VM and do one of the following depending on the status of your VM.

   To migrate a running or suspended VM using XenMotion or Storage XenMotion: On the **VM** menu, click **Migrate to Server** and then **Migrate VM wizard**. This opens the Migrate VM wizard.

To move a stopped VM: On the VM menu, select **Move VM**. This opens the **Move VM wizard**.

2. Select a standalone server or a pool from the **Destination** drop-down list.

3. Select a server from the **Home Server** list to assign a home server for the VM and click **Next**.

4. On the **Storage** page, specify a storage repository where you would like to place the virtual disks of the migrated VM and click **Next**.

   - The **Place all migrated virtual disks on the same SR** option is selected by default and displays the default shared SR on the destination pool.
   - Click **Place migrated virtual disks onto specified SRs** to specify an SR from the **Storage Repository** drop-down list. This option allows you to select different SR for each virtual disk on the migrated VM.

5. On the **Networking** page, map the virtual network interfaces in the selected VM to networks in the destination pool or standalone server. Specify your options using the **Target Network** drop-down list and click **Next**.

6. Select a storage network on the destination pool that will be used for the migration of the VM's virtual disks and then click **Next**.

   > **Note:**
   >
   > Due to performance reasons, it is recommended that you do not use management network for live migration.

7. Review the configuration settings and click **Finish** to start migrating or moving the VM.

Delete a VM

September 30, 2020

Deleting a virtual machine removes its configuration and its filesystem from the server. When you delete a VM, you can choose to also delete or preserve any virtual disks that are attached to the VM, as well as any snapshots of the VM.

To delete a VM:

1. Shut down the VM.

2. Select the stopped VM in the **Resources** panel, right-click and click **Delete** on the shortcut menu. Alternatively, on the **VM** menu, click **Delete**.

3. To delete an attached virtual disk, select its check box.

> **Important:**
>
> Any data stored in the VM's virtual disk drives will be lost.

4. To delete a snapshot of the VM, select its check box.

5. Click **Delete**.

   When the delete operation is completed, the VM is removed from the **Resources** pane.

> **Note:**
>
> VM snapshots whose parent VM has been deleted ("orphan snapshots") can still be accessed from the **Resources** pane and can be exported, deleted, or used to create new VMs and templates. To view snapshots in the **Resources** pane, click **Objects** in the Navigation pane and then expand the **Snapshots** group in the Resources pane.

Importing and Exporting VMs

June 3, 2019

Topics

- About VM Import and Export
- Open Virtualization Format (OVF and OVA)
- Disk Image Formats (VHD and VMDK)
- Import VMs From OVF/OVA
- Import Disk Images
- Import VMs From XVA
- Export VMs as OVF/OVA
- Export VMs as XVA

About VM Import and Export

June 3, 2019

VMs can be imported from OVF/OVA packages, from disk images, and from XenServer XVA files. VMs can be exported as OVF/OVA packages and as XenServer XVA files. You import and export VMs in XenCenter using the Import and Export wizards. XenCenter uses the Transfer VM to transfer the content of a disk image.

When importing VMs created on hypervisors other than XenCenter, for example Hyper-V or VMware, it is necessary to use the Operating System Fixup tool to ensure that imported VMs can boot on XenServer.

Supported Import and Export Formats

Format	Description
Open Virtualization Format (OVF and OVA)	OVF is an open standard for packaging and distributing a virtual appliance consisting of one or more virtual machines. For more information about XenCenter support for OVF and OVA file formats, see Open Virtualization Format (OVF and OVA).
Disk image formats (VHD and VMDK)	Virtual Hard Disk (VHD) and Virtual Machine Disk (VMDK) format disk image files can be imported using the Import wizard. You might want to import a disk image when only a virtual disk image is available, but there is no OVF metadata associated with it. For more information about supported disk image formats, see Disk Image Formats (VHD and VMDK)
XenServer XVA format	XVA is a format specific to Xen-based hypervisors for packaging a single VM as a single file archive of a descriptor and disk images. Its file extension is .xva.

XenServer XVA Version 1 (ova.xml)	XVA version 1 is a format for packaging a single VM as a set of files including a descriptor and disk images. The descriptor, named ova.xml , specifies the virtual hardware of a single VM. The disk image format is a directory of files. The directory name corresponds to a reference name in the descriptor. There is one file for each one GB chunk of the disk image. The base name of each file includes the chunk number in decimal. It contains one block of the disk image in raw binary format compressed with gzip.

Which Format to Use?

Use OVF/OVA to:

- Share XenServer vApps and VMs with other hypervisors that support OVF.
- Save more than one VM.
- Secure a vApp or VM from corruption and tampering.
- Include a license agreement.
- Simplify vApp distribution by storing an OVF package in an OVA.

Use XVA to:

- Share VMs with versions of XenServer earlier than 6.0.
- Import and export VMs from a script with a command line interface (CLI).

Operating System Fixup

XenCenter includes an advanced hypervisor interoperability feature – Operating System Fixup – which aims to ensure a basic level of interoperability for VMs that are imported to XenServer. You will need to use Operating System Fixup when importing VMs created on other hypervisors from OVF/OVA packages and disk images.

Operating System Fixup configures a guest operating system to boot in XenServer by enabling boot devices critical for booting in XenServer and disabling any services, also called tools, for hypervisors other than XenServer. Guest operating systems include all versions of Windows that XenServer supports and some Linux distributions. Note that Operating System Fixup does not convert the guest operating system from one hypervisor to another.

Operating System Fixup is supplied as an automatically booting ISO image that is attached to the imported VM's DVD drive and which performs the necessary configuration changes when the VM is first started, and then shuts down the VM. The next time the new VM is started, the boot device is reset and the VM starts normally.

To use Operating System Fixup on imported disk images and OVF/OVA packages, you enable the feature on the **OS Fixup Settings** page of the XenCenter Import wizard and specify a location where the Fixup ISO should be copied so that XenServer can use it.

Operating System Fixup requirements

Operating System Fixup requires an ISO SR with 40 MB of free space and 256 MB of virtual memory.

Transfer VM

The Transfer VM is a built-in VM that only runs during the import or export of a virtual disk image to transfer its contents between the disk image file location and A XenServer storage repository (SR). One Transfer VM runs for each import or export of a disk image. When importing or exporting VMs or vApps with more than one disk image, only one disk image transfers at a time.

You will need to configure networking settings for the Transfer VM in the **Import** and **Export** wizards.

Transfer VM requirements

The XenServer requirements to run one Transfer VM are:

Virtual CPU	1
Virtual memory	256 MB
Storage	8 MB
Networking	A network that is reachable from the XenServer host - usually the management network. Static or dynamic IP address. (Citrix recommends the use of a dynamic IP address.)

The default transfer protocol is iSCSI, and this requires an iSCSI Initiator on the XenCenter host. RawVDI can also be used as an alternate transfer protocol; see the XenServer Virtual Machine User's Guide for details.

Importing VMs: overview

When you import a VM, you are effectively creating a new VM, which involves many of the same steps as creating and provisioning a new VM using the **New VM** wizard, such as nominating a home server, and configuring storage and networking for the new VM; see Creating a New VM for detailed information about each of these steps.

The **Import** wizard takes you through the following steps to import a VM:

1. **Select the import file.**

   The first step is to locate and select the file containing the VM or VMs you want to import.

   For files that are not currently located on your local XenCenter host, you can enter a URL location (http, https, file, ftp) in the **Filename** box. On clicking **Next** , a **Download File** dialog box opens and you can specify a folder on your XenCenter host where the file will be copied. The **Import** wizard will continue to the next page when the file has been downloaded.

2. **(VHD and VMDK import only) Specify the new VM's name and allocate vCPU and memory resources.**

   When importing from VHD or VMDK file, you will need to specify a name for the new VM and allocate it some virtual CPUs (vCPUs) and memory. All of these values can be adjusted later, after the new VM has been created. See VM CPU and Memory Allocation for more information on this step. VM names are not checked for uniqueness within XenCenter, so it will make it easier for you to manage different VMs if you give them meaningful, memorable names; see VM Name and Description for more information.

3. **(OVF/OVA only) Review/accept EULA(s).**

   If the package you are importing includes any EULAs, accept them and then click **Next** to continue. If no EULAs are included in the package, the wizard will skip this step and move straight on to the next page.

4. **Choose the location/home server.**

   Select the destination pool or standalone server where you want to place the imported VM(s). To nominate a Home Server for the incoming VM(s), select a server in the list.

5. **Configure storage.**

   Next, choose the storage repositories (SRs) where the virtual disks in the imported VM(s) will be placed:

   For VMs in XVA format, you select an SR where all of the imported VM's virtual disks will be placed.

   For VMs in OVF/OVA packages or in disk image files, you can place all of the imported virtual disks onto the same SR, or you can place individual virtual disks onto specific SRs.

6. **Configure networking.**

   Next, map the virtual network interfaces in the imported VM(s) to target networks in the destination pool/standalone server.

7. **(OVF/OVA only) Security validation.**

   If the selected OVF/OVA package is configured with security features such as certificates or a manifest, you will need to specify the necessary information.

8. **(OVF/OVA and disk image only) Configure OS Fixup and Transfer VM settings.**

   If the VMs you are importing were built on a hypervisor other than XenServer, you will need to configure the Operating System Fixup feature to enable the imported VM to boot correctly on XenServer. You will also need to configure networking for the Transfer VM, a temporary VM used when importing/exporting VMs from disk images and OVF/OVA packages.

9. **Complete new VM creation.**

   On the final page of the **Import** wizard, you can review all the configuration options you have chosen. When importing from XVA, you can select the **Start VM automatically** check box to have the new VM start automatically as soon as it is created.

   Click **Finish** to finish importing the selected VMs and close the wizard.

Exporting VMs: overview

Select the VM or VMs you want to export and then open **Export** wizard: on the **VM** menu, click **Export**.

1. **Specify export file details.**

   On the first page of the wizard, enter the name of the export file, specify the folder where you want the file to be saved, and choose the export file format from the **Format** list:

   Choose **XVA File (*.xva)** to export the selected VM to an XVA file. Only single VMs can be exported in this format.

   Choose **OVF/OVA Package (*.ovf, *.ova)** to export the selected VM(s) as an OVF or OVA package.

2. **Confirm VMs selected for export.**

   On the next page of the wizard, you can modify the VM selection set, for example, to add more VMs to the set of VMs to be exported or to deselect others. Note that when exporting to XVA, only one VM may be selected.

3. **(OVF/OVA only) Configure EULA, Advanced Options, and Transfer VM settings**

   When exporting VMs as an OVF of OVA package, a number of additional settings can be configured. See Export VMs as OVF/OVA for details.

4. **Complete VM export.**

On the final page of the wizard, review the settings you have selected on the previous wizard pages. To have the wizard verify the export file, select the **Verify export on completion** check box.

Click **Finish** to begin exporting the selected VM(s) and close the wizard.

Open Virtualization Format (OVF and OVA)

June 3, 2019

OVF is an open standard, specified by the Distributed Management Task Force (DMTF), for packaging and distributing a virtual appliance consisting of one or more virtual machines (VMs).

An **OVF Package** is composed of metadata and file elements that describe virtual machines, plus additional information that is important to the deployment and operation of the applications in the OVF package. Its file extension is .ovf .

An **Open Virtual Appliance (OVA)** is an OVF Package in a single file archive with the .ova extension.

Note that in a XenServer environment where Role-Based Access Control (RBAC) is implemented, only users with the RBAC role of Pool Admin can import and export OVF and OVA packages. See RBAC overview for more information about RBAC roles.

What's in an OVF package?

An **OVF Package** always includes a descriptor file (*.ovf) and may also include a number of other files.

File type	Description
Descriptor	The descriptor specifies the virtual hardware requirements of the service and can also include other information such as descriptions of virtual disks, the service itself, and guest operating systems, a license agreement (EULA), instructions to start and stop VMs in the appliance, and instructions to install the service. The descriptor file extension is .ovf .

Manifest	The manifest is an SHA-1 digest of every file in the package, allowing the package contents to be verified by detecting any corruption. The manifest file extension is .mf .
Signature	The signature is the digest of the manifest signed with the public key from the X.509 certificate included in the package, and allows the package author to be verified. The signature file extension is .cert .
Virtual disks	OVF does not specify a disk image format. An OVF package includes files comprising virtual disks in the format defined by the virtualization product that exported the virtual disks. XenServer produces OVF packages with disk images in Dynamic VHD format; VMware products and Virtual Box produce OVF packages with virtual disks in Stream-Optimized VMDK format.

An **OVA package** is a single archive file, in the Tape Archive (tar) format, containing the files that comprise an OVF Package.

Which format should I use?

OVF packages contain a series of uncompressed files that make it handier for users who may want to access individual disk images in the file, while OVA packages are just one large file. While you can compress this file, it doesn't have the flexibility of a series of files like OVF.

OVA is better for specific applications where it is beneficial to have just one file, making the package easier to handle, such as when creating packages for Web downloads. Exporting and importing OVA packages takes longer than OVF.

Using Operating System Fixup

Using OVF as a method of packaging does not guarantee cross-hypervisor compatibility of the virtual machines contained in the package. An OVF package created on one hypervisor might not automatically work on a different hypervisor. This happens for a variety of reasons including; different inter-

pretations of the OVF specification, guest operating system devices, drivers, and implementations inherent to a hypervisor.

XenCenter includes an advanced hypervisor interoperability feature – **Operating System Fixup** which aims to ensure a basic level of interoperability for OVF packages that are imported to XenServer. It is necessary to run Operating System Fixup on imported VMs that were created on other hypervisors to ensure that they boot correctly on XenServer.

For more information about the Operating System Fixup feature, see About VM Import and Export.

More information about OVF

Refer to the following document on the DMTF website for more information about OVF:

Open Virtualization Format Specification

Disk Image Formats (VHD and VMDK)

June 3, 2019

Using the **Import** wizard, you can import a disk image into a resource pool or into a specific host as a VM.

You might want to import a disk image when only a virtual disk image is available, but there is no OVF metadata associated with it. Situations when this might occur include:

- The OVF metadata is not readable. However, it is still possible to import the disk image.
- You have a virtual disk that is not defined in an OVF package.
- You are moving from a platform that does not let you create an OVF appliance (for example, older platforms or images).
- You want to import an older VMware appliance that does not have any OVF information.
- You want to import a standalone VM that does not have any OVF information.

When available, Citrix recommends importing appliance packages that contain OVF metadata and not just importing an individual disk image. The OVF data provides information that the **Import** wizard needs to recreate a VM from its disk image, including the number of disk images associated with the VM, the processor, storage, and memory requirements and so on. Without this information, it can be much more complex and error-prone trying to recreate the VM.

Note that in A XenServer environment where Role-Based Access Control (RBAC) is implemented, only users with the RBAC role of Pool Admin can import disk images. See RBAC overview for more information about RBAC roles.

Supported disk image formats

The following disk image formats can be imported using the XenCenter **Import** wizard:

Format	Description
Virtual Hard Disk (VHD)	VHD is a group of virtual disk image formats specified by Microsoft as part of their Open Specification Promise. Their file extension is .vhd . XenCenter imports and exports the Dynamic VHD format – a thinly provisioned virtual disk image that allocates space only when used.
Virtual Machine Disk (VMDK)	VMDK is a group of virtual disk image formats specified by VMware. Their file extension is .vmdk . XenCenter imports stream-optimized and monolithic flat VMDK formats. Stream-optimized VMDK is the format used by OVF packages produced for VMware and Virtual Box hypervisors. Monolithic flat VMDK is a common format of a virtual disk available for download by VMware management clients.

Using Operating System Fixup

XenCenter includes an advanced hypervisor interoperability feature Operating System Fixup which aims to ensure a basic level of interoperability for VMs created on hypervisors other than XenServer. It is necessary to run Operating System Fixup when importing VMs from disk images that were created on other hypervisors to ensure that they boot correctly on XenServer.

To find out more, see About VM Import and Export.

Import VMs From OVF/OVA

September 30, 2020

You can import virtual machines (VMs) that have been saved as OVF/OVA files using the **Import** wizard. The wizard will take you through many of the usual steps needed to create a new VM in XenCenter: nominating a home server, and configuring storage and networking for the new VMs, plus some additional steps required as part of the OVF import process, including:

- Specifying security settings if the OVF package includes a certificate or a manifest.
- Configuring networking for the About VM Import and Export.
- Specifying **Operating system fixup** settings if importing VMs that were built on a hypervisor other than XenServer.

See Creating a New VM for an overview of the steps involved in creating a new VM.

Imported OVF packages will appear as vApps when imported using XenCenter. When the import is complete, the new VMs will appear in the XenCenter Resources pane and the new vApp will appear in the Managing vApps dialog box.

Prerequisites

- You need an RBAC role of Pool Admin to import OVF/OVA packages. The Import wizard will perform checks to ensure that you have a Pool Admin role in the destination pool before allowing you to continue. See RBAC overview for more information about RBAC roles.
- When importing an OVF Package that was compressed or contains compressed files, additional free disk space is necessary on the XenServer host to decompress the files.

To import an OVF package

1. Open the **Import** wizard: on the **File** menu, select **Import**.

2. On the first page of the wizard, locate the package you want to import (with a .ovf , .ova or .ova.gz file extension), then click **Next** to continue.

   - If you select a compressed OVA file (*.ova.gz), on clicking **Next** , the file will be decompressed to an OVA file and the old *.ova.gz file will be deleted.

   - If you enter a URL location (http, https, file, ftp) in the **Filename** box, on clicking **Next** , a **Download Package** dialog box opens and you must specify a folder on your XenCenter host where the package will be copied.

3. **Review/accept EULA(s).**

   If the package you are importing includes any EULAs, accept them and then click **Next** to continue. If no EULAs are included in the package, the wizard will skip this step and move straight on to the next page.

4. **Specify the VM location and home server.** On the Location page, choose the pool or standalone server where you want to place the VMs you are importing from the Import VM(s) to list, and (optionally) assign them a home server:

   - To nominate a home server for a VM, select the server from the list in the **Home Server** column. XenServer will always attempt to start up a VM on its home server if it can. See Home Server for more information on assigning a home server to new VMs.

- If you do not want to nominate a home server, select **Don't assign a home server** from the list in the **Home Server** column.

  Click **Next** to continue.

5. **Configure storage for the imported VMs.** On the **Storage** page, select one or more storage repositories (SRs) where the disk images for the imported will be placed, then click **Next** to continue.

   - To place all the imported disk images on the same SR, click **Place all imported VMs on this target SR** and select an SR from the list.

   - To place the disk images of incoming VMs onto different SRs, click **Place imported VMs on the specified SR targets** then, for each virtual disk, select the target SR from the list in the **SR** column.

6. **Configure networking for the imported VMs.** On the **Networking** page, map the virtual network interfaces in the VMs you are importing to target networks in the destination pool. The Network and MAC address shown in the list of incoming VMs on this page are stored as part of the definition of the original (exported) VM in the export file. To map an incoming virtual network interface to a target network, select a network from the list in the **Target network** column.

   Click **Next** to continue.

7. **Specify security settings.** If the selected OVF/OVA package is configured with security features such as certificates or a manifest, specify the necessary information on the Security page and then click Next to continue. Different options appear on this page depending on which security features have been configured on the OVF package:

   - If the package is signed, a **Verify digital signature** check box appears here; click this check box if you want to verify the signature. Click **View Certificate** to display the certificate used to sign the package. If the certificate appears as untrusted, it is likely that the either the Root Certificate or the Issuing Certificate Authority is not trusted on the local computer.

   - If the package includes a manifest, a **Verify manifest content** check box appears here. Select this check box to have the wizard verify the list of files in the package.

   Note that when packages are digitally signed, the associated manifest is verified automatically and so the **Verify manifest content** check box does not appear on the **Security** page.

   > **Important:**
   >
   > VMware Workstation 7.1 produces an OVF appliance with a manifest that has invalid SHA-1 hashes. Choosing to verify the manifest when importing an appliance from this source will cause the import fail.

8. **Enable Operating System Fixup.** If the VMs in the package you are importing were built on a hypervisor other than XenServer, select the **Use Operating System Fixup** check box and select

an ISO SR where the Fixup ISO can be copied so that XenServer can use it. For more information about this feature, see the Operating System Fixup information in the *VM User's Guide*.

If the ISO library you want is not listed, click **New ISO Library** to create a new ISO SR; see ISO Storage for more information about using this option.

Click **Next** to continue.

9. **Configure Transfer VM networking.** On the **Transfer VM Settings** page, configure the networking settings for the About VM Import and Export , then click **Next** to continue.

   - Select a network from the list of network interfaces available in the destination pool or standalone server.

   - To use automated Dynamic Host Configuration Protocol (DHCP) to automatically assign networking settings including the IP address, subnet mask and gateway, select **Automatically obtain network settings using DHCP.**

   - To configure networking settings manually, click **Use these network settings** and enter the IP address, subnet mask and gateway.

10. On the **Finish** page, review all the import settings and then click **Finish** to begin the import process and close the wizard.

    The import progress is displayed on the status bar at the bottom of the XenCenter window and also on the **Events** view under **Notifications**.

    The import process may take some time, depending on the size of the imported virtual disks, the available network bandwidth, and the disk interface speed of the XenCenter host. When the import is finished, the newly imported VMs will appear in the **Resources** pane and the new vApp will appear in the Managing vApps dialog box.

    > **Note:**
    >
    > After using XenCenter to import an OVF package that contains Windows operating systems, you must set the **platform** parameter. This will vary according to the version of Windows contained in the OVF package:
    >
    > - For Windows Vista, Server 2008, and later, set the platform parameter to device_id=0002. For example:

    ```
    1  xe vm-param-set uuid=&lt;VM UUID&gt; platform:device\_id
       =0002
    2  <!--NeedCopy-->
    ```

    > - For all other versions of Windows, set the platform parameter to viridian=true. For

---

> example:

```
1  xe vm-param-set uuid=&lt;VM UUID&gt; platform:viridian=true
2  <!--NeedCopy-->
```

Errors when trying to start an imported VM

If you cannot boot the VMs imported from an OVF package, try importing the package again without using the Operating System Fixup feature: in the **OS Fixup Settings** page of the **Import** wizard, clear the **Use Operating System Fixup** check box (see About VM Import and Export for more information about this feature.)

Import Disk Images

September 30, 2020

Using the XenCenter **Import** wizard, you can import a disk image into a resource pool or a specific host, creating a new VM. The wizard will take you through many of the usual steps needed to create a new VM in XenCenter: nominating a home server, and configuring storage and networking for the new VM, plus some additional steps required as part of the import process, including:

- Configuring networking for the About VM Import and Export.
- Specifying About VM Import and Export settings if importing disk images that were built on a hypervisor other than XenServer.

Requirements for importing disk images

You need an RBAC role of Pool Admin to import disk images. The **Import** wizard will perform checks to ensure that you have a Pool Admin role in the destination pool before allowing you to continue. See RBAC overview for more information about RBAC roles.

Procedure

- Open the **Import** wizard: on the **File** menu, select **Import**.

- On the first page of the wizard, locate the disk image file you want to import, then click **Next** to continue.

If you enter a URL location (http, https, file, ftp) in the **Filename** box, on clicking **Next** , a **Download Package** dialog box opens and you must specify a folder on your XenCenter host where the disk image will be copied.

- **Specify the VM name and allocate CPU and memory resources.**

  On the **VM Definition** page, enter the name of the new VM to be created from the imported disk image, and allocate CPU and initial memory resources. See VM CPU and Memory Allocation for guidelines on these fields.

  Click **Next** to continue.

- **Specify where to place the new VM and choose a home server.**

  On the **Location** page, choose where you want to place the new VM, and (optionally) assign it a home server, then click **Next** to continue.

    - Click on a pool or standalone server in the **Import VM(s) to** list.
    - To nominate the home server for the VM, select a server from the list in the **Home Server** column. XenServer will always attempt to start up a VM on its home server if it can. See Home Server for more information on assigning a home server to new VMs.
    - If you do not want to nominate a home server, select **Don't assign a home server** from the list in the **Home Server** column.

- **Configure storage for the new VM.**

  On the **Storage** page, select a storage repository (SR) where the imported virtual disk will be placed, then click **Next** to continue.

- **Configure networking for the new VM.**

  On the **Networking** page, select a target network in the destination pool/standalone server for the new VM's virtual network interface.

  Click **Next** to continue.

- **Enable Operating System Fixup.**

  If the disk image that you are importing was built on a hypervisor other than XenServer, select the **Use Operating System Fixup** check box and select an ISO SR where the Fixup ISO can be copied so that XenServer can use it. See About VM Import and Export for more information about this feature.

- **Configure Transfer VM networking.**

  On the **Transfer VM Settings** page, configure the networking settings for the About VM Import and Export , then click **Next** to continue.

    - Select a network from the list of network interfaces available in the destination pool or standalone server.

- To use automated Dynamic Host Configuration Protocol (DHCP) to automatically assign networking settings including the IP address, subnet mask and gateway, select **Automatically obtain network settings using DHCP.**
- To configure networking settings manually, click **Use these network settings** and enter the IP address, subnet mask and gateway.

- On the **Finish** page, review all the import settings and then click **Finish** to begin the import process and close the wizard.

The import progress is displayed on the status bar at the bottom of the XenCenter window and also on the **Events** view under **Notifications**.

The import process may take some time, depending on the size of the imported virtual disks, the available network bandwidth, and the disk interface speed of the XenCenter host. When the import is finished, the newly-imported VMs will appear in the **Resources** pane.

> **Note:**
>
> After using XenCenter to import a disk image that contains Windows operating systems, you must set the platform parameter. This will vary according to the version of Windows contained in the disk image:
>
> - For Windows Vista, Server 2008, and later, set the platform parameter to device_id=0002. For example:
>
> ```
> 1  xe vm-param-set uuid=&lt;VM UUID&gt; platform:device\_id
>       =0002
> 2  <!--NeedCopy-->
> ```
>
> - For all other versions of Windows, set the platform parameter to viridian=true. For example:
>
> ```
> 1  xe vm-param-set uuid=&lt;VM UUID&gt; platform:viridian=true
> 2  <!--NeedCopy-->
> ```

Import VMs From XVA

September 30, 2020

You can import VMs, templates and snapshots that have previously been exported and stored locally

in XVA format (with a .xva file extension) or XVA version 1 format (ova.xml and associated files) using the XenCenter **Import** wizard.

Importing a VM from an XVA or ova.xml file involves the same steps as creating and provisioning a new VM using the **New VM** wizard, such as, nominating a home server, and configuring storage and networking for the new VM; see Creating a New VM for detailed information about each of these steps.

> **Important:**
>
> It may not always be possible to run an imported VM that was exported from another server with a different CPU type. For example, a Windows VM created on a server with an Intel VT Enabled CPU, then exported, may not run when imported to a server with an AMD-V CPU.

Procedure

1. Open the **Import** wizard by doing one of the following:

   - In the **Resources** pane, right-click, and then select **Import** on the shortcut menu.
   - On the **File** menu, select **Import**.

2. On the first page of the wizard, locate the XVA (or ova.xml) file you want to import, then click **Next**. If you enter a URL location (http, https, file, ftp) in the **Filename** box, on clicking **Next** , a **Download Package** dialog box opens and you must specify a folder on your XenCenter host where the file(s) will be copied.

3. On the **Home Server** page, specify where to put the new VM:

   - To place the imported VM in a pool without assigning it a home server, select the destination pool in the list, then click **Next**.
   - To place the imported VM in a pool and assign it to a specific home server (or to place it on a standalone server), select a server and then click **Next** to continue.

4. On the **Storage** page, select a storage repository (SR) where the imported virtual disks will be placed, then click **Next** to continue.

5. On the **Networking** page, map the virtual network interfaces in the VM you are importing to target networks in the destination pool. The Network and MAC address shown in the list on this page are stored as part of the definition of the original (exported) VM in the export file. To map an incoming virtual network interface to a target network, select a network from the list in the **Target network** column.

   Click **Next** to continue.

6. On the last page of the **Import** wizard, review the configuration options you have selected. To have the imported VM start up as soon as the import process has finished and the new VM is provisioned, select the **Start VM after import** check box.

---

7. Click **Finish** to begin importing the selected file and close the wizard.

   The import progress is displayed on the status bar at the bottom of the XenCenter window and also on the **Events** view under **Notifications**.

   The import process may take some time, depending on the size of the imported VM's virtual disks, the available network bandwidth, and the disk interface speed of the XenCenter host. When the newly imported VM is available, it appears in the **Resources** pane.

Export VMs as OVF/OVA

June 3, 2019

You can export one or more VMs as an OVF or OVA package using the XenCenter **Export** wizard. To open the wizard, select the VM you want to export and on the **VM** menu, click **Export**. Note that the VMs must be shut down or suspended before they can be exported.

Note that you need an RBAC role of Pool Admin to export to OVF/OVA. The **Export** wizard will perform checks when it starts up to ensure that you have a Pool Admin role before allowing you to continue. See RBAC overview for more information about RBAC roles.

Procedure

1. Open the **Export** wizard: select the pool or server containing the VM(s) you want to export, then on the **VM** menu, click **Export**.

2. On the first page of the wizard, enter the name of the export file, specify the folder where you want the file(s) to be saved, and select **OVF/OVA Package (*.ovf, *.ova)** from the **Format** list, and click **Next**.

3. Select the VMs you want to export and then click **Next**.

4. On the **EULAs** page, you can add previously prepared End User Licensing Agreement (EULA) documents (.rtf, .txt) in the package. To view the contents of a EULA in a text editor, select it in the list and click **View**. If you do not wish to include a EULA in the package, just click **Next** to continue.

5. On the Advanced options page, specify any manifest, signature, and output file options, or just click Next to continue:

   a) To create a manifest for the package, select the **Create a manifest** check box. The manifest provides an inventory or list of the other files in a package and is used to ensure the files originally included when the package was created are the same files present when the

package arrives. When the files are imported, a checksum is used to verify that the files have not changed since the package was created.

b) To add a digital signature to the package, select the **Sign the OVF package** check box, browse to locate a certificate, and enter the private key associated with the certificate in the **Private key password** box. When a signed package is imported, the user can verify the package creator's identity by using the certificate's public key to validate the digital signature. Use a X.509 certificate which you have already created from a Trusted Authority and exported as either a .pem or .pfx file that contains the signature of the manifest file and the certificate used to create that signature.

c) To output the selected VMs as a single (tar) file in OVA format, select the **Create OVA Package** check box. See Open Virtualization Format (OVF and OVA) for more on the different file formats.

d) To compress the virtual hard disk images (.VHD files) included in the package, select the **Compress OVF files** check box. By default when you create an appliance package, the virtual hard disk images that are exported consume the same amount of space that was allocated to the VM. For example, a VM that is allocated 26 GB of space will have a hard disk image that consumes 26 GB of space, regardless of whether or not the VM actually requires that much space. Note that compressing the VHD files makes the export process take longer to complete, and importing a package containing compressed VHD files will also take longer as the Import wizard must extract all of the VHD images as it imports them.

If both the **Create OVA Package** and **Compress OVF files** options are checked, the result is a compressed OVA file *.ova.gz .

6. On the **Transfer VM Settings** page, configure networking options for the temporary VM used to perform the export process (the "Transfer VM"). Click **Next** to continue.

7. On the final page of the wizard, review the settings you have selected on the previous pages. To have the wizard verify the exported package, select the **Verify export on completion** check box. Click **Finish** to begin exporting the selected VMs and close the wizard.

The export process may take some time, depending on the size of the virtual disks, the available network bandwidth, and the disk interface speed of the XenCenter host. Progress is displayed on the status bar at the bottom of the XenCenter window and on the **Events** view under **Notifications**.

To cancel an export in progress, click **Notifications** and then **Events** , find the export in the list of events, and click **Cancel**.

Export VMs as XVA

September 30, 2020

You can export a single VM as an XVA file using the **Export** wizard. VMs must be shut down or suspended before they can be exported.

> **Note:**
>
> It may not always be possible to run an imported VM that was exported from another server with a different CPU type. For example, a Windows VM created on a server with an Intel VT enabled CPU, then exported, may not run when imported to a server with an AMD-V CPU.

Procedure

1. Select the VM you want to export and on the **VM** menu, click **Export**.

2. On the first page of the wizard, enter the name of the export file, specify the folder where you want the file to be saved, select **XVA File (*.xva)** from the **Format** list, and then click **Next**.

3. On the **Virtual Machines** page, the VM to be exported is selected in the list. When exporting as XVA, only one VM may be selected in this list. Click **Next** to continue.

4. On the final page of the wizard, review the settings you have selected on the previous pages. To have the wizard verify the exported XVA file, select the **Verify export on completion** check box. Click **Finish** to begin exporting the selected VM and close the wizard.

   The export process may take some time, depending on the size of the VM's virtual disks, the available network bandwidth, and the disk interface speed of the XenCenter host. Progress is displayed in the status bar at the bottom of the XenCenter window and on the **Events** view under **Notifications**.

   To cancel an export in progress, click **Notifications** and then **Events** , find the export in the list of events, and click **Cancel**.

VM Snapshots

June 3, 2019

Topics

- About Snapshots
- Take a VM Snapshot
- Revert to a Snapshot

- [Create a New VM From a Snapshot](#)
- [Create a New Template From a Snapshot](#)
- [Export a Snapshot to a File](#)
- [Delete a Snapshot](#)

About Snapshots

September 30, 2020

A virtual machine (VM) snapshot is a record of a running virtual machine at a point in time. When you take a snapshot of a VM, its storage information (the data on the hard drive) and metadata (configuration information) is also saved. Where necessary, I/O is temporarily halted while the snapshot is being taken to ensure that a self-consistent disk image can be captured.

Unlike VM exports, snapshots can be created without first shutting down the VM. A snapshot is similar to a normal VM template but it contains all the storage and configuration information for the original VM, including networking information. Snapshots provide a fast way of creating templates that can be exported for backup purposes and then restored, or that can be used to quickly create new VMs.

Snapshots are supported on all storage types, though for LVM-based storage types (XenServer version 5.5 onwards) the storage repository must have been upgraded if it was created on an older version of XenServer, and the volume must be in the default format; see [Upgrading older SRs](#).

Types of snapshots

XenCenter supports all three types of VM snapshots: disk-only, quiesced, and disk and memory. See [Take a VM Snapshot](#) for more information.

Disk-only snapshots

Disk-only snapshots store a VM's configuration information (metadata) and disks (storage), allowing them to be exported and restored for backup purposes. This type of snapshot is crash-consistent and can be performed on all VM types, including Linux VMs.

Quiesced snapshots

Quiesced snapshots take advantage of the Windows Volume Shadow Copy Service (VSS) to generate application-consistent point-in-time snapshots. The VSS framework helps VSS-aware applications (for example Microsoft Exchange or Microsoft SQL Server) flush data to disk and prepare for the snapshot before it is taken. Quiesced snapshots are therefore safer to restore, but can have a greater per-

formance impact on a system while they are being taken. They may also fail under load, so more than one attempt to take the snapshot may be required.

See the *XenServer Administrator's Guide* for information about operating system support for quiesced snapshots and for detailed information about how quiesced snapshots are implemented in XenServer.

> **Important:**
>
> The XenServer VSS provider used to quiesce the guest filesystem in preparation for a snapshot is installed as part of the XenServer Tools, but is not enabled by default. Follow the instructions in the XenServer Virtual Machine User's Guide to enable the VSS provider.

Disk and memory snapshots

In addition to saving the VM's metadata and disks, disk and memory snapshots also save the VM's memory state (RAM). Reverting back to a disk and memory snapshot does not require a reboot of the VM, and VMs can be running or suspended when the snapshot is taken. Disk and memory snapshots can be useful if you are upgrading or patching software, or want to test a new application, but also want the option to be able to get back to the current, pre-change state (RAM) of the VM.

Accessing orphaned snapshots

If you take snapshots of a VM and subsequently delete the original VM, you can still access those snapshots in the Resources pane. Switch to **Objects** view in the Navigation pane and then expand the **Snapshots** group to see all available snapshots.

Take a VM Snapshot

September 30, 2020

Perform the following steps to take a snapshot of a Virtual Machine:

1. Select the VM in the **Resources** pane and then click on the **Snapshots** tab.

2. Click the **Take Snapshot** button. Alternatively,

   - Right-click in the **Properties** pane and click **Take Snapshot**.
   - On the **VM** menu, click **Take Snapshot**.

3. Enter the name of the new snapshot and an optional description.

4. Under **Snapshot** mode, choose the type of snapshot to create:

   - To create a disk-only snapshot, select **Snapshot the virtual machine's disks**.

- To create a quiesced snapshot, select **Snapshot the virtual machine's disks** and then click **Quiesce the VM before taking the snapshot**.

- To create a disk and memory snapshot, select **Snapshot the virtual machine's disks and memory**.

5. Click **OK** to begin creating the new snapshot. Progress is displayed on the status bar and on the **Events** view under **Notifications**.

When the new snapshot has been created, it will appear on the VM's **Snapshots** tab and under the Snapshots group in the **Resources** pane in **Folder** View:

- A disk-only snapshot
- A disk and memory snapshot

Revert to a Snapshot

June 3, 2019

Reverting to a snapshot restores the VM to the state it was in at the point in time when the snapshot was created. All changes made to the VM since the snapshot was taken will be discarded and the current state of the VM will be lost.

The **Revert to Snapshot** dialog box includes an option to take a new snapshot of the current VM state before reverting back the earlier snapshot, allowing you to easily restore the VM to its current state again if you need to.

1. On the **Snapshots** tab, select the snapshot and click **Revert To**.

   If the snapshot you want to revert to is a scheduled snapshot (created automatically using the VM Protection and Recovery feature), you will need to make scheduled snapshots visible on the **Snapshots** tab before you can select it: click **View** and then **Scheduled Snapshots**.

2. To take a new snapshot of the current state of VM before reverting it back to the earlier snapshot, select the check box.

3. Click **Yes**.

Create a New VM From a Snapshot

June 3, 2019

Creating a new VM from a snapshot works in exactly the same way as creating one from a regular VM template, using the **New VM** wizard.

To create a new VM from a snapshot

1. On the **Snapshots** tab, select the snapshot you want to use, then right click and click **New VM from Snapshot** on the shortcut menu.

   The **New VM** wizard opens, with your snapshot pre-selected on the **Templates** page.

2. Follow the steps in the **New VM** wizard to create the new VM. See Creating a new VM.

To create a new VM from an orphan snapshot

If the original VM used to create the snapshot has been deleted, you can select the snapshot and start the **New VM** wizard as follows:

1. In the **Resources** pane, switch to Folder View.
2. Click to expand the **Types** group and then expand the **Snapshots** group.
3. Select the snapshot, then right-click and click **New VM from Snapshot** on the shortcut menu.

Create a New Template From a Snapshot

June 3, 2019

While it is not possible to copy a VM snapshot directly, you can create a new VM template from a snapshot and then use that to make copies of the snapshot. Templates are a "gold image" - ordinary VMs which are intended to be used as master copies from which to create new VMs. Once you have set up a VM the way you want it and taken a snapshot of it, you can then save the snapshot as a new template and use it to create copies of your specially-configured VM in the same resource pool. Note that the snapshot's memory state will not be saved when you do this.

To save a snapshot as a new template

1. On the **Snapshots** tab, select the snapshot, right click and then click **Create Template from Snapshot** on the shortcut menu.
2. Enter the name of the new template and then click **OK**. Once the new template has been successfully created, it appears as a custom template in the **Resources** pane and on the **Templates** page in the **New VM** wizard.

To save an orphan snapshot as a new template

If the original VM used to create the snapshot has been deleted, you can save it as a new template as follows:

---

1. In the **Resources** pane, switch to Folder View.
2. Click to expand the **Types** group and then expand the **Snapshots** group.
3. Select the snapshot, then right-click and click **Create Template from Snapshot** on the shortcut menu.

Export a Snapshot to a File

June 3, 2019

When you export a VM snapshot, it is saved as a VM template in a single XVA file on your local machine. This template contains a complete copy of the snapshot (including disk images) and it can then be imported and used to create a new VM in the same or in a different resource pool.

To export a snapshot to a file

1. On the **Snapshots** tab, select the snapshot, click **Actions** and then click **Export to File**.
2. Browse to locate the folder where you want to create the XVA file, enter the filename, then click **Save** to begin the export.

To export an orphan snapshot

If the original VM used to create the snapshot has been deleted, you can export the snapshot as follows:

1. In the **Resources** pane, switch to Folder View.
2. Click to expand the **Types** group and then expand the **Snapshots** group.
3. Select the snapshot then right-click and select **Export Snapshot as Template** on the shortcut menu.
4. Browse to locate the folder where you want to create the XVA file, enter the filename, then click **Save** to begin the export.

Delete a Snapshot

June 3, 2019

To delete a snapshot

1. On the **Snapshots** tab, select the snapshot and click **Delete**.

2. Click **OK** to confirm.

To delete an orphan snapshot

If the original VM used to create the snapshot has been deleted, you can delete the snapshot as follows:

1. In the **Resources** pane, switch to Folder View.
2. Click to expand the **Types** group and then expand the **Snapshots** group.
3. Select the snapshot, right-click and then click **Delete Snapshot** on the shortcut menu.

XenServer vApps

June 3, 2019

Topics

- Managing vApps
- Create a vApp
- Modify vApps
- Delete a vApp
- Start and Shut Down vApps
- Export and Import vApps

Managing vApps

September 30, 2020

A XenServer vApp is a logical group of one or more related virtual machines (VMs) which can be managed as a single entity. The VMs within a vApp do not have to reside on one host server and will be distributed within a pool using the normal rules.

When a vApp is started, the VMs contained within it will start in a user-defined order, allowing VMs which depend upon one another to be automatically sequenced. This means that you do not need to manually sequence the startup of dependent VMs should a whole service require restarting, for instance in the case of a software update.

Using the Manage vApps dialog box

Use the XenCenter **Manage vApps** dialog box you can create, delete and modify vApps, start and shutdown vApps, and import and export vApps within the selected pool. When you select a vApp in the

list, the VMs it contains are listed in the details pane on the right.

Control	Function
New vApp	Opens the **New vApp** wizard. See Create a vApp.
Delete	Deletes the selected vApp. The VMs in the vApp will not be deleted.
Properties	Opens a **Properties** dialog box for the selected vApp. Here you can change its name or description, add or remove VMs from the vApp, and change their start order and delay interval. See Modify vApps.
Start	Starts up all of the VMs in the selected vApp in the sequence specified by the start order and delay interval values set on each individual VM. See Start and shut down vApps.
Shut Down	Shut down all of the VMs in the selected vApp. See Start and shut down vApps.
Import	Open **Import** wizard and import an OVF/OVA package as a vApp. See Export and import vApps.
Export	Open the Export wizard and export a vApp as an OVF/OVA package. See Export and import vApps.

Create a vApp

June 3, 2019

To create a new vApp, use the Manage vApps dialog box.

1. Select the pool and, from the **Pool** menu, click **Manage vApps**.

   Alternatively, right-click in the **Resources** pane and click **Manage vApps** on the shortcut menu.

2. Click **New vApp...** on the top left corner of the Manage vApps dialog box.

3. Enter the name of the new vApp and (optionally) a description, then click **Next**. You can choose any name you like, but a descriptive name is usually best. Although it is advisable to avoid

having multiple vApps with the same name, it is not a requirement, and XenCenter does not enforce any uniqueness constraints on vApp names. It is not necessary to use quotation marks for names that include spaces.

4. Choose which virtual machines to include in the new vApp and then click **Next**. You can use the **Search** box to list only VMs with names that include the specified string.

5. Specify the startup sequence for the VMs in the vApp, and then click **Next**.

Value	Description
Start order	Specifies the order in which individual VMs will be started up within the vApp, allowing certain VMs to be restarted before others. VMs with a start order value of 0 (zero) will be started first, then VMs with a start order value of 1, then VMs with a start order value of 2, and so on.
Attempt to start next VM after	This is a delay interval that specifies how long to wait after starting the VM before attempting to start the next group of VMs in the startup sequence, that is, VMs with a lower start order.

Note that the shutdown order of VMs in a vApp is always the reverse of the configured start order.

6. On the final page of the wizard, you can review the vApp configuration. Click **Previous** to go back and modify any settings, or **Finish** to create the new vApp and close the wizard.

Modify vApps

June 3, 2019

To change the name or description of a vApp, add or remove VMs from the vApp, and change the startup sequence of the VMs in the vApp, use the Manage vApps dialog box.

1. Select the pool and, on the **Pool** menu, click **Manage vApps**.

   Alternatively, right-click in the **Resources** pane and click **Manage vApps** on the shortcut menu.

2. Select the vApp and click **Properties** to open its **Properties** dialog box

3. Click the **General** tab to change the vApp name or description.

4. Click the **Virtual Machines** tab to add or remove VMs from the vApp.

5. Click the **VM Startup Sequence** tab to change the start order and delay interval values for individual VMs in the vApp.

Control	Description
Start order	Specifies the order in which individual VMs will be started up within the vApp, allowing certain VMs to be restarted before others. VMs with a start order value of 0 (zero) will be started first, then VMs with a start order value of 1, then VMs with a start order value of 2, and so on.
Attempt to start next VM after	This is a delay interval that specifies how long to wait after starting the VM before attempting to start the next group of VMs in the startup sequence, that is, VMs with a lower start order.

Note that the shutdown order of VMs in a vApp is always the reverse of the configured start order.

6. Click **OK** to save your changes and close the **Properties** dialog box.

Delete a vApp

September 30, 2020

To delete a vApp from a pool, use the **Manage vApps** dialog box.

1. Select the pool and, on the **Pool** menu, click **Manage vApps**.

   Alternatively, right-click in the **Resources** pane and click **Manage vApps** on the shortcut menu.

2. Select the vApp you want to delete from the list, then click 🗑 **Delete**.

   The VMs in the vApp will not be deleted.

Start and Shut Down vApps

September 30, 2020

To start or shut down a vApp, use the **Manage vApps** dialog box, accessed from the **Pool** menu.

When you start a vApp, all the VMs within it are started up automatically in sequence. The start order and delay interval values specified for each individual VM control the startup sequence; these values can be set when you first create the vApp and changed at any time from the vApp Properties dialog box or from the individual VM's Properties dialog box.

The shutdown order of VMs in a vApp is always the reverse of the configured start order.

To start a vApp

1. Open the **Manage vApps** dialog box: select the pool where the VMs in the vApp are located and, on the **Pool** menu, click **Manage vApps**.

   Alternatively, right-click in the **Resources** pane and click **Manage vApps** on the shortcut menu.

2. Select the vApp and click ⏻ **Start** to start all of the VMs it contains.

To shut down a vApp

1. Open the **Manage vApps** dialog box: select the pool where the VMs in the vApp are located and, on the **Pool** menu, click **Manage vApps**.

   Alternatively, right-click in the Resources pane and click **Manage vApps** on the shortcut menu.

2. Select the vApp and click ⏻ **Shut Down** to shut down all of the VMs in the vApp. A soft shut down will be attempted on all VMs; if this is not possible, then a forced shut down will be performed. See Shut down a VM to find out more about soft and forced VM shutdowns.

Export and Import vApps

September 30, 2020

To export a vApp

vApps can be exported as OVF/OVA packages.

1. Open the **Manage vApps** dialog box: on the **Pool** menu, click **Manage vApps**.

2. Select the vApp you want to export in the list and click 🖳 **Export**.

3. Follow the procedure described in Export VMs as OVF/OVA.

   Exporting a vApp may take some time.

To import a vApp

OVF/OVA packages are imported as vApps.

1. Open the **Manage vApps** dialog box: on the **Pool** menu, click **Manage vApps**.

2. Click 🖥️ **Import** to open the **Import** wizard.

3. Follow the procedure described in Import VMs from OVF/OVA.

   When the import is complete, the new vApp appears in the list of vApps in the **Manage vApps** dialog box.

Protecting VMs and vApps

September 30, 2020

XenServer offers a range of features to enable you to protect your VMs and vApps.

VM Protection and Recovery (VMPR) - scheduled snapshots

XenServer VM Protection and Recovery (VMPR) provides protection against loss of individual VMs by taking regular scheduled VM snapshots which you can import back into XenCenter and use to recreate the original VM if it is corrupted or inaccessible.

Read more: About VM Protection and Recovery ; Create a VM Protection policy ; Manage VM Protection policies ; Recovering VMs from snapshots

> **Important:**
>
> VMPR and its functionality has been removed from XenServer version 6.2.0 and higher. Applications, code, or usage that depend on VMPR will not function in XenServer version 6.2.0 and higher. The VM snapshot functionality and functionalities that depend on VM snapshot (other than VMPR) remain unaffected.

High Availability (HA)

XenServer High Availability (HA) provides protection against downtime of critical VMs caused by the failure of individual servers in a pool by guaranteeing that VMs are automatically restarted on an alternate server in the same pool, with minimal service interruption. The pool database is constantly replicated across all nodes and is also backed up to shared storage on the heartbeat SR for additional safety.

Read more: About XenServer HA ; HA requirements ; VM startup settings ; Configure HA ; Disable HA ; Change HA settings

---

Disaster Recovery (DR)

XenServer Disaster Recovery (DR) provides protection against the loss of multiple servers at your primary data site. With DR enabled, the pool database is constantly replicated via mirrored storage, and in the event of a disaster at your primary site, VMs and vApps can be recovered from the mirrored storage to a pool on a secondary (DR) site.

Read more: About XenServer DR ; Configuring DR ; Failover ; Failback ; Test Failover

VM Protection and Recovery (VMPR)

September 30, 2020

> **Important:**
>
> VMPR and its functionality has been removed from XenServer version 6.2.0 and higher. Applications, code, or usage that depend on VMPR will not function in XenServer version 6.2.0 and higher. The VM snapshot functionality and functionalities that depend on VM snapshot (other than VMPR) remain unaffected.

Topics:

- About VM Protection and Recovery
- Create a VM Protection Policy
- Manage VM Protection Policies
- Recover VMs From Scheduled Snapshots

About VM Protection and Recovery

September 30, 2020

> **Important:**
>
> VMPR and its functionality has been removed from XenServer version 6.2.0 and higher. Applications, code, or usage that depend on VMPR will not function in XenServer version 6.2.0 and higher. The VM snapshot functionality and functionalities that depend on VM snapshot (other than VMPR) remain unaffected.

The VM Protection and Recovery (VMPR) feature provides a simple backup and restore utility for your critical service VMs. Regular scheduled snapshots are taken automatically and can be used to restore individual VMs.

VMPR works by having pool-wide VM protection policies that define snapshot schedules for selected VMs in the pool. When a policy is enabled, snapshots are taken of the specified VMs at the scheduled time each hour, day or week. If configured, these snapshots can also be archived automatically to a remote CIFS or NFS share, providing an additional level of security.

Several VMPR policies may be enabled in a pool, covering different VMs and with different schedules. A VM can be assigned to only one policy at a time.

XenCenter provides a range of tools to help you use the VMPR feature:

- To define VMPR policies, use the **New VM Protection Policy** wizard.
- To enable, disable, edit and delete VMPR policies for a pool, you use the **VM Protection Policies** dialog box.
- To edit a policy, open its Properties dialog box from the **VM Protection Policies** dialog box.
- To recover a VM from a scheduled snapshot, you simply select the snapshot on the **Snapshots** tab and revert the VM to it.
- To recover a VM from an archived scheduled snapshot, you import the snapshot, creating a new custom VM template, and then create a new VM from the template.

Create a VM Protection Policy

September 30, 2020

> **Important:**
>
> VMPR and its functionality has been removed from XenServer version 6.2.0 and higher. Applications, code, or usage that depend on VMPR will not function in XenServer 6.2.0 and higher. The VM snapshot functionality and functionalities that depend on VM snapshot (other than VMPR) remain unaffected.

Use the **New VM Protection Policy** wizard to create a VM Protection policy that identifies the VMs in the pool that you want to snapshot, the type of snapshot to be created (disk-only or disk and memory), and the snapshot and archive schedules. You can also configure email notification for snapshot and archive event alerts in the wizard.

To open the **New VM Protection Policy** wizard: on the **Pool** menu, click VM Protection Policies to open the **VM Protection Policies** dialog box, and then click **New** to start the wizard.

Selecting the VMs to protect

All the VMs in the pool are listed on this page of the wizard. If a VM is already included in a VM Protection policy, the name of the policy is shown. A VM can be assigned to only one policy at a time, so if you assign it to the new policy, it will be removed from its current policy (you will be asked to confirm this).

Selecting snapshot type

Scheduled VM snapshots can be either disk-only snapshots or disk and memory snapshots.

- Disk-only snapshots store the VM's disks (storage) and metadata. They are crash-consistent and can be performed on all VM types, including Linux VMs.
- Disk and memory snapshots save the VM's disks (storage), metadata, and its current memory state (RAM). Note that this type of snapshot can be very large.

Setting the snapshot schedule and retention period

Snapshot schedule options:

- **Hourly snapshots** A snapshot of the specified VM or VMs will be taken each hour at the specified time.
- **Daily snapshots** A snapshot of the specified VM or VMs will be taken each day at the specified time.
- **Weekly snapshots** A snapshot of the specified VM or VMs will be taken at the specified time on the specified days of the week. You can select one or more days.

Snapshot retention policy

Under **Number of snapshots to keep** , specify how many scheduled snapshots you want to retain. When the number of scheduled snapshots taken exceeds this value, the oldest one will be deleted automatically.

> **Note:**
>
> The upper limit for the **Maximum number of scheduled snapshots** is 10.

Setting an archive schedule

Scheduled snapshots can be automatically archived to a remote CIFS or NFS share, if required. You can define a daily or weekly archive schedule that is independent of the snapshot schedule, or simply have every scheduled snapshot archived as soon as it is taken (the **Archive after every scheduled snapshot** option). You cannot define an archive schedule that is more frequent than the snapshot schedule.

Archive naming conventions

The archive folder name combines the VM name and the first 16 characters of the VM UUID. For example:

VM Name	VM UUID	Archive folder name
Win7_Test_1	cb53200c-bbd8-4c12-a076-e2eb29b38f06	Win7_Test_1-cb53200c-bbd8-4c

Archived scheduled snapshots placed in the archive folder are named `YYYYMMDD-HHMM.xva`. For example:

```
1  20100624-1830.xva
2  20100624-1836.xva
3  20100625-0031.xva
4  <!--NeedCopy-->
```

Configuring email alerts

Email notifications can be sent when system alerts are generated for any scheduled snapshot or archive operations. By default, VMPR event email notification is turned off.

Manage VM Protection Policies

September 30, 2020

> **Important:**
>
> VMPR and its functionality has been removed from XenServer version 6.2.0 and higher. Applications, code, or usage that depend on VMPR will not function in XenServer version 6.2.0 and higher. The VM snapshot functionality and functionalities that depend on VM snapshot (other than VMPR) remain unaffected.

To enable, disable, edit and delete VM Protection policies for a pool, you use the **VM Protection Policies** dialog box: on the Pool menu, click **VM Protection Policies**.

Enabling a VM Protection policy

When you enable a VM Protection policy, you turn it "on": automated snapshots of the specified VMs will then be generated at the scheduled time and archived, if this is also configured. Scheduled snapshots will continue being taken until the policy is disabled.

To enable a policy:

---

1. Select the pool or any server or VM in the pool in the **Resources** pane and on the **Pool** menu, select **VM Protection Policies**.
2. Select a policy from the list of policies defined in the pool and click **Enable**.

Disabling a VM Protection policy

If you want to stop automated VM snapshots from being taken, you can disable the VM Protection policy using the VM Protection Policies dialog box: a disabled policy can be enabled again at any time.

To disable a policy:

1. Select the pool or any server or VM in the pool in the **Resources** pane and on the **Pool** menu, select **VM Protection Policies**.
2. Select the policy from the list of policies defined in the pool and click **Disable**.

Editing a VM Protection policy

You can change a VMPR policy, for example, to add more VMs or to change the snapshot schedule:

1. Select the pool or any server or VM in the pool in the **Resources** pane and on the **Pool** menu, select **VM Protection Policies**.

2. Select the policy from the list of policies defined in the pool and click **Properties** and go to the tab you need:

   - **Protected VMs**: Add or remove VMs on this tab. A VM can be assigned to only one policy at a time, so if you assign it to a new policy, it will be removed from its current policy.
   - **Snapshot Type**: Change the type of snapshot that will be created.
   - **Snapshot Schedule**: Change the schedule used to take VM snapshots and change the snapshot retention policy.
   - **Archive Options**: Turn automatic archiving on and off, and set the archive schedule.
   - **Email Alerts**: Configure email notifications for scheduled snapshot or archive operation system alerts.

3. Click **OK** to save your changes and close the **Properties** window.

Delete a VM Protection policy

To delete a VM Protection policy:

1. Select the pool, or any server or VM in the pool, in the **Resources** pane and on the **Pool** menu, select **VM Protection Policies**.
2. Select the policy from the list of policies defined in the pool and click **Delete**.

Recovering VMs From Snapshots

September 30, 2020

To recover a VM from a scheduled snapshot

Important:

VMPR and its functionality has been removed from XenServer version 6.2.0 and higher. Applications, code, or usage that depend on VMPR will not function in XenServer version 6.2.0 and higher. The VM snapshot functionality and functionalities that depend on VM snapshot (other than VMPR) remain unaffected.

To recover a VM from a scheduled snapshot, you simply revert the VM to the specified snapshot.

1. Select the VM and click on the **Snapshots** tab.
2. To show scheduled snapshots (by default, they are not shown on this tab): click **View** and then **Scheduled Snapshots**.
3. Select the scheduled snapshot you want to revert the VM to and then click **Revert To**.
4. To take a new snapshot of the current state of VM before reverting it back to the scheduled snapshot, select the check box.
5. Click **Yes** to revert the VM back to the selected snapshot.

To recover a VM from an archived snapshot

To recover a VM from an archived scheduled snapshot, follow these steps:

1. Mount the CIFS or NFS archive share target locally.

2. Import the archive XVA file from the target using the **Import** wizard, as described in Import VMs from XVA. This will create a new custom VM template.

3. Create a new VM based on the new template using the **New VM** wizard, as described in Creating a new VM.

   See Archive naming conventions for information on how archive folders and archived snapshots are named.

High Availability

June 3, 2019

Topics:

About XenServer HA

June 3, 2019

XenServer High Availability (HA) allows virtual machines to automatically be restarted in the event of an underlying hardware failure or loss of any managed server. HA is about making sure that important VMs are always running in a resource pool. With HA enabled, if one of your servers fails, its VMs will be intelligently restarted on other servers in the same pool, allowing essential services to be restored in the event of system or component failure with minimal service interruption. If the pool master server fails, XenServer HA automatically selects a new server to take over as master, so you can continue to manage the pool. Any server in a pool can be a master server, and the pool database is constantly replicated across all nodes and also backed up to shared storage on the heartbeat SR for additional safety.

There are two key aspects to XenServer HA: reliably detecting server failure, and computing a failure plan to enable swift recovery, and these are covered in detail below.

Heartbeats for availability

Detecting server failure reliably is difficult since you need to remotely distinguish between a server disappearing for a while versus catastrophic failure. If we mistakenly decide that a master server has broken down and elect a new master in its place, there may be unpredictable results if the original server were to make a comeback. Similarly, if there is a network issue and a resource pool splits into two equal halves, we need to ensure that only one half accesses the shared storage and not both simultaneously. XenServer solves all these problems by having two mechanisms: a storage heartbeat and a network heartbeat.

When you enable HA in a pool, you nominate an iSCSI, Fibre Channel or NFS storage repository to be the heartbeat SR. XenServer automatically creates a couple of small virtual disks in this SR. The first disk is used by every server in the resource pool as a shared quorum disk. Each server allocates itself a unique block in the shared disk and regularly writes to the block to indicate that it is alive. When HA starts up, all servers exchange data over both network and storage channels, indicating which servers they can see over both channels - that is, which I/O paths are working and which are not. This

information is exchanged until a fixed point is reached and all of the servers in the pool are satisfied that they are in agreement about what they can see. When this happens, HA is enabled and the pool is protected. This HA arming process can take a few minutes to settle for larger pools, but is only required when HA is first enabled.

Once HA is active, each server regularly writes storage updates to the heartbeat virtual disk, and network packets over the management interface. It is vital to ensure that network adapters are bonded for resilience, and that storage interfaces are using dynamic multipathing where supported. This will ensure that any single adapter or wiring failures do not result in any availability issues.

Server fencing

The worst-case scenario for HA is the situation where a server is thought to be off-line but is actually still writing to the shared storage, since this can result in corruption of persistent data. In order to prevent this situation, XenServer uses server fencing , that is, the server is automatically powered off and isolated from accessing any shared resources in the pool. This prevents the failing server from writing to any shared disks and damaging the consistency of the stored data during automated failover, when protected virtual machines are being moved to other, healthy servers in the pool.

Servers will self-fence (that is, power off and restart) in the event of any heartbeat failure unless any of the following hold true:

- The storage heartbeat is present for all servers but the network has partitioned (so that there are now two groups of servers). In this case, all of the servers that are members of the largest network partition stay running, and the servers in the smaller network partition self-fence. The assumption here is that the network outage has isolated the VMs, and they ought to be restarted on a server with working networking. If the network partitions are exactly the same size, then only one of them will self-fence according to a stable selection function.
- If the storage heartbeat goes away but the network heartbeat remains, then the servers check to see if they can see all other servers over the network. If this condition holds true, then the servers remain running on the assumption that the storage heartbeat server has gone away. This doesn't compromise VM safety, but any network glitches will result in fencing, since that would mean both heartbeats have disappeared.

Capacity planning for failure

The heartbeat system gives us reliable notification of server failure, and so we move onto the second step of HA: capacity planning for failure.

A resource pool consists of several servers (say, 16), each with potentially different amounts of memory and a different number of running VMs. In order to ensure that no single server failure will make it impossible to restart its VMs on another server (for example, due to insufficient memory on any other

215

server), XenServer HA dynamically computes a failure plan which calculates the actions that would be taken on any server failure. In addition to dealing with failure of a single server, XenServer HA can deal with the loss of multiple servers in a pool, for example when failure of a network partition takes out an entire group of servers.

In addition to calculating what actions will be taken, the failure plan considers the number of server failures that can be tolerated in the pool. There are two important considerations involved in calculating the HA plan for a pool:

- Maximum failure capacity . This is the maximum number of servers that can fail before there are insufficient resources to run all the protected VMs in the pool; this value is calculated by XenServer by taking account of the restart priorities of the VMs in the pool, and the pool configuration (the number of servers and their CPU and memory capacity).
- Server failure limit . This is a value that you can define as part of HA configuration which specifies the number of server failures that you want to allow in the pool, within the HA plan. For example, in a resource pool of 16 servers, when you set the server failure limit to 3, XenServer calculates a failover plan that allows for any 3 servers to fail and still be able to run all protected VMs in the pool. You can configure the server failure limit to a value that is lower than the maximum failure capacity, making it is less likely that the pool will become overcommitted. This can be useful in an environment with RBAC enabled, for example, to allow RBAC users without Pool Operator permissions to bring more VMs online without breaking the HA plan; see HA and Role-Based Access Control (RBAC) below.

A system alert will be generated when the maximum failure capacity value falls below the value specified for the server failure limit.

Overcommit protection

When HA is first enabled on a pool, a failure plan is calculated based on the resources available at that time. XenServer HA dynamically calculates a new failure plan in response to events which would affect the pool, for example, starting a new VM. If a new plan cannot be calculated due to insufficient resources across the pool (for example not enough free memory or changes to virtual disks and networks that affect which VMs may be restarted on which servers), the pool becomes overcommitted .

HA restart priority is used to determine which VMs should be started when a pool is overcommitted. When you configure the restart priority for the VMs you want to protect in the **HA Configuration** dialog box or in the **Configure HA** wizard, you can see the maximum failure capacity for the pool being recalculated dynamically, allowing you to try various combinations of VM restart priorities depending on your business needs, and see if the maximum failure capacity is appropriate to the level of protection you need for the critical VMs in the pool.

If you attempt to start or resume a VM and that action would cause the pool to be overcommitted, a

warning will be displayed in XenCenter. The message may also be sent to an email address, if configured. You will then be allowed to cancel the operation, or proceed anyway, causing the pool to become overcommitted.

Working with an HA-enabled pool

The best practice for HA is not to make configuration changes to the pool while HA is enabled. Instead, it is intended to be the "2am safeguard" which will restart servers in the event of a problem when there isn't a human administrator nearby. If you are actively making configuration changes in the pool such as applying software updates, then HA should be disabled for the duration of these changes.

- If you try to shut down a protected VM from XenCenter, XenCenter will offer you the option of removing the VM from the pool failure plan first and then shutting it down. This ensures that accidental VM shutdowns will not result in downtime, but that you can still stop a protected VM if you really want to.
- If you need to reboot a server when HA is enabled, XenCenter automatically uses the VM restart priorities to determine if this would invalidate the pool failure plan. If it doesn't affect the plan, then the server is shut down normally. If the plan would be violated, but the maximum failure capacity is greater than 1, XenCenter will give you the option of lowering the pool's server failure limit by 1. This reduces the overall resilience of the pool, but always ensures that at least one server failure will be tolerated. When the server comes back up, the plan is automatically recalculated and the original server failure limit is restored if appropriate.
- When you install software updates using the **Install Update** wizard, you must disable HA on the pool by clicking on the **Turn HA off** option until after the update has been installed. If you do not disable HA, the update will not proceed. You will need to monitor the pool manually while updates are being installed to ensure that server failures do not disrupt the operation of the pool.
- When HA is enabled, some operations that would compromise the plan for restarting VMs may be disabled, such as removing a server from a pool. To perform these operations, you should temporarily disable HA or you can shut down the protected VMs before continuing.

HA and Role-Based Access Control (RBAC)

In a XenServer environment where Role-Based Access Control (RBAC) is implemented, not all users will be permitted to change a pool's HA configuration settings. Users who can start VMs (VM Operators), for example, will not have sufficient permissions to adjust the failover capacity for an HA-enabled pool. For example, if starting a VM reduces the maximum number of server failures allowed to a value lower than the current maximum failure capacity, the VM Operator will not be able to start the VM. Only Pool Administrator or Pool Operator-level users can configure the number of server failures allowed.

In this case, the user who enables HA (with a Pool Administrator or Pool Operator role) can set the

server failure limit to a number that is actually lower than the maximum number of failures allowed. This creates slack capacity and so ensures that less privileged users can start up new VMs and reduce the pool's failover capacity without threatening the failure plan.

HA Requirements

June 3, 2019

Before you can configure HA on a resource pool, you must ensure that the following requirements are satisfied for all servers and virtual machines in the pool:

- Shared storage must be available, including at least one iSCSI, Fibre Channel or NFS LUN of 356 MiB or greater which will be used for the heartbeat SR. If you are using a NetApp or EqualLogic storage array, then you should manually provision an iSCSI LUN on the array to use for the heartbeat SR.

- We strongly recommend the use of a bonded management interface on the servers in the pool if HA is enabled, and multipath storage for the heartbeat SR.

- Adequate licenses must be installed on all servers.

- All the virtual machines you want to protect with HA must be agile . This means:

  – Virtual disks must be on shared storage. Any type of shared storage may be used; the iSCSI, Fibre Channel or NFS LUN is only required for the storage heartbeat and can be used for virtual disk storage if you prefer, but this is not necessary.
  – Virtual network interfaces must be on pool-wide networks.
  – No connection should be configured to any local DVD drive.

  VMs that are not agile (for example, with a physical CD drive mapped in from a server) can only be assigned a **Restart if possible** HA restart priority since they are tied to one server.

VM Restart Settings

June 3, 2019

If more servers fail than have been planned for, then an HA recovery operation begins. The **HA restart priority** is used to determine which VMs are restarted, while the order in which individual VMs are started is determined by their start order and delay interval values, ensuring that the most important VMs are restarted first.

HA restart priority

The **HA restart priority** specifies which VMs will be restarted under the HA failure plan for a pool:

Value	Description
Restart	VMs with this priority are guaranteed to be restarted if sufficient resources are available within the pool. They will be restarted before VMs with a **Restart if possible** priority. All VMs with this restart priority are considered when calculating a failure plan. If no plan exists for which all VMs with this priority can be reliably restarted, then the pool is considered to be overcommitted.
Restart if possible	VMs with this restart priority are not considered when calculating a failure plan, but one attempt to restart them will be made if a server that is running them fails. This restart is attempted after all higher-priority VMs are restarted, and if the attempt to start them fails, then it will not be retried. This is a useful setting for test/development VMs which aren't critical to keep running, but would be nice to do so in a pool which also has some important VMs which absolutely must run.
Do not restart	No attempts will be made to restart VMs with this priority.

Start order

The **Start order** specifies the order in which individual VMs will be started up during an HA recovery operation, allowing certain VMs to be started before others. VMs with a start order value of 0 (zero) will be started first, then VMs with a start order value of 1, followed by VMs with a start order value of 2, and so on.

Delay interval (Attempt to start next VM after)

The VM property **Attempt to start next VM after** specifies how long to wait after starting the VM before attempting to start the next group of VMs in the startup sequence, that is, VMs with a later start order.

Configure HA

June 3, 2019

You enable HA for a resource pool using the **Configure HA** wizard. The wizard takes you through the HA configuration process, step-by-step, calculating the server failure limit for the pool given the available resources and the HA restart priorities you specify.

To open the **Configure HA** wizard: in XenCenter, select the pool, click on the **HA** tab, and then click **Configure HA**.

Alternatively:

- On the **Pool** menu, click **High Availability**.
- Right-click in the **Resources** pane and then click **High Availability** on the shortcut menu.

To configure HA on a pool:

1. Ensure that the HA requirements identified in HA requirements are satisfied.

2. Open the **Configure HA** wizard and click **Next** on the first page of wizard to scan the pool for a shared iSCSI, Fibre Channel or NFS LUN that can be used as the pool's heartbeat SR. If no suitable SR is found, you will need to configure some appropriate new storage before you can continue.

3. On the **Heartbeat SR** page, choose an SR from the list and then click **Next**.

4. On the **HA Plan** page, select one or more VMs in the list and set the required VM startup settings. For more on these options, see VM startup settings. To select more than one virtual machine, click at the start of the selection, scroll to the end of the selection, and then hold down SHIFT while you click where you want the selection to end. To select several virtual machines that are not next to each other in the list, click the first one, press CTRL, and then click the additional virtual machines that you want to select.

   Set the following options:

   - **HA restart priority** : Choose a restart priority for each VM:
     - Choose **Restart** to ensure the selected VM(s) are restarted if sufficient resources are available within the pool.
     - Choose **Restart if Possible** if it is not essential to restart the VM automatically.
     - Choose **Do Not Restart** if you never want the VM to be restarted automatically.
   - **Start order** : Specifies the order in which individual VMs will be started up during the HA recovery operation, allowing certain VMs to be started before others. VMs with a start order value of 0 (zero) will be started first, then VMs with a start order value of 1, followed by VMs with a start order value of 2, and so on.
   - **Attempt to start next VM after** : This is a delay interval that specifies how long to wait after starting the VM before attempting to start the next group of VMs in the startup sequence,

that is, VMs with a lower start order.

5. Also on the **HA Plan** page, under **Server failure limit** , you can set the number of server failures to allow within this HA plan. This value should be less than or equal to the maximum failure capacity for the pool, shown here as max . If max is 0 (zero), the pool is overcommitted, and you will not be able to continue to the next page of the wizard without either adjusting the HA restart priorities or making more resources available within the pool; see To increase the maximum failure capacity for a pool. Click **Next** when you have finished HA plan configuration.

6. On the last page of the wizard, review your HA configuration settings. Click **Back** to go back and change any of the settings or click **Finish** to enable HA and close the wizard.

Disable HA

June 3, 2019

When HA is enabled, some operations that would compromise the plan for restarting virtual machines may be disabled, such as removing a server from a pool. To perform these operations, you can temporarily disable HA.

To disable HA:

1. Select the pool in the **Resources** pane, click on the **HA** tab, and then click **Disable HA**.
2. Click **OK** to confirm. The VM startup settings specified for each VM in the pool are stored and will be remembered if you turn HA back on again later.

Change HA Settings

June 3, 2019

Once HA has been enabled on a pool, you can change VM startup settings and adjust the server failure limit for the pool from the **Configure HA** dialog box.

To change HA restart priority and VM startup sequence settings

1. Select the pool in the **Resources** pane, click on the **HA** tab, and then click **Configure HA**. Alternatively:

   - On the **Pool** menu, click **High Availability**.
   - Right-click in the **Resources** pane and then click **High Availability** on the shortcut menu.

2. Select one or more VMs in the list and set the required VM startup settings. For more on these options, see VM startup settings. To select more than one virtual machine, click at the start of the selection, scroll to the end of the selection, and then hold down SHIFT while you click where you want the selection to end. To select several virtual machines that are not next to each other in the list, click the first one, press CTRL, and then click the additional virtual machines that you want to select.

Set the following options:

- **HA restart priority** : Choose a restart priority for each VM:
    - Choose **Restart** to ensure the selected VM(s) are restarted if sufficient resources are available within the pool.
    - Choose **Restart if Possible** if it is not essential to restart the VM automatically.
    - Choose **Do Not Restart** if you never want the VM to be restarted automatically.
- **Start order** : Specifies the order in which individual VMs will be started up during the HA recovery operation, allowing certain VMs to be started before others. VMs with a start order value of 0 (zero) will be started first, then VMs with a start order value of 1, followed by VMs with a start order value of 2, and so on.
- **Attempt to start next VM after** : This is a delay interval that specifies how long to wait after starting the VM before attempting to start the next group of VMs in the startup sequence, that is, VMs with a lower start order.

3. Click **OK** to apply the changes and close the dialog box.

To change the server failure limit for a pool

1. Select the pool in the **Resources** pane, click on the **HA** tab, and then click **Configure HA**. Alternatively:
    - On the **Pool** menu, click **High Availability**.
    - Right-click in the **Resources** pane and then click **High Availability** on the shortcut menu.
2. Under **Server failure limit** , enter the number of server failures to allow. This value should be less than or equal to the maximum failure capacity for the pool, shown here as max . If max is 0 (zero), the pool is overcommitted, and you will not be able to save the change without either adjusting the HA restart priorities or making more resources available within the pool; see To increase the maximum failure capacity for a pool.
3. Click **OK** to apply the changes and close the dialog box.

To increase the maximum failure capacity for a pool

To increase the maximum failure capacity for a pool, you need to do one or more of the following:

- Reduce the HA restart priority of some VMs.

- Increase the amount of RAM on your servers or add more servers to the pool to increase its capacity.
- Reduce the amount of memory configured on some VMs.
- Shut down non-essential VMs.

Disaster Recovery (DR)

June 3, 2019

Topics:

- About XenServer DR
- Configuring DR
- Failover
- Failback
- Test Failover

About XenServer DR

September 30, 2020

The XenServer Disaster Recovery (DR) feature is designed to allow you to recover virtual machines (VMs) and vApps from a catastrophic failure of hardware which disables or destroys a whole pool or site. For protection against single server failures, you can use XenServer High Availability to have VMs restarted on an alternate server in the same pool.

Understanding XenServer DR

XenServer DR works by storing all the information needed to recover your business-critical VMs and vApps on storage repositories (SRs) that are then replicated from your primary (production) environment to a backup environment. When a protected pool at your primary site goes down, the VMs and vApps in that pool can be recovered from the replicated storage and recreated on a secondary (DR) site, with minimal application or user downtime.

Once the recovered VMs are up and running in the DR pool, the DR pool metadata must also be saved on storage that is replicated, allowing recovered VMs and vApps to be restored back to the primary site when it is back online.

> **Note:**
>
> XenServer DR can only be used with LVM over HBA or LVM over iSCSI storage types.

XenServer VMs consists of two components:

- Virtual disks that are being used by the VM, stored on configured storage repositories (SRs) in the pool where the VMs are located.
- Metadata describing the VM environment. This is all the information needed to recreate the VM if the original VM is unavailable or corrupted. Most metadata configuration data is written when the VM is created and is updated only when you make changes to the VM configuration. For VMs in a pool, a copy of this metadata is stored on every server in the pool.

In a DR environment, VMs are recreated on a secondary (DR) site from the pool metadata - configuration information about all the VMs and vApps in the pool. The metadata for each VM includes its name, description and Universal Unique Identifier (UUID), and its memory, virtual CPU and networking and storage configuration. It also includes the VM's startup options - start order, delay interval and HA restart priority - which are used when restarting the VM in an HA or DR environment. For example, when recovering VMs during disaster recovery, the VMs within a vApp will be restarted in the DR pool in the order specified in the VM metadata, and with the specified delay intervals.

XenServer DR requirements

Software version	XenServer version 6.0 or later
Access	You must be logged in as root or have a role of Pool Operator or higher.

Disaster recovery infrastructure

To use XenServer DR, the appropriate DR infrastructure needs to be set up at both the primary and secondary sites:

- The storage used for both the pool metadata and the virtual disks used by the VMs must be replicated from your primary (production) environment to a backup environment. Storage replication, for example using mirroring, is best handled by your storage solution, and will vary from device to device.
- Once VMs and vApps have been recovered to a pool on your DR site and are up and running, the SRs containing the DR pool metadata and virtual disks must also be replicated to allow the recovered VMs and vApps to be restored back to the primary site (failed back) once the primary site is back online.

- The hardware infrastructure at your DR site does not have to match the primary site, but the XenServer environment must be at the same release and patch level, and sufficient resources should be configured in the target pool to allow all the failed over VMs to be re-created and started.

Important:

XenCenter and the **Disaster Recovery** wizard do not control any storage array functionality. Users of the Disaster Recovery feature must ensure that the pool metadata and the storage used by the VMs which are to be restarted in the event of a disaster are replicated to a backup site. Some storage arrays contain mirroring features to achieve the copy automatically: if these features are used, then it is essential that the mirror functionality is disabled (the mirror is broken) before VMs are restarted on the recovery site.

Failover, Failback and Test Failover with the Disaster Recovery wizard

The **Disaster Recovery** wizard makes failover (recovery of protected VMs and vApps to a secondary site) and failback (restoration of VMs and vApps back to the primary site) simple. The steps involved in the two processes are outlined here:

Failover

1. First, you choose a target pool on your secondary DR site to which you want to recover your VMs and vApps.

2. Next, you provide details of the storage targets containing the replicated SRs from your primary site.

3. The wizard scans the targets and lists all SRs found there.

   Now you select the SRs containing the metadata and virtual disks for the VMs and vApps you want to recover.

4. The wizard scans the SRs and lists all the VMs and vApps found.

   Now you select which VMs and vApps you want to recover to the DR site, and specify whether you want the wizard to start them up automatically as soon as they have been recovered, or whether you prefer to wait and start them up manually yourself.

5. The wizard then performs a number of prechecks to ensure that the selected VMs and vApps can be recovered to the target DR pool, for example, it checks that all the storage required by the selected VMs and vApps is available.

6. Finally, when the prechecks are complete and any issues resolved, the failover process begins. The selected VMs and vApps are exported from the replicated storage to the DR pool.

   Failover is now complete.

Failback

1. First, you choose the target pool on your primary site to which you want to restore the VMs and vApps currently running on the DR site.

2. Next, you provide details of the storage targets containing the replicated SRs from your DR site.

3. The wizard scans the targets and lists all SRs found.

   Now you select the SRs containing the metadata and virtual disks for the VMs and vApps you want to restore.

4. The wizard scans the SRs and lists all the VMs and vApps found.

   Now you select which VMs and vApps you want to restore back to the primary site and specify whether you want the wizard to start them up automatically as soon as they have been recovered, or whether you prefer to wait and start them up manually yourself.

5. The wizard then performs a number of prechecks to ensure that the selected VMs and vApps can be recovered to the target pool on the primary site, for example, it checks that all the storage required by the selected VMs and vApps is available.

6. Finally, when the prechecks are complete and any issues resolved, the failback process begins. The selected VMs and vApps running on your DR site are exported from the replicated storage back to the selected pool at your primary site.

   Failback is now complete.

If the **Disaster Recovery** wizard finds information for the same VM present in a two or more places (for example, storage from the primary site, storage from the DR site and also in the pool that the data in to be imported into) then it will ensure that only the most recent information per VM is used.

> **Tip:**
>
> Recovering VMs and vApps from replicated storage will be easier if your SRs are named in a way that captures how your VMs and vApps are mapped to SRs, and the SRs to LUNs.

You can also use the **Disaster Recovery** wizard to run test failovers for non-disruptive testing of your disaster recovery system. In a test failover, all the steps are the same as for failover, but the VMs and vApps are started up in a paused state after they have been recovered to the DR site, and cleanup is performed when the test is finished to remove all VMs, vApps and storage recreated on the DR site. See Test Failover.

XenServer DR terminology

vApp: A logical group of related VMs which are managed as a single entity.

Site: A physical group of XenServer resource pools, storage and hardware equipment.

Primary site: A physical site that runs VMs or vApps which must be protected in the event of disaster.

Secondary site, DR site: A physical site whose purpose is to serve as the recovery location for the primary site, in the event of a disaster.

Failover: Recovery of VMs and vApps on a secondary (recovery) site in the event of disaster at the primary site.

Failback: Restoration of VMs and vApps back to the primary site from a secondary (recovery) site.

Test failover: A "dry run" failover where VMs and vApps are recovered from replicated storage to a pool on a secondary (recovery) site but not actually started up. Test failovers can be run to check that DR is correctly configured and that your processes are effective.

Pool metadata: Information about the VMs and vApps in the pool, such as their name and description, and, for VMs, configuration information including UUID, memory, virtual CPU, networking and storage configuration, and startup options - start order, delay interval and HA restart priority. Pool metadata is used in DR to re-create the VMs and vApps from the primary site in a recovery pool on the secondary site.

Configuring DR

September 30, 2020

Use the XenCenter **Configure DR** dialog box to identify storage repositories (SRs) where the metadata for a pool - configuration information about all the VMs and vApps in the pool - will be stored. This metadata will be updated whenever you make changes to VM or vApp configuration within the pool.

To configure DR on the primary pool:

1. On your primary site, select the pool that you want to protect and on the **Pool** menu, point to **Disaster Recovery** and then click **Configure**.

2. Select up to 8 SRs where the pool metadata will be stored. A small amount of space will be required on this storage for a new LUN which will contain the pool recovery information.

   > **Note:**
   >
   > XenServer DR can only be used with LVM over HBA or LVM over iSCSI storage types.

3. Click **OK**.

   > **Important:**
   >
   > Your XenServer DR environment will not be fully configured until the SRs containing the pool metadata and the virtual disks used by the VMs in the pool are replicated from your production environment to a backup environment. Storage replication, for example using

> mirroring, is best handled by your storage solution, and will vary from device to device; it cannot be configured from within XenCenter.

Failover

September 30, 2020

Failover is the process of recovering VMs and vApps to a secondary (recovery) site in the event of disaster at your primary site. To failover your critical VMs and vApps, use the **Disaster Recovery** wizard.

> **Important:**
>
> The **Disaster Recovery** wizard does not control any storage array functionality. Duplication (mirroring) of the metadata storage and the storage used by the VMs which are to be restarted must be turned off before you attempt failover to your recovery site.

To fail over VMs and vApps to a secondary site:

1. In XenCenter, select the secondary pool, and on the **Pool** menu, click **Disaster Recovery** to open the **Disaster Recovery** wizard.

2. Select **Failover** and then click **Next**.

   > **Note:**
   >
   > If you use Fibre Channel shared storage with LUN mirroring to replicate the data to the secondary site, before you attempt to recover data, mirroring must be broken so that the secondary site has Read/Write access.

3. Select the storage repositories (SRs) containing the pool metadata for the VMs and vApps that you want to recover. By default, the list on this wizard page shows all SRs that are currently attached within the pool. To scan for more SRs, choose **Find Storage Repositories** and then select the storage type to scan for:

   * To scan for all the available Hardware HBA SRs, select **Find Hardware HBA SRs**.
   * To scan for software iSCSI SRs, select **Find Software iSCSI SRs** and then enter the target host, IQN and LUN details in the dialog box.

   When you have selected the required SRs in the wizard, click **Next** to continue.

4. Select the VMs and vApps that you wish to recover and choose the appropriate **Power state after recovery** option to specify whether you want the wizard to start them up automatically as soon as they have been recovered, or whether you prefer to wait and start them up manually yourself after failover is complete. Click **Next** to progress to the next wizard page and begin failover prechecks.

5. The wizard performs a number of pre-checks before starting failover, for example, to ensure that all the storage required by the selected VMs and vApps is available. If any storage is missing at this point, you can click **Attach SR** on this page to find and attach the relevant SR. Resolve any issues on the prechecks page, and then click **Failover** to begin the recovery process.

6. A progress page is displayed showing whether recovery was successful for each VM and vApp. Failover may take some time depending on the number of VMs and vApps you are recovering, as the metadata for the VMs and vApps are exported from the replicated storage, the VMs and vApps are re-created in the primary pool, the SRs containing the virtual disks are attached to the re-created VMs, and VMs are started, if specified.

7. When the failover is complete, click **Next** to see the summary report. Click **Finish** on the summary report page to close the wizard.

   Once your primary site is available again and you want to return to running your VMs on that site, you will need to work through the **Disaster Recovery** wizard again, but this time selecting the **Failback** option on the first page; see Failback.

Failback

September 30, 2020

Failback is the process of restoring VMs and vApps from replicated storage back to a pool on your primary (production) site when the primary site comes back up after a disaster event. To fail back VMs and vApps to your primary site, use the **Disaster Recovery** wizard.

> **Important:**
>
> The **Disaster Recovery** wizard does not control any storage array functionality. Duplication (mirroring) of the metadata storage and the storage used by the VMs which are to be restored must be turned off before you attempt failback to your primary site.

To fail back VMs and vApps to your primary site:

1. In XenCenter, select the target pool on your primary site, and on the **Pool** menu, click **Disaster Recovery** to open the **Disaster Recovery** wizard.

2. Select **Failback** and then click **Next**.

   > **Note:**
   >
   > If you use Fibre Channel shared storage with LUN mirroring to replicate the data on the DR site, before you attempt to recover data, mirroring must be broken so that the primary site has Read/Write access.

3. Select the storage repositories (SRs) containing the pool metadata for the VMs and vApps that you want to restore back to your primary site. By default, the list on this wizard page shows all SRs that are currently attached within the pool. To scan for more SRs, choose **Find Storage Repositories** and then select the storage type to scan for:

   - To scan for all the available Hardware HBA SRs, select **Find Hardware HBA SRs**.
   - To scan for software iSCSI SRs, select **Find Software iSCSI SRs** and then enter the target host, IQN and LUN details in the dialog box.

   When you have selected the required SRs in the wizard, click **Next** to continue.

4. Select the VMs and vApps that you wish to restore and choose the appropriate **Power state after recovery** option to specify whether you want the wizard to start them up automatically as soon as they have been restored, or whether you prefer to wait and start them up manually yourself after failback is complete. Click **Next** to progress to the next wizard page and begin failback prechecks.

5. The wizard performs a number of pre-checks before starting failback, for example, to ensure that all the storage required by the selected VMs and vApps is available. If any storage is missing at this point, you can click **Attach SR** on this page to find and attach the relevant SR. Resolve any issues on the prechecks page, and then click **Failback** to begin the recovery process.

6. A progress page is displayed showing whether restoration was successful for each VM and vApp. Failback may take some time depending on the number of VMs and vApps you are restoring.

7. When the failback is complete, click **Next** to see the summary report. Click **Finish** on the summary report page to close the wizard.

Test Failover

September 30, 2020

Failover testing is an essential component in disaster recovery planning. You can use the **Disaster Recovery** wizard to perform non-disruptive testing of your disaster recovery system. During a test failover operation, all the steps are the same as for failover, but instead of be started after they have been recovered to the DR site, the VMs and vApps are placed in a paused state. At the end of a test failover operation, all VMs, vApps and storage recreated on the DR site are automatically removed.

After initial DR configuration, and after you make significant configuration changes in a DR-enabled pool, we recommend that you verify that failover still works correctly by performing a test failover.

To perform a test failover of VMs and vApps to a secondary site:

1. In XenCenter, select the secondary pool, and on the **Pool** menu, click **Disaster Recovery** to open the **Disaster Recovery** wizard.

2. Select **Test Failover** and then click **Next**.

> **Note:**
>
> If you use Fibre Channel shared storage with LUN mirroring to replicate the data to the secondary site, before you attempt to recover data, mirroring must be broken so that the secondary site has Read/Write access.

3. Select the storage repositories (SRs) containing the pool metadata for the VMs and vApps that you want to recover. By default, the list on this wizard page shows all SRs that are currently attached within the pool. To scan for more SRs, choose **Find Storage Repositories** and then select the storage type to scan for:

   - To scan for all the available Hardware HBA SRs, select **Find Hardware HBA SRs**.
   - To scan for software iSCSI SRs, select **Find Software iSCSI SRs** and then enter the target host, IQN and LUN details in the dialog box.

   When you have selected the required SRs in the wizard, click **Next** to continue.

4. Select the VMs and vApps that you wish to recover then click **Next** to progress to the next wizard page and begin failover prechecks.

5. Before beginning the test failover process, the wizard performs a number of pre-checks, for example, to ensure that all the storage required by the selected VMs and vApps is available.

   a) **Check that storage is available.** If any storage is missing, you can click **Attach SR** on this page to find and attach the relevant SR.

   b) **Check that HA is not enabled on the target DR pool.** To avoid having the same VMs running on both the primary and DR pools, HA must be disabled on the secondary pool to ensure that the recovered VMs and vApps are not started up automatically by HA after recovery. To disable HA on the secondary pool, you can simply click **Disable HA** on the this page. (If HA is disabled at this point, it will be enabled again automatically at the end of the test failover process.)

   Resolve any issues on the pre-checks page, and then click **Failover** to begin the test failover.

6. A progress page is displayed showing whether recovery was successful for each VM and vApp. Failover may take some time depending on the number of VMs and vApps you are recovering, as the metadata for the VMs and vApps are recovered from the replicated storage, the VMs and vApps are re-created in the DR pool, the SRs containing the virtual disks are attached to the re-created VMs. **The recovered VMs are placed in a paused state: they will not be started up on the secondary site during a test failover.**

7. After you are satisfied that the test failover was performed successfully, click **Next** in the wizard to have the wizard clean up on the DR site:

   - VMs and vApps that were recovered during the test failover will be removed.

- Storage that was recovered during the test failover will be detached.
- If HA on the DR pool was disabled at the prechecks stage to allow the test failover to take place, it will be enabled again automatically.

The progress of the cleanup process is displayed in the wizard.

8. Click **Finish** to close the wizard.

Access Control (AD & RBAC)

June 3, 2019

Topics

- Managing Users
- RBAC Overview
- Definitions of RBAC Roles and Permissions
- Join a Domain and Add Users
- Assign Roles to Users and Groups
- Calculating RBAC Roles
- Audit Changes to XenServer

Managing Users

April 9, 2020

When you first install XenServer, a user account is added to XenServer automatically. This account is the local super user (LSU) , or root , which is authenticated locally by the XenServer computer. You can create additional users by adding Active Directory accounts from the **Users** tab in XenCenter. (Note that the term "user" refers to anybody with a XenServer account, that is, anyone administering XenServer hosts, regardless of level of their role.) If you want to have multiple user accounts on a server or a pool, you must use Active Directory user accounts for authentication. This allows XenServer users to log in to the servers in a pool using their Windows domain credentials.

> **Note:**
>
> Mixed-authentication pools are not supported (that is, you cannot have a pool where some servers in the pool are configured to use Active Directory and some are not).

When you create a new user in XenServer you must first assign a role to the newly created user before they can use the account. Note that XenServer **does not** automatically assign a role to the newly

created user. As a result, these accounts will not have any access to the XenServer pool until you assign them a role.

Using the Role Based Access Control (RBAC) feature, you can assign the Active Directory accounts different levels of permissions depending on the user's role. If you do not use Active Directory in your environment, you are limited to the LSU account.

AD authentication in a XenServer environment

Even though XenServer is Linux-based, XenServer lets you use Active Directory accounts for XenServer user accounts. To do so, it passes Active Directory credentials to the Active Directory domain controller.

> **Note:**
>
> You can enable LDAP channel binding and LDAP signing on your AD domain controllers. For more information, see Microsoft Security Advisory.

When added to XenServer, Active Directory users and groups become XenServer subjects , generally referred to as simply users in XenCenter. When a subject is registered with XenServer, users/groups are authenticated with Active Directory on login and do not need to qualify their user name with a domain name.

> **Note:**
>
> By default, if you did not qualify the user name (for example, enter either mydomain\myuser or myuser@mydomain.com), XenCenter always attempts to log users in to Active Directory authentication servers using the domain to which it is currently joined. The exception to this is the LSU account, which XenCenter always authenticates locally (that is, on the XenServer) first.

The external authentication process works as follows:

1. The credentials supplied when connecting to a server are passed to the Active Directory domain controller for authentication.
2. The domain controller checks the credentials. If they are invalid, the authentication fails immediately.
3. If the credentials are valid, the Active Directory controller is queried to get the subject identifier and group membership associated with the credentials.
4. If the subject identifier matches the one stored in the XenServer, the authentication is completed successfully.

When you join a domain, you enable Active Directory authentication for the pool. However, when a pool is joined to a domain, only users in that domain (or a domain with which it has trust relationships) can connect to the pool.

RBAC overview

September 30, 2020

The Role Based Access Control (RBAC) feature lets you assign predefined roles , or sets of XenServer permissions, to Active Directory users and groups. These permissions control the level of access XenServer users (that is, people administering XenServer) have to servers and pools: RBAC is configured and deployed at the pool level. Because users acquire permissions through their assigned role, you simply need to assign a role to a user or their group.

Using Active Directory accounts for XenServer user accounts

RBAC lets you restrict which operations different groups of users can perform, which reduces the likelihood of inexperienced users making disastrous, accidental changes. Assigning RBAC roles also helps prevent unauthorized changes to your resource pools for compliance reasons. To facilitate compliance and auditing, RBAC also provides an Audit Log feature and its corresponding Workload Balancing Pool Audit Trail report.
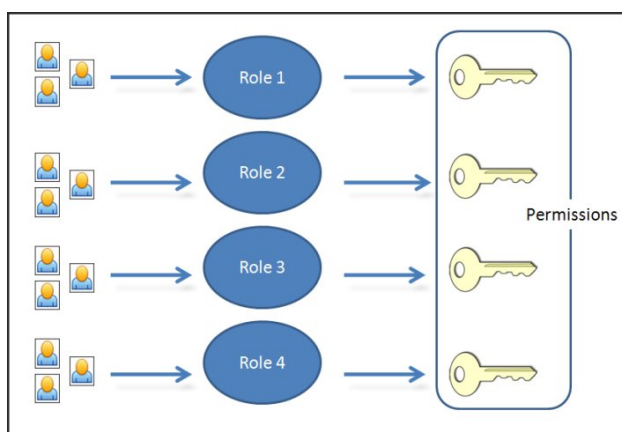
RBAC depends on Active Directory for authentication services. Specifically, XenServer keeps a list of authorized users based on Active Directory user and group accounts. As a result, you must join the pool to the domain and add Active Directory accounts before you can assign roles.

RBAC process

This is the standard process for implementing RBAC and assigning a user or group a role:

1. Join the domain.
2. Add an Active Directory user or group to the pool.
3. Assign (or modify) the user or group's RBAC role.

Local Super User

The local super user (LSU) , or root, is a special user account used for system administration and has all rights or permissions. In XenServer, the local super user is the default account at installation. The LSU is authenticated by XenServer and not an external authentication service. This means that if the external authentication service fails, the LSU can still log in and manage the system. The LSU can always access the XenServer physical server through SSH.

RBAC Roles

XenServer comes with six pre-established roles that are designed to align with different functions in an IT organization.

- **Pool Administrator (Pool Admin)**. This role is the most powerful role available. Pool Admins have full access to all XenServer features and settings. They can perform all operations, including role and user management. They can grant access to the XenServer console. As a best practice, Citrix recommends assigning this role to an extremely limited number of users.

  > **Note:**
  >
  > The local super user (root) always has the Pool Admin role. The Pool Admin role has the same permissions as the local root.

- **Pool Operator (Pool Operator)**. This role is designed to let the assignee manage pool-wide resources, including creating storage, managing servers, managing patches, and creating pools. Pool Operators can configure pool resources. They also have full access to the following features: High Availability (HA), Workload Balancing, and patch management. Pool Operators cannot add users or modify roles.

- **Virtual Machine Power Administrator (VM Power Admin)**. This role has full access to VM and Template management. They can choose where to start VMs. They have full access to the dynamic memory control features and the VM snapshot feature. In addition, they can set the Home Server and choose where to run workloads. Assigning this role grants the assignee sufficient permissions to provision virtual machines for VM Operator use.

- **Virtual Machine Administrator (VM Admin)**. This role can manage VMs and Templates and access the storage necessary to complete these tasks. However, this role relies on XenServer to choose where to run workloads and must use the settings in templates for dynamic memory control and the Home Server. (This role cannot access the dynamic memory control features, make snapshots, set the Home Server or choose where to run workloads.)

- **Virtual Machine Operator (VM Operator)**. This role can use the VMs in a pool and manage their basic lifecycle. VM Operators can interact with the VM consoles and start or stop VMs, provided sufficient hardware resources are available. Likewise, VM Operators can perform start and stop

lifecycle operations. The VM Operator role cannot create or destroy VMs, alter VM properties, or server resources.

- **Read-only (Read Only)**. This role can only view resource pool and performance data.

For information about the permissions associated with each role, see Definitions of RBAC roles and permissions. For information about how RBAC calculates which roles apply to a user, see Calculating RBAC roles.

> **Note:**
>
> When you create a new user, you must first assign a role to the newly created user before they can use the account. Note that XenServer **does not** automatically assign a role to the newly created user.

Definitions of RBAC roles and permissions

June 15, 2020

Permissions available for each role

The following table summarizes which permissions are available for each role. For details on the operations available for each permission, see the next section.

Permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Assign/modify roles	X					
Log in to (physical) server consoles (through SSH and XenCenter)	X					
Server backup/restore	X					

Permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Import/export OVF/OVA packages; import disk images	X					
Set cores per socket	X					
Convert VMs using XenServer Conversion Manager	X					
Switch-port locking	X	X				
Log out active user connec-tions	X	X				
Create and dismiss alerts	X	X				
Cancel task of any user	X	X				
Pool man-agement	X	X				
Storage XenMotion	X	X	X			
VM advanced operations	X	X	X			
VM cre-ate/destroy operations	X	X	X	X		

Permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
VM change CD media	X	X	X	X	X	
VM change power state	X	X	X	X	X	
View VM consoles	X	X	X	X	X	
XenCenter view management operations	X	X	X	X	X	
Cancel own tasks	X	X	X	X	X	X
Read audit logs	X	X	X	X	X	X
Configure, initialize, enable, disable WLB	X	X				
Apply WLB optimization recommendations	X	X				
Modify WLB report subscriptions	X	X				
Accept WLB placement recommendations	X	X	X			
Display WLB configuration	X	X	X	X	X	X

Permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Generate WLB reports	X	X	X	X	X	X
Connect to pool and read all pool metadata	X	X	X	X	X	X
Configure virtual GPU	X	X				
View virtual GPU configuration	X	X	X	X	X	X
Access the config drive (CoreOS VMs only)	X					
Container management	X					
Configure Health Check	X	X				
View Health Check results and settings	X	X	X	X	X	X
vCPU Hotplug	X	X	X	X		
Configure PVS-Accelerator	X	X				

Permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
View PVS-Accelerator configuration	X	X	X	X	X	X

Definitions of permissions

This section provides additional details about permissions:

Assign/modify roles

- Add and remove users
- Add and remove roles from users
- Enable and disable Active Directory integration (being joined to the domain)

This permission lets the user grant themself any permission or perform any task.

Warning: This role lets the user disable the Active Directory integration and all subjects added from Active Directory.

Log in to server consoles

- Server console access through ssh
- Server console access through XenCenter

Warning: With access to a root shell, the assignee could arbitrarily reconfigure the entire system, including RBAC.

Server backup/restore; VM create/destroy operations

- Back up and restore servers
- Back up and restore pool metadata

The ability to restore a backup lets the assignee revert RBAC configuration changes.

Import/export OVF/OVA packages; import disk images

- Import OVF and OVA packages
- Import disk images
- Export VMs as OVF/OVA packages

Set cores-per-socket

- Set the number of cores per socket for the VM's virtual CPUs

This permission enables the user to specify the topology for the VM's virtual CPUs.

Convert VMs using XenServer Conversion Manager

- Convert VMware VMs to XenServer VMs

This permission lets the user convert workloads from VMware to XenServer by copying batches of VMware VMs to XenServer environment.

Switch-port locking

- Control traffic on a network

This permission lets the user to block all traffic on a network by default, or define specific IP addresses from which a VM is allowed to send traffic.

Log out active user connections

- Ability to disconnect logged in users

Create/dismiss alerts

- Configure XenCenter to generate alerts when resource usage crosses certain thresholds
- Remove alerts from the Alerts view

Warning: A user with this permission can dismiss alerts for the entire pool.

Note: The ability to view alerts is part of the **Connect to Pool and read all pool metadata permission**.

Cancel task of any user

- Cancel any user's running task

This permission lets the user request XenServer cancel an in-progress task initiated by any user.

Pool management

- Set pool properties (naming, default SRs)
- Enable, disable, and configure HA
- Set per-VM HA restart priorities
- Configure DR and perform DR failover, failback and test failover operations.

- Enable, disable, and configure Workload Balancing (WLB)
- Add and remove server from pool
- Emergency transition to master
- Emergency master address
- Emergency recover pool members
- Designate new master
- Manage pool and server certificates
- Patching
- Set server properties
- Configure server logging
- Enable and disable servers
- Shut down, reboot, and power-on servers
- Restart toolstack
- System status reports
- Apply license
- Live migration of all other VMs on a server to another server, due to either WLB, Maintenance Mode, or HA
- Configure server management interfaces
- Disable server management
- Delete crashdumps
- Add, edit, and remove networks
- Add, edit, and remove PBDs/PIFs/VLANs/Bonds/SRs

Storage XenMotion

- Migrate VMs from one host to another host
- Move Virtual Disk (VDIs) from one SR to another SR

VM advanced operations

- Adjust VM memory (through Dynamic Memory Control)
- Create a VM snapshot with memory, take VM snapshots, and roll-back VMs
- Migrate VMs
- Start VMs, including specifying physical server
- Resume VMs

Log in to server consoles

VM create/destroy operations

- Install and delete VMs

- Clone/copy VMs
- Add, remove, and configure virtual disk/CD devices
- Add, remove, and configure virtual network devices
- Import/export XVA files
- VM configuration change

VM change CD media

- Eject current CD
- Insert new CD

VM change power state

- Start VMs (automatic placement)
- Shut down VMs
- Reboot VMs
- Suspend VMs
- Resume VMs (automatic placement)

Log out active user connections

View VM consoles

- See and interact with VM consoles

Create/dismiss alerts

Configure, Initialize, Enable, Disable WLB

- Configure WLB
- Initialize WLB and change WLB servers
- Enable WLB
- Disable WLB

Apply WLB Optimization Recommendations

- Apply any optimization recommendations that appear in the WLB tab

Modify WLB Report Subscriptions

- Change the WLB report generated or its recipient

Accept WLB Placement Recommendations

- Select one of the servers Workload Balancing recommends for placement ("star" recommendations)

Display WLB Configuration

- View WLB settings for a pool as shown on the WLB tab

Generate WLB Reports

- View and run WLB reports, including the Pool Audit Trail report

XenCenter view management operations

- Create and modify global XenCenter folders
- Create and modify global XenCenter custom fields
- Create and modify global XenCenter searches

Cancel own tasks

- Enables users to cancel their own tasks

Read audit log

- Download the XenServer audit log

Connect to pool and read all pool metadata

- Log in to pool
- View pool metadata
- View historical performance data
- View logged in users
- View users and roles
- View tasks
- View messages
- Register for and receive events

Configure virtual GPU

- Specify a pool-wide placement policy
- Assign a virtual GPU to a VM
- Remove a virtual GPU from a VM
- Modify allowed virtual GPU types
- Create, destroy, or assign a GPU group

View virtual GPU configuration

- View GPUs, GPU placement policies, and virtual GPU assignments.

Access the config drive (CoreOS VMs only)

- Access the config driver of the VM
- Modify the cloud-config parameters

Container management

- Start
- Stop
- Pause
- Resume
- access information about the container

Configure Health Check

- Enable Health Check
- Disable Health Check
- Update Health Check settings
- Manually upload a server status report

View Health Check results and settings

- View the results of a Health Check upload
- View Health Check enrollment settings

Configure PVS-Accelerator

- Enable PVS-Accelerator
- Disable PVS-Accelerator

- Update PVS-Accelerator cache configuration
- Add or Remove PVS-Accelerator cache configuration

View PVS-Accelerator configuration

- View the status of PVS-Accelerator

Join a domain and add users

September 30, 2020

Before you can assign a user or group account an RBAC role, you must add the account to XenServer through RBAC. This requires two tasks:

1. Join the pool or server to the domain. The domain can be either the domain in which the user or group belongs or a domain that is in the same Active Directory forest or that has a trust relationship with the user's domain.
2. Add the user's Active Directory account or group to XenServer.

After adding the user's Active Directory account or group to XenServer, in both the free XenServer product and in XenServer Advanced, the user is automatically assigned a fixed role of Pool Admin. In XenServer Enterprise and higher, you must assign a role to the user or group manually.

To change domains, leave the current domain and then join the new domain.

To join the XenServer or pool to a domain

1. In the **Resources Pane** , select the pool or server for which you want to grant somebody permissions.
2. Click the **Users** tab.
3. Click **Join Domain**.
4. Enter Active Directory credentials with sufficient privileges to add servers to the domain you want to join. The domain to be joined must be specified as a fully qualified domain name (FQDN) rather than a NetBIOS name. For example, enter your_domain.net instead of your_domain .

To add an Active Directory user or group to a pool

1. After joining the user's domain, in the **Users** tab, click **Add**.
2. In the **Add Users** dialog box, enter one or more user or group names. Separate multiple names by commas. To specify a user in a different, trusted domain (other than the one currently

joined), supply the domain name with the user name (for example, other_domain\jsmith) or enter a fully qualified domain name (FQDN) (for example, jsmith@other_domain.com).

3. Click **Grant Access**.
4. Follow Assign roles to users and groups to assign the account a role and grant access.

To leave the domain

> **Note:**
>
> When you leave the domain (that is, disable Active Directory authentication and disconnect a pool or server from its domain), any users who authenticated to the pool or server with Active Directory credentials are disconnected.

1. In the **Resources Pane** , select the pool or server that you want to disconnect from its Active Directory domain.
2. Click **Leave Domain** and select **Yes** to continue.
3. Enter Active Directory credentials with sufficient privileges to disable servers in the domain you want to leave.
4. Decide whether to disable the computer accounts in the Active Directory server, and then click one of the following:
   - **Disable**. Removes the pool or server from the domain and disables the computer account for the server or pool master in the Active Directory database.
   - **Ignore**. Select this option if you did not fill the username/password or you do not know an account with sufficient privileges to remove the server or pool master's computer account from the Active Directory database. (This option removes the pool or server from the domain, but leaves the computer account for the server or pool master in the Active Directory.)

Assign roles to users and groups

September 30, 2020

All XenServer users must have an RBAC role. In XenServer, you must first assign a role to the newly created user before they can use the account. Note that XenServer does not automatically assign a role to the newly created user. As a result, these accounts will not have any access to the XenServer pool until you assign them a role.

> **Note:**
>
> Before you can assign a role to a user or group, you must add the user or group's Active Directory account to XenServer after joining the associated domain as described in Join a domain and add

> RBAC users.

You can assign a user a different role by one of the following methods:

1. Change the role assigned to the user in the **Select Roles** dialog in XenCenter. This requires the **Assign/modify role** permission, which is only available to a Pool Administrator.
2. Modify the user's containing group membership in Active Directory (so that the user becomes part of a group that is assigned a different role).

If, on separate occasions, an administrator indirectly applies multiple roles to a user (for example, by the user being a member of multiple groups), XenServer grants the user the permissions from the highest role to which he or she was assigned.

To change or assign a role to a user or group

1. In the **Resources** pane, select the pool or server that contains the user or group.

2. Click the **Users** tab.

3. In the **Users and Groups with Access** pane, select the user or group to which you want to assign permissions.

4. Click **Change Role**.

5. In the **Select Roles** dialog, select the role you want to apply and click **Save**. For information about the permissions associated with each role, see Definitions of RBAC roles and permissions.

   > **Tip:**
   >
   > When you are assigning a role, you can select multiple users simultaneously by pressing the CTRL key and selecting the user accounts.

6. (Optional.) When changing a role, if the user is currently logged on to the pool and you want them to receive their new permissions immediately, click **Logout User**. This disconnects all of the user's sessions on the pool so the user receives a new session with the modified role.

   > **Note:**
   >
   > When changing a role, for the new role to take effect, the user must log out and log back in again. You can force this log out by clicking the **Logout User** button. (Forcing logouts requires the **Logout active user connections** permission, which is available to a Pool Administrator or Pool Operator).

Assigning roles to users in XenServer 6.1.0 and earlier

In XenServer v6.1.0 and earlier, when new users are added, they are automatically assigned the Pool Administrator role. In XenServer Enterprise and higher, when you add new users, XenServer does

not assign newly added user accounts roles automatically. You must assign roles to new accounts separately.

Calculating RBAC roles

June 20, 2019

When I log in, how does XenServer compute the roles for the session?

1. The Active Directory server authenticates the subject. During authentication, Active Directory also determines if the subject belongs to any other containing groups in Active Directory.
2. XenServer then verifies which roles have been assigned to (a) the subject and (b) to any Active Directory groups to which it is a member.
3. XenServer applies the highest level of permissions to the subject. Because subjects can be members of multiple Active Directory groups, they will inherit all of the permissions of the associated roles.
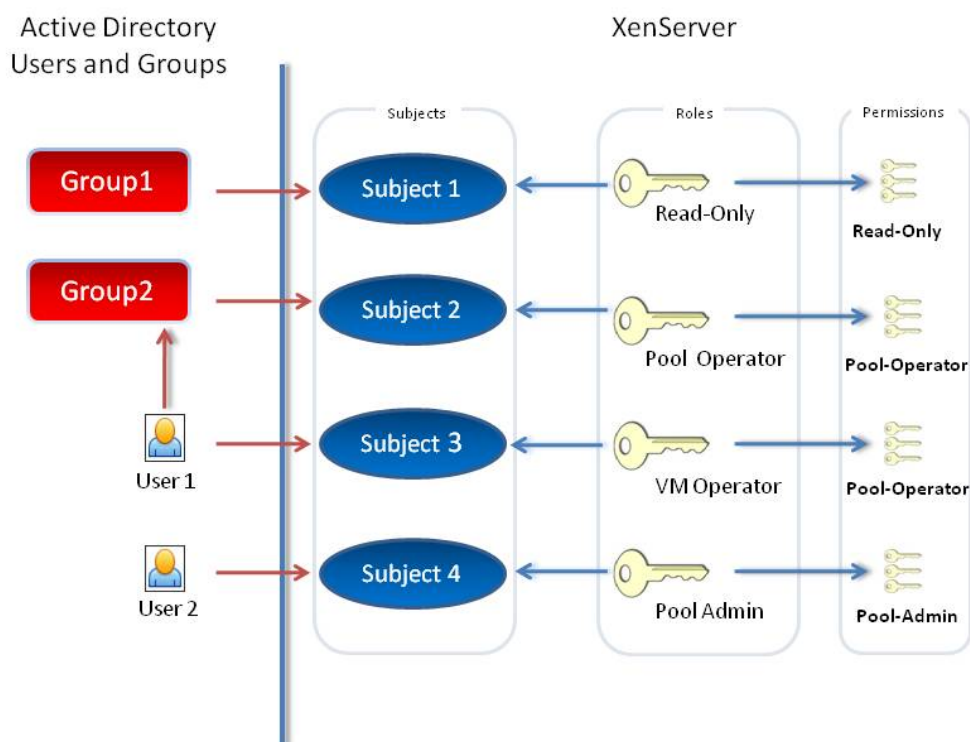
In this illustration, since Subject 2 (Group 2) is the Pool Operator and User 1 is a member of Group 2, when Subject 3 (User 1) tries to log in, he or she inherits both Subject 3 (VM Operator) and Group 2 (Pool Operator) roles. Since the Pool Operator role is higher, the resulting role for Subject 3 (User 1) is Pool Operator and not VM Operator.

Audit changes

June 3, 2019

The XenServer audit log, which is enabled in XenServer by default, records any operation with side-effects (successful or unsuccessful) performed by a known user. This includes:

- The user's name who performed the action. If the user's name is not available, XenServer logs the user ID instead.
- The server name that was targeted by the action.
- The status of the action - if it was successful or unsuccessful and if it was authorized. if the operation failed then the error code is logged.

The audit logging feature is enabled by default. The audit log can be backed up by using the XenServer syslog command to duplicate the audit log to a safe box. The syslog command is available from the CLI and documented in *XenServer Administrator's Guide*.

While Citrix strongly recommends that customers concerned with auditing implement Role Based Access Control, the audit log itself does not require that users be assigned RBAC roles nor does it require Active Directory integration.

XenServer logs actions on the pool level, and creates a log for each pool on the pool master.

To display the audit log, you have two choices. You can:

- Generate the Pool Audit Trail report, provided you have Workload Balancing enabled.
- Display the audit log by opening it in any text editor. The log is stored on the pool master.

Workload Balancing Overview

September 30, 2020

Workload Balancing is a XenServer appliance that balances your pool by relocating virtual machines onto the best possible servers for their workload in a resource pool. For example, Workload Balancing can:

- Balance virtual-machine workloads across hosts in a XenServer resource pool

- Determine the best host on which to start a virtual machine
- Determine the best host on which to power on a virtual machine that you powered off
- Determine the best host for each of the host's virtual machines when you put that host into Maintenance mode

Note:

Workload Balancing is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. For more information about licensing, see About XenServer Licensing.

Depending on your preference, Workload Balancing can accomplish these tasks automatically or prompt you to accept its optimization, consolidation, and placement recommendations. You can also configure Workload Balancing to power off hosts automatically during periods of low usage (for example, to save power at night).

Workload Balancing evaluates the utilization of virtual-machine workloads across a pool and, when a host reaches one of its thresholds, relocates the virtual machine to a different host in the pool.

To ensure the rebalancing and placement recommendations align with your environment's needs, you can configure Workload Balancing to optimize workloads for resource performance or to maximize density. These optimization modes can be configured to change automatically at predefined times or stay the same at all times. For additional granularity, you can fine-tune the weighting of individual resource metrics (CPU, network, memory, and disk).

To help you perform capacity planning, Workload Balancing provides historical reports about host and pool health, optimization and virtual-machine performance, and virtual-machine motion history.

Getting Started with Workload Balancing

September 30, 2020

You can download the Workload Balancing virtual appliance and get it up and running using the following process:

1. Download the Workload Balancing virtual appliance package from **My Account** at www.citrix.com.

2. In XenCenter, click **File** and then **Import** and follow the on-screen instructions to import the Workload Balancing virtual appliance.

3. Configure the Workload Balancing virtual appliance using the Workload Balancing Configuration wizard, which appears in the appliance's **Console** tab in XenCenter after you import the virtual appliance.

4. Connect your pool to the Workload Balancing virtual appliance as described in Connecting to Workload Balancing.

More information about these steps is in the *Workload Balancing Quick Start*.

Note:

The WLB tab appears on the Properties pane after licensing your hosts with XenServer Enterprise or a Citrix Virtual Apps and Desktops license. For more information about licensing, see About XenServer Licensing.

Workload Balancing Basic Concepts

June 3, 2019

Workload Balancing captures data for resource performance on virtual machines and physical hosts. It uses this data, combined with the preferences you set, to provide optimization and placement recommendations. Workload Balancing stores performance data in an internal database: the longer Workload Balancing runs, the better its recommendations become.

Workload Balancing recommends moving virtual-machine workloads across a pool to get the maximum efficiency, which means either performance or density depending on your goals. Within a Workload Balancing context:

- **Performance** refers to the usage of physical resources on a host (for example, the CPU, memory, network, and disk utilization on a host). When you set Workload Balancing to maximize performance, it recommends placing virtual machines to ensure the maximum amount of resources are available for each virtual machine.
- **Density** refers to the number of virtual machines on a host. When you set Workload Balancing to maximize density, it recommends placing virtual machines on as few hosts as possible, while ensuring they maintain adequate computing power, so you can reduce the number of hosts powered on in a pool.

Workload Balancing lets you modify settings for placement (performance or density), power management, automation, metric weightings, and performance thresholds.

Workload Balancing will not conflict with High Availability settings; High Availability settings always take precedence.

Connecting to Workload Balancing

September 30, 2020

After importing the Workload Balancing virtual appliance and running the Workload Balancing Configuration wizard, you must connect the pool you want monitored to Workload Balancing. To do so, use either the Connect to WLB Server dialog box in XenCenter or the xe commands.

> **Note:**
>
> The WLB tab appears on the Properties pane after licensing your hosts with XenServer Enterprise or a Citrix Virtual Apps and Desktops license. For more information about licensing, see About XenServer Licensing.

Prerequisites

To complete the XenCenter procedure that follows, you need the:

- IP address or Fully Qualified Domain Name (FQDN) and port of the Workload Balancing virtual appliance.
- Credentials for the resource pool (that is, the pool master) you want Workload Balancing to monitor.
- Credentials for the account you created on the Workload Balancing appliance. This is often known as the Workload Balancing user account. XenServer uses this account to communicate with Workload Balancing. (You created this account on the Workload Balancing virtual appliance during Workload Balancing Configuration.)

If you want to specify the Workload Balancing virtual appliance's FQDN when connecting to the Workload Balancing server (that is, in the Connect to WLB Server dialog), you must first manually add its host name to your DNS. If you want to configure Trusted Authority certificates, Citrix recommends specifying either an FQDN or an IP address that does not expire.

When you first connect to Workload Balancing, it uses the default thresholds and settings for balancing workloads. Automatic features, such as Automated Optimization Mode, Power Management, and Automation, are disabled by default.

> **Important:**
>
> After Workload Balancing is running for a period of time, if you do not receive optimal placement recommendations, Citrix strongly recommends you evaluate your performance thresholds as described in Evaluating the Effectiveness of Your Optimization Thresholds. It is critical to set the correct thresholds for your environment for Workload Balancing recommendations to be optimal.

To connect to the Workload Balancing virtual appliance

1. In the **Resources** pane of XenCenter, select **XenCenter** > **your-resource-pool**.

2. In the **Properties** pane, click on the **WLB** tab. Note that the WLB tab appears in the Properties pane after licensing your XenServer hosts.

3. In the **WLB** tab, click **Connect**. The **Connect to WLB Server** dialog box appears.

4. In the **Server Address** section, dialog box, enter the following:

   - In the **Address** box, type the IP address or FQDN of the Workload Balancing server. An example of an FQDN is WLB-appliance-computername.yourdomain.net .

   - Enter the port number in the **Port** box. XenServer uses this port to communicate with Workload Balancing.

     By default, XenServer connects to Workload Balancing (specifically the Web Service Host service) on port 8012. If you changed the port number during Workload Balancing Configuration, you must enter that port number in the **Port** box.

     > **Important:**
     >
     > Do not change the default port number unless you changed the default port during Workload Balancing Configuration. The port number specified during Configuration, in any firewalls, and in the **Connect to WLB Server** dialog must match.

5. In the **WLB Server Credentials** section, enter the user name (for example, wlbuser) and password the XenServer will use to connect to the Workload Balancing virtual appliance. This must be the account you created during Workload Balancing Configuration. By default, the user name for this account is **wlbuser**.

6. In the **XenServer Credentials** section, enter the user name and password for the pool you are configuring. Workload Balancing will use these credentials to connect to each of the hosts in that pool. To use the credentials with which you are currently logged into XenServer, select the **Use the current XenCenter credentials** check box. If you have assigned permissions to the account using the Role Based Access Control feature (RBAC), be sure they are sufficient to use Workload Balancing. See Definitions of RBAC roles and permissions.

7. After connecting to the Workload Balancing appliance, if you want to change the settings for thresholds or the priority given to specific resources, see Editing Workload Balancing Settings.

Introduction to Basic Tasks

September 30, 2020

Workload Balancing is a powerful XenServer component that includes many features designed to optimize the workloads in your environment, such as host power management, the ability to schedule optimization-mode changes, and running reports. In addition, you can fine-tune the criteria Workload Balancing uses to make optimization recommendations.

However, when you first begin using Workload Balancing, there are two main tasks you will probably use Workload Balancing for on a daily (or regular) basis:

- Determining the best host on which to start a virtual machine
- Accepting Workload Balancing optimization recommendations

For information about obtaining and configuring the Workload Balancing virtual appliance, see the *Workload Balancing Quick Start*.

Determining the best host on which to start a VM

See also Choosing an Optimal Server for VM Initial Placement, Migrate, and Resume.

Workload Balancing can provide recommendations about the host Determining the host on which to start a VM (VM placement) is handy when you want to restart a powered off virtual machine and when you want to migrate a virtual machine to a different host. It may also be useful in Citrix Virtual Desktops environments.

Accepting Workload Balancing recommendations

See also Accepting Optimization Recommendations.

After Workload Balancing is running for a while, it begins to make recommendations about ways in which you can improve your environment. For example, if your goal is to improve VM density on hosts, with the appropriate settings, Workload Balancing will issue a recommendation to consolidate virtual machines on a host. Assuming you are not running in automated mode, you can choose to either apply this recommendation or just simply ignore it.

Both of these tasks, and the way in which you perform them in XenCenter, are explained in more depth in the sections that follow. Another frequently used task is running reports about the workloads in your environment, which is described in Generating and Managing Workload Balancing Reports.

> **Important:**
>
> After Workload Balancing is running for a period of time, if you do not receive optimal placement recommendations, Citrix strongly recommends you evaluate your performance thresholds as described in the *Workload Balancing Administrator's Guide*. It is critical to set Workload Balancing to the correct thresholds for your environment or its recommendations might not be appropriate or occur at the correct times.

Choosing an Optimal Server for VM Initial Placement, Migrate, and Resume

June 20, 2019

When Workload Balancing is enabled and you start a virtual machine that is offline, XenCenter provides recommendations to help you determine the optimal physical host in the resource pool on which to start a virtual machine. The recommendations are also known as star ratings since stars are used to indicate the best host.
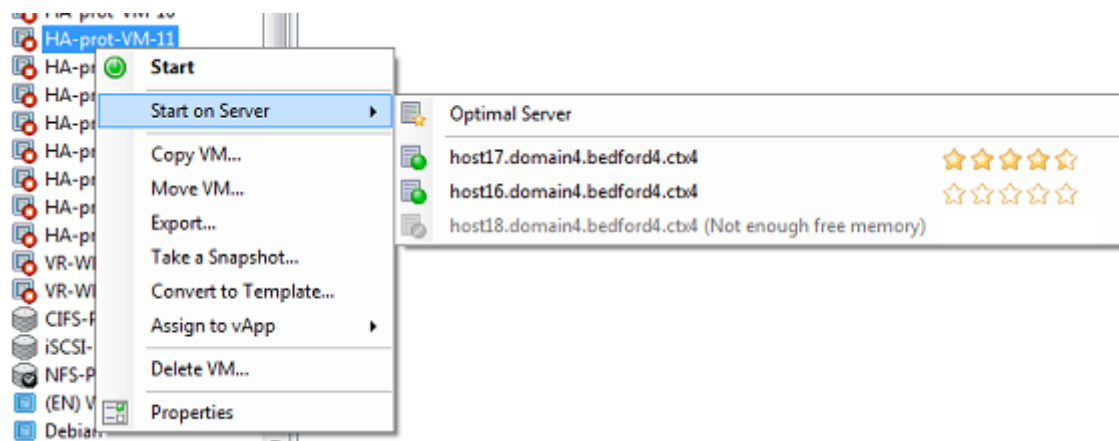
More stars appear beside host17 since this is the optimal host on which to start the VM. host16 does not have any stars beside it, which indicates that host is not recommended, but since it is enabled the user can select that host. host18 is grayed out due to insufficient memory, so the user cannot select it.

How do placements work?

When Workload Balancing is enabled, XenCenter provides star ratings to indicate the optimal hosts for starting a virtual machine, including when you want to start the virtual machine when it is powered off or suspended and when you want to migrate the virtual machine to a different server (Migrate and Maintenance Mode).

When you use these features with Workload Balancing enabled, host recommendations appear as star ratings beside the name of the physical host. Five empty stars indicates the lowest-rated (least optimal) server. When it is not possible to start or move a virtual machine to a host, the host name is grayed out and the reason, it cannot accept the VM, appears beside it.

In general, Workload Balancing functions more effectively and makes better, less frequent optimization recommendations if you start virtual machines on the hosts it recommends (that is, by using one of the placement features to select the host with the most stars beside it).

What does optimal mean?

The term optimal refers to the physical server best suited to hosting your workload. There are several factors Workload Balancing uses when determining which host is optimal for a workload:

- The amount of resources available on each host in the pool. When a pool is running in Maximum Performance mode, Workload Balancing tries to balance the virtual machines across the hosts in the pool so that all virtual machines have good performance. When a pool is running in Maximum Density mode, Workload Balancing tries to place virtual machines onto hosts as densely as possible while ensuring the virtual machines have sufficient resources.

- The optimization mode in which the pool is running (Maximum Performance or Maximum Density). When a pool is running in Maximum Performance mode, Workload Balancing tries to place virtual machines on hosts with the most resources available of the type the VM requires. In Maximum Density mode, Workload Balancing tries to place virtual machines on hosts that already have virtual machines running so that virtual machines are running on as few hosts as possible.

- The amount and type of resources the VM requires. After Workload Balancing has monitored a virtual machine for a while, it uses the VM metrics it gathered to make placement recommendations according to the type of resources the virtual machine requires. For example, Workload Balancing may select a host with less available CPU but more available memory if it is what the VM requires (based on its past performance history). It should, however, be noted that Workload Balancing only makes a recommendation if it determines the current host is under resource pressure.

To start a virtual machine on the optimal server

1. In the **Resources** pane of XenCenter, select the virtual machine you want to start.
2. From the **VM** menu, select **Start on Server** and then select one of the following:
   - **Optimal Server**. The optimal server is the physical host that is best suited to the resource demands of the virtual machine you are starting. Workload Balancing determines the optimal server based on its historical records of performance metrics and your placement strategy. The optimal server is the server with the most stars.
   - One of the servers with star ratings listed under the Optimal Server command. Five stars indicates the most-recommended (optimal) server and five empty stars indicates the least-recommended server.

To resume a virtual machine on the optimal server

1. In the **Resources** pane of XenCenter, select the suspended virtual machine you want to resume.
2. From the **VM** menu, select **Resume on Server** and then select one of the following:
   - **Optimal Server**. The optimal server is the physical host that is best suited to the resource demands of the virtual machine you are starting. Workload Balancing determines the optimal server based on its historical records of performance metrics and your placement strategy. The optimal server is the server with the most stars.
   - One of the servers with star ratings listed under the Optimal Server command. Five stars indicates the most-recommended (optimal) server and five empty stars indicates the least-

recommended server.

Accepting Optimization Recommendations

September 30, 2020

Workload Balancing provides recommendations about ways you can migrate virtual machines to optimize your environment. Optimization recommendations appear in the **WLB** tab in XenCenter.

This illustration shows a screen capture of the Optimization Recommendations list, which appears on the WLB tab. The Reason column displays the purpose of the recommendation. The Operation column displays the behavior change suggested for that optimization recommendation. This screen capture shows an optimization recommendation for a virtual machine, HA-prot-VM-7, and a host, host17.domain4.bedford4.ctx4.

Basis for optimization recommendations

Optimization recommendations are based on the:

- Placement strategy you select (that is, the placement optimization mode), as described in Adjusting the Optimization Mode
- Performance metrics for resources such as a physical host's CPU, memory, network, and disk utilization
- The role of the host in the resource pool. When making placement recommendations, Workload Balancing considers only the pool master if no other host can accept the workload. (Likewise, when a pool is operating in Maximum Density mode, Workload Balancing considers the pool master last when determining the order in which to fill hosts with virtual machines.)

The optimization recommendations display the name of the virtual machine that Workload Balancing recommends relocating, the host it currently resides on, and the host Workload Balancing recommends as the machine's new location. The optimization recommendations also display the reason Workload Balancing recommends moving the virtual machine (for example, "CPU" to improve CPU utilization).

After you accept an optimization recommendation, XenServer relocates all virtual machines listed as recommended for optimization.

> **Tip:**
>
> You can find out the optimization mode for a resource pool by selecting the pool in XenCenter and checking the **Configuration** section of the **WLB** tab.

To accept an optimization recommendation

1. Select the pool for which you want to display recommendations in the **Resources** pane and then click on the **WLB** tab. If there are any recommended optimizations for any virtual machines on the selected resource pool, they display on the **WLB** tab.

2. To accept the recommendations, click **Apply Recommendations**. XenServer begins moving all virtual machines listed in the **Optimization Recommendations** section to their recommended servers.

   After you click **Apply Recommendations** , you can click **Notifications** and then **Events** tab to display the progress of the virtual machine migration.

Understanding WLB recommendations under High Availability

If you have Workload Balancing and the XenServer High Availability feature enabled in the same pool, it is helpful to understand how the two features interact. Workload Balancing is designed not to interfere with High Availability. If there is a conflict between a Workload Balancing recommendation and a High Availability setting, the High Availability setting always takes precedence. In practice, this means:

- Workload Balancing will not automatically power off any hosts beyond the number specified in the **Failures allowed** box in the **Configure HA** dialog.
    - However, Workload Balancing may still make recommendations to power off more hosts than the number of host failures to tolerate. (For example, Workload Balancing will still make a recommendation to power off two hosts when High Availability is only configured to tolerate one host failure.) However, when you attempt to apply the recommendation, XenCenter may display an error message stating that High Availability will no longer be guaranteed.
    - When Workload Balancing is running in automated mode and has power management enabled, any recommendations that exceed the number of host failures to tolerate are ignored. In this situation, if you look in the Workload Balancing log, you will see a message that says a power-management recommendation was not applied because High Availability is enabled.

Working with Workload Balancing Reports

June 3, 2019

This topic provides general information about Workload Balancing historical reports and an overview of where to find additional information about these reports.

To generate a Workload Balancing report, you must have imported the Workload Balancing appliance and connected the pool to it.

Introduction

Workload Balancing provides reporting on three types of objects: physical hosts, resource pools, and virtual machines. At a high level, Workload Balancing provides two types of reports:

- Historical reports that display information by date
- "Roll up" style reports

Workload Balancing provides some reports for auditing purposes, so you can determine, for example, the number of times a virtual machine moved.

Types of Reports

Workload Balancing offers several different reports about the pool, hosts, and VMs. For more information, see Workload Balancing Report Glossary.

Generating reports

Workload Balancing lets you generate reports, export them as PDFs or spreadsheets, and print them out. For more information, see Generating and Managing Workload Balancing reports.

Using Workload Balancing Reports for Tasks

June 3, 2019

The Workload Balancing reports can help you perform capacity planning, determine virtual-machine health, and evaluate the effectiveness of your configured threshold levels.

Evaluating the effectiveness of your performance thresholds

You can use the Pool Health report to evaluate the effectiveness of your optimization thresholds. Work-load Balancing provides default threshold settings. However, you might need to adjust these defaults for them to provide value in your environment. If you do not have the thresholds adjusted to the correct level for your environment, Workload Balancing recommendations might not be appropriate for your environment.

Troubleshooting administrative changes

You can use the Pool Audit Trail report to determine not only the source (that is, user account) of problematic changes but also the event or task that user performed.

Generating and Managing Workload Balancing Reports

September 30, 2020

This topic provides basic instructions for using Workload reports, including how to generate, navigate in, print, and export reports.

To generate a Workload Balancing report

1. In the **Resources** pane of XenCenter, select **your-resource-pool**.

2. From the **Pool** menu, select **View Workload Reports**.

   > **Tip:**
   >
   > You can also display the Workload Reports screen from the **WLB** tab by clicking the **Reports** button.

3. From the Workload Reports screen, select a report from the left pane.

4. Select the **Start Date** and the **End Date** for the reporting period. Depending on the report you select, you might need to specify other parameters such as **Host** , **User** , and **Object**.

5. Click **Run Report**. The report displays in the report window.

To navigate in a Workload Balancing report

After generating a report, you can use the toolbar buttons in the report to navigate and perform certain tasks. To display the name of a toolbar button, pause your mouse over toolbar icon.

Report Toolbar Buttons:

	Document Map. Displays a document map that helps you navigate through long reports.
	Page Forward/Back. Moves one page ahead or back in the report.
	Back to Parent Report. Returns to the parent report when working with drill-through reports. **Note:** This button is only available in drill-through reports, such as the Pool Health report.
	Stop Rendering. Stops the report generation.
	Print. Prints a report and specify general printing options, such as the printer, the number of pages, and the number of copies.
	Print Layout. Displays a preview of the report before you print it. To exit Print Layout, click the **Print Layout** button again.
	Page Setup. Specifies printing options such as the paper size, page orientation, and margins.
	Export. Exports the report as an Acrobat (.pdf) file or as an Excel file with a .xls extension.
	Find. Searches for a word in a report, such as the name of a virtual machine.

To print a Workload Balancing report

Before you can print a report, you must first generate it.

1. (Optional.) To preview the printed document, click **Print Layout**.
2. (Optional.) To change the paper size/source, page orientation, or margins, click **Page Setup**.
3. Click **Print**.

To export a Workload Balancing report

You can export a report in Microsoft Excel and Adobe Acrobat (.pdf) formats.

After generating the report, click 💾 **Export** and select one of the following:

- **Excel**
- **Acrobat (.pdf) file**

> **Note:**
>
> The amount of data included when you export a report may differ depending on the export for‑
> mat. Reports exported to Excel include all the data available for the report, including "drilldown"
> data, whereas reports displayed in XenCenter or exported as .pdf only contain the data that you
> selected when you generated the report.

Workload Balancing Report Glossary

September 30, 2020

This topic provides information about the following Workload Balancing reports.

Chargeback Utilization Analysis

You can use the Chargeback Utilization Analysis report ("chargeback report") to determine how much
of a resource (such as a physical server) a specific department within your organization used. Specif‑
ically, the report shows information about all the virtual machines in your pool, including their avail‑
ability and resource utilization. Since this report shows virtual machine availability ("up time"), it can
help you demonstrate Service Level Agreements compliance and availability.

The chargeback report can help you implement a simple chargeback solution and facilitate billing. To
bill customers for usage of a specific resource, generate the report, save it as Excel, and then customize
the spreadsheet data to include your price per unit or import the Excel data into your billing system.

If you know that you will want to bill internal or external customers for virtual machine usage, consider
incorporating department or customer names in your virtual machine naming conventions. This will
make reading chargeback reports easier.

The resource reporting in the chargeback report is, in some cases, based on the allocation of physical
resources to individual virtual machines.

Likewise, because XenServer lets you allocate fixed or automatic allocations of memory, the average
memory data in this report is based on the amount of memory currently allocated to the virtual ma‑
chine, whether it is through a fixed memory allocation or an automatically adjusting memory alloca‑
tion (Dynamic Memory Control).

The chargeback report contains the following columns of data:

- **VM Name**. The name of the virtual machine to which the data in the columns in that row applies.

- **VM Uptime**. The number of minutes the virtual machine was powered on (or, more specifically, appears with a green icon beside it in XenCenter).

- **vCPU Allocation**. The number of virtual CPUs configured on the virtual machine. Each virtual CPU receives an equal share of the physical CPU(s) on the host. For example, if you configured eight virtual CPUs on a host that contains two physical CPUs and this column had "1" in it, then this value is equal to 2/16 of the total processing power on the host.

- **Minimum CPU Usage (%)**. The lowest recorded value for virtual CPU utilization in the reporting period. This value is expressed as a percentage of the virtual machine's virtual CPU capacity, and the capacity is based on the number of virtual CPUs allocated to the virtual machine. For example, if, when you created the virtual machine, you allocated one virtual CPU to it, Minimum CPU Usage represents the lowest percentage of virtual CPU usage XenServer recorded, even if it was only for a second. If you allocated two virtual CPUs to the virtual machine, the value in this column represents the lowest usage of the combined capacity of both virtual CPUs, expressed as a percentage.

  Ultimately, the percentage of CPU usage represents the lowest recorded workload that virtual CPU handled. For example, if you allocate one virtual CPU to a virtual machine and the physical CPU on the host is 2.4 GHz, you are allocating one-eighth of 2.4 GHz to the virtual machine. This means that if the virtual machine's allocated capacity is 0.3GHz, or 300MHz, and the Minimum CPU Usage for the virtual machine was 20%, the virtual machine's lowest usage of the physical host's CPU during the reporting period was 60MHz.

- **Maximum CPU Usage (%)**. The highest percentage of the virtual machine's virtual CPU capacity that the virtual machine consumed during the reporting period. The CPU capacity consumed is a percentage of the virtual CPU capacity you allocated to the virtual machine. For example, if, when you created the virtual machine, you allocated one virtual CPU to it, the Maximum CPU Usage represents the highest recorded percentage of virtual CPU usage during the time reported. If you allocated two virtual CPUs to the virtual machine, the value in this column represents the highest utilization from the combined capacity of both virtual CPUs.

- **Average CPU Usage (%)**. The average amount, expressed as a percentage, of the virtual machine's virtual CPU capacity that was in use during the reporting period. The CPU capacity is the virtual CPU capacity you allocated to the virtual machine. If you allocated two virtual CPUs to the virtual machine, the value in this column represents the average utilization from the combined capacity of both virtual CPUs.

- **Total Storage Allocation (GB)**. The amount of disk space that is currently allocated to the virtual machine at the time the report was run. Frequently, unless you modified it, this disk space is the amount of disk space you allocated to the virtual machine when you created it.

- **Virtual NIC Allocation**. The number of virtual interfaces (VIFs) allocated to the virtual machine.

- **Current Minimum Dynamic Memory (MB)**.

  – **Fixed memory allocation**. If you assigned a virtual machine a fixed amount of memory (for example, 1024MB), the same amount of memory appears in the following columns: Current Minimum Dynamic Memory (MB), Current Maximum Dynamic Memory (MB), Current Assigned Memory (MB), and Average Assigned Memory (MB).

  – **Dynamic memory allocation**. If you configured XenServer to adjust a virtual machine's memory automatically based on a range (that is, you enabled Dynamic Memory Control), the minimum amount of memory specified in the range appears in this column. For example, if in the Memory Settings dialog box in XenCenter, you selected the Automatically allocate memory within this range option for this virtual machine and then specified the range values as 1024MB as the minimum memory and 2048MB as the maximum memory, then 1024MB will appear in the Current Minimum Dynamic Memory (MB) column.

- **Current Maximum Dynamic Memory (MB)**.

  – **Dynamic memory allocation**. If XenServer is set to adjust a VM's memory automatically based on a range, the maximum amount of memory specified in the range appears in this column. For example, if the memory range you provided was 1024MB minimum and 2048MB maximum, then 2048MB will appear in the Current Maximum Dynamic Memory (MB) column.

  – **Fixed memory allocation**. If you assign a VM a fixed amount of memory (for example, 1024MB), the same amount of memory appears in the following columns: Current Minimum Dynamic Memory (MB), Current Maximum Dynamic Memory (MB), Current Assigned Memory (MB), and Average Assigned Memory (MB).

- **Current Assigned Memory (MB)**.

  – **Dynamic memory allocation**. When Dynamic Memory Control is configured, this value indicates the amount of memory XenServer is currently allocating to the virtual machine at the time at which the report was run.

  – **Fixed memory allocation**. If you assign a virtual machine a fixed amount of memory (for example, 1024MB), the same amount of memory appears in the following columns: Current Minimum Dynamic Memory (MB), Current Maximum Dynamic Memory (MB), Current Assigned Memory (MB), and Average Assigned Memory (MB).

  > **Note:**
  >
  > If you change the virtual machine's memory allocation immediately before running this report, the value reflected in this column reflects the new memory allocation you configured.

- **Average Assigned Memory (MB)**.

  – **Dynamic memory allocation**. When Dynamic Memory Control is configured, this value

indicates the average amount of memory XenServer allocated to the virtual machine over the reporting period.

– **Fixed memory allocation**. If you assign a virtual machine a fixed amount of memory (for example, 1024MB), the same amount of memory appears in the following columns: Current Minimum Dynamic Memory (MB), Current Maximum Dynamic Memory (MB), Current Assigned Memory (MB), and Average Assigned Memory (MB).

> **Note:**
>
> If you change the virtual machine's memory allocation immediately before running this report, the value displayed in this column may not change from what it would have previously displayed. The value in this column reflects the average over the time period.

- **Average Network Reads (BPS)**. The average amount of data (in bits per second) the virtual machine received during the reporting period.

- **Average Network Writes (BPS)**. The average amount of data (in bits per second) the virtual machine sent during the reporting period.

- **Average Network Usage (BPS)**. The combined total (in bits per second) of the Average Network Reads and Average Network Writes. For example, if a virtual machine sent, on average, 1,027 bits per second and received, on average, 23,831 bits per second over the reporting period, then the Average Network Usage would be the combined total of these two values: 24,858 bits per second.

- **Total Network Usage (BPS)**. The total of all network read and write transactions in bits per second over the reporting period.

Host Health History

This report displays the performance of resources (CPU, memory, network reads, and network writes) on specific host in relation to threshold values.

The colored lines (red, green, yellow) represent your threshold values. You can use this report with the Pool Health report for a host to determine how a particular host's performance might be affecting overall pool health. When you are editing the performance thresholds, you can use this report for insight into host performance.

You can display resource utilization as a daily or hourly average. The hourly average lets you see the busiest hours of the day, averaged, for the time period.

To view report data grouped by hour, expand **+ Click to view report data grouped by house for the time period** under the **Host Health History** title bar.

Workload Balancing displays the average for each hour for the time period you set. The data point is based on a utilization average for that hour for all days in the time period. For example, in a report for

May 1, 2009 to May 15, 2009, the Average CPU Usage data point represents the resource utilization of all fifteen days at 12:00 hours combined together as an average. That is, if CPU utilization was 82% at 12 PM on May 1st, 88% at 12 PM on May 2nd, and 75% on all other days, the average displayed for 12 PM is 76.3%.

Pool Optimization Performance History

The optimization performance report displays optimization events (that is, when you optimized a resource pool) against that pool's average resource usage. Specifically, it displays resource usage for CPU, memory, network reads, and network writes.

The dotted line represents the average usage across the pool over the period of days you select. A blue bar indicates the day on which you optimized the pool.

This report can help you determine if Workload Balancing is working successfully in your environment. You can use this report to see what led up to optimization events (that is, the resource usage before Workload Balancing recommended optimizing).

This report displays average resource usage for the day; it does not display the peak utilization, such as when the system is stressed. You can also use this report to see how a resource pool is performing if Workload Balancing is not making optimization recommendations.

In general, resource usage should decline or be steady after an optimization event. If you do not see improved resource usage after optimization, consider readjusting threshold values. Also, consider whether or not the resource pool has too many virtual machines and whether or not new virtual machines were added or removed during the time frame you specified.

Pool Audit Trail

This report displays the contents of the XenServer Audit Log , A XenServer feature designed to log attempts to perform unauthorized actions and select authorized actions, including import/export, host and pool backups, and guest and host console access. The report gives more meaningful information when XenServer administrators are given their own user accounts with distinct roles assigned to them using the Role-Based Access Control feature. For information about the Audit Log feature, see the audit log documentation in the *Workload Balancing Administrator's Guide*.

> **Important:**
>
> To run the audit log report, the Audit Logging feature must be enabled. By default, Audit Log is always enabled in the Workload Balancing virtual appliance.

The enhanced Pool Audit Trail feature allows you to specify the granularity of the audit log report. You can also search and filter the audit trail logs by specific users, objects, and by time. The Pool Audit Trail Granularity is set to **Minimum** by default. This option captures limited amount of data for

specific users and object types. You can modify the setting at any time based on the level of detail you would require in your report. For example, set the granularity to **Medium** for a user-friendly report of the audit log. If you require a detailed report, set the option to **Maximum**.

To modify the Pool Audit Trail Granularity setting:

1. Select the pool in the **Infrastructure** view, click on the **WLB** tab, and then click **Settings**.

2. In the left pane, click **Advanced**.

3. On the Advanced page, click the **Pool Audit Trail Report Granularity** drop-down list and select an option from the list.

   > **Important:**
   >
   > Select the granularity based on the your audit log requirements. For example, if you set your audit log report granularity to **Minimum** , the audit report only captures limited amount of data for specific users and object types. If you set the granularity to **Medium** , the report provides a user-friendly report of the audit log. If you choose to set the granularity to **Maximum** , the report will contain detailed information about the audit log report. Note that setting the audit log report to Maximum can cause the Workload Balancing server to use more disk space and memory.

4. Click **OK** to confirm your changes.

This report displays the following:

- **Time**. The time XenServer recorded the user's action.
- **User Name**. The name of the person who created the session in which the action was performed. In some cases, this may be the User ID.
- **Event Object**. The object that was the subject of the action (for example, a virtual machine).
- **Event Action**. The action that occurred. For definitions of these actions, see Audit Log Event Names.
- **Access**. Whether or not the user had permission to perform the action.
- **Object Name**. The name of the object (for example, the name of the virtual machine).
- **Object UUID**. The UUID of the object (for example, the UUID of the virtual machine).
- **Succeeded**. This provides the status of the action (that is, whether or not it was successful).

Pool Health

The pool health report displays the percentage of time a resource pool and its hosts spent in four different threshold ranges: Critical, High, Medium, and Low. You can use the Pool Health report to evaluate the effectiveness of your performance thresholds.

A few points about interpreting this report:

- Resource utilization in the Average Medium Threshold (blue) is the optimum resource utilization regardless of the placement strategy you selected. Likewise, the blue section on the pie chart indicates the amount of time that host used resources optimally.
- Resource utilization in the Average Low Threshold Percent (green) is not necessarily positive. Whether Low resource utilization is positive depends on your placement strategy. For example, if your placement strategy is Maximum Density and most of the time your resource usage was green, Workload Balancing might not be fitting the maximum number of virtual machines possible on that host or pool. If this is the case, you should adjust your performance threshold values until the majority of your resource utilization falls into the Average Medium (blue) threshold range.
- Resource utilization in the Average Critical Threshold Percent (red) indicates the amount of time average resource utilization met or exceeded the Critical threshold value.

If you double-click on a pie chart for a host's resource usage, XenCenter displays the Host Health History report for that resource (for example, CPU) on that host. Clicking the **Back to Parent Report** toolbar button returns you to the Pool Health history report. Note: This button is only available in drill-through reports, such as the Pool Health report.

If you find the majority of your report results are not in the Average Medium Threshold range, you probably need to adjust the Critical threshold for this pool. While Workload Balancing provides default threshold settings, these defaults are not effective in all environments. If you do not have the thresholds adjusted to the correct level for your environment, Workload Balancing's optimization and placement recommendations might not be appropriate. For more information, see Changing the Critical Thresholds.

> **Note:**
>
> The High, Medium, and Low threshold ranges are based on the Critical threshold value.

Pool Health History

This report provides a line graph of resource utilization on all physical hosts in a pool over time. It lets you see the trend of resource utilization - if it tends to be increasing in relation to your thresholds (Critical, High, Medium, and Low). You can evaluate the effectiveness of your performance thresholds by monitoring trends of the data points in this report.

Workload Balancing extrapolates the threshold ranges from the values you set for the Critical thresholds. Although similar to the Pool Health report, the Pool Health History report displays the average utilization for a resource on a specific date rather than the amount of time overall the resource spent in a threshold.

With the exception of the Average Free Memory graph, the data points should never average above the Critical threshold line (red). For the Average Free Memory graph, the data points should never

average below the Critical threshold line (which is at the bottom of the graph). Because this graph displays free memory, the Critical threshold is a low value, unlike the other resources.

A few points about interpreting this report:

- When the Average Usage line in the chart approaches the Average Medium Threshold (blue) line, it indicates the pool's resource utilization is optimum regardless of the placement strategy configured.
- Resource utilization approaching the Average Low Threshold (green) is not necessarily positive. Whether Low resource utilization is positive depends on your placement strategy. For example, if your placement strategy is Maximum Density and most days the Average Usage line is at or below the green line, Workload Balancing might not be placing virtual machines as densely as possible on that pool. If this is the case, you should adjust the pool's Critical threshold values until the majority of its resource utilization falls into the Average Medium (blue) threshold range.
- When the Average Usage line intersects with the Average Critical Threshold Percent (red), this indicates the days when the average resource utilization met or exceeded the Critical threshold value for that resource.

If you find the data points in the majority of your graphs are not in the Average Medium Threshold range, but you are satisfied with the performance of this pool, you might need to adjust the Critical threshold for this pool. For more information, see Changing the Critical Thresholds.

Pool Optimization History

The Pool Optimization History report provides chronological visibility into Workload Balancing optimization activity.

Optimization activity is summarized graphically and in a table. Drilling into a date field within the table displays detailed information for each pool optimization performed for that day.

This report lets you see the following information:

- **VM Name**. The name of the virtual machine that Workload Balancing optimized.
- **Reason**. The reason for the optimization.
- **Status**. If the optimization was successful.
- **From Host**. The physical server where the virtual machine was originally hosted.
- **To Host**. The physical server where the virtual machine was moved.
- **Time**. The time when the optimization occurred.

> **Tip:**
>
> You can also generate a Pool Optimization History report from the **WLB** tab, by clicking the **View History** link.

Virtual Machine Motion History

This line graph displays the number of times virtual machines moved on a resource pool over a period of time. It indicates if a move resulted from an optimization recommendation and to which host the virtual machine moved. This report also indicates the reason for the optimization. You can use this report to audit the number of moves on a pool.

Some points about interpreting this report:

- The numbers on the left side of the chart correspond with the number of moves possible, which is based on how many virtual machines are in a resource pool.
- You can look at details of the moves on a specific date by expanding the + sign in the **Date** section of the report.

Virtual Machine Performance History

This report displays performance data for each virtual machine on a specific host for a time period you specify. Workload Balancing bases the performance data on the amount of virtual resources allocated for the virtual machine. For example, if the Average CPU Usage for your virtual machine is 67%, this means that your virtual machine was using, on average, 67% of its virtual CPU for the period you specified.

The initial view of the report displays an average value for resource utilization over the period you specified.

Expanding the + sign displays line graphs for individual resources. You can use these graphs to see trends in resource utilization over time.

This report displays data for CPU Usage, Free Memory, and Network Reads/Writes.

Audit Log Events

June 3, 2019

The Audit Log report logs XenServer events, event objects and actions, including import/export, host and pool backups, and guest and host console access. The following table defines some of the typical events that appear frequently in the XenServer Audit Log and Pool Audit Trail report, and it also specifies the granularity of these events.

In the Pool Audit Trail report, the events listed in the Event Action column apply to a pool, VM, or host. To determine what the events apply to, see the **Event Object** and **Object Name** columns in the report. For additional event definitions, see the *Citrix XenServer Management API Guide*. To know more about the Pool Audit Trail granularity settings, see Advanced Settings.

Pool Audit Trail Granularity	Event Action	User Action
Minimum	VM.start	Started a virtual machine.
Minimum	VM.copy	Copied the specified VM, making a new VM.
Minimum	host.reboot	Restarted a XenServer host.
Minimum	host.disable	Put the host into a state in which no new VMs can be started.
Minimum	pool.join	Instructed host to join a new pool.
Minimum	pool.join_force	Instructed (forced) host to join a new pool.
Medium	SR.destroy	Destroyed the storage repository.
Medium	SR.create	Created a new storage repository.
Medium	VDI.snapshot	Took a read-only snapshot of the VDI, returning a reference to the snapshot.
Medium	VDI.clone	Took an exact copy of the VDI and returned a reference to the new disk.
Medium	VIF.plug	Hotplugged the specified VIF, dynamically attaching it to the running VM.
Medium	VIF.unplug	Hot-unplugged the specified VIF, dynamically detaching it from the running VM.
Maximum	auth.get_subject_identifier	Queried the external directory service to obtain the subject identifier as a string from the human-readable subject name.
Maximum	task.cancel	Requested that a task be cancelled.

Pool Audit Trail Granularity	Event Action	User Action
Maximum	VBD.insert	Inserted new media into the device.
Maximum	VIF.get_by_uuid	Got a reference to the VIF instance with the specified UUID.
Maximum	VDI.get_sharable	Got the sharable field of the given VDI.
Maximum	SR.get_all	Returned a list of all the SRs known to the system.
Maximum	pool.create_new_blob	Created a placeholder for a named binary blob of data that is associated with this pool.
Maximum	host.send_debug_keys	Injected the given string as debugging keys into Xen.
Maximum	VM.get_boot_record	Returned a record describing the VM's dynamic state, initialised when the VM boots and updated to reflect runtime configuration changes, for example,CPU hotplug.

Editing Workload Balancing Settings

September 30, 2020

After connecting to the Workload Balancing virtual appliance, you can edit the settings Workload Balancing uses to calculate placement and optimization recommendations. You can perform tasks such as the following:

- Adjusting the Optimization Mode
- Setting Automation and Power Management
- Changing the Critical Thresholds
- Tuning Metric Weightings

- Excluding Hosts from Recommendations
- Advanced Settings

Note:

After connecting or reconnecting to Workload Balancing, wait at least sixty seconds (until the Workload Balancing (/var/log/wlb/LogFile.log) shows discovery is finished) before changing settings.

How Workload Balancing settings apply

Workload Balancing settings apply collectively to all virtual machines and hosts in the pool.

Provided the network and disk thresholds align with the hardware in your environment, you might want to consider using most of the defaults in Workload Balancing initially.

After Workload Balancing is enabled for a while, Citrix recommends evaluating your performance thresholds and determining if you need to edit them. For example, consider if you are:

- Getting optimization recommendation when they are not yet required. If this is the case, try adjusting the thresholds until Workload Balancing begins providing suitable optimization recommendations.
- Not getting recommendations when you think your network has insufficient bandwidth. If this is the case, try lowering the network critical thresholds until Workload Balancing begins providing optimization recommendations.

Before you edit your thresholds, you might find it handy to generate a Host Health History report for each physical host in the pool.

You can use either the **Workload Balancing Configuration** properties in XenCenter or the XE commands to modify the configuration settings.

To update the credentials XenServer and the Workload Balancing server use to communicate, see Updating Workload Balancing Credentials.

For more detailed guidance about tuning Workload Balancing settings, see the *Workload Balancing Administrator's Guide*.

To display the Workload Balancing Settings dialog box

1. In the **Resources** pane of XenCenter, select **your-resource-pool**.
2. In the **Properties** pane, click the **WLB** tab.
3. In the **WLB** tab, click **Settings**.

Adjusting the Optimization Mode

June 3, 2019

Workload Balancing makes recommendations to rebalance, or optimize , the virtual machine workload in your environment based on a strategy for placement you select known as the optimization mode .

Maximum Density and Maximum Performance modes

You can select one of two optimization modes:

- **Maximize Performance**. (Default.) Workload Balancing attempts to spread workload evenly across all physical hosts in a resource pool. The goal is to minimize CPU, memory, network, and disk pressure for all hosts. When Maximize Performance is your placement strategy, Workload Balancing recommends optimization when a virtual machine reaches the High threshold.

- **Maximize Density**. Workload Balancing attempts to fit as many virtual machines as possible onto a physical host. The goal is to minimize the number of physical hosts that must be online.

  When you select Maximize Density as your placement strategy, you can specify rules similar to the ones in Maximize Performance. However, Workload Balancing uses these rules to determine how it can pack virtual machines onto a host. When Maximize Density is your placement strategy, Workload Balancing recommends optimization when a virtual machine reaches the Critical threshold.

Workload Balancing also lets you apply these optimization modes all of the time, Fixed , or switch between modes for specified time periods, Scheduled .

Fixed

Fixed optimization modes set Workload Balancing to a specific optimization behavior - either to try to create the best performance or the highest density - at all times.

Scheduled

Scheduled optimization modes let you schedule for Workload Balancing to apply different optimization modes depending on the time of day. For example, you might want to configure Workload Balancing to optimize for performance during the day when you have users connected and then, to save energy, specify for Workload Balancing to optimize for Maximum Density at night.

When you configure Scheduled optimization modes, Workload Balancing automatically changes to the optimization mode at the beginning of the time period—day of the week and time—you specified.

To set an optimization mode for all time periods

- Select the pool in the **Resources** pane, click the **WLB** tab, and then click **Settings**.
- In the left pane, click **Optimization Mode**.
- Select **Fixed** , and select one of these optimization modes:
  - **Maximize Performance**. (Default.) Attempts to spread workload evenly across all physical hosts in a resource pool. The goal is to minimize CPU, memory, network, and disk pressure for all hosts.
  - **Maximize Density**. Attempts to fit as many virtual machines as possible onto a physical host. The goal is to minimize the number of physical hosts that must be online. (Note that Workload Balancing considers the performance of consolidated VMs and issues a recommendation to improve performance if a resource on a host reaches a Critical threshold.)

To specify times when the optimization mode will change automatically

- Select the pool in the **Resources** pane, click on the **WLB** tab, and then click **Settings**.
- In the left pane, click **Optimization Mode**.
- Select **Scheduled**.
- Click **Add New** to open the **Optimization Mode Scheduler** dialog box.
- Select an optimization mode in the **Change to** list box:
  - **Maximize Performance**. Attempts to spread workload evenly across all physical hosts in a resource pool. The goal is to minimize CPU, memory, network, and disk pressure for all hosts.
  - **Maximize Density**. Attempts to fit as many virtual machines as possible onto a physical host. The goal is to minimize the number of physical hosts that must be online.
- Select the day of the week and the time when you want Workload Balancing to begin operating in this mode.
- Create more scheduled mode changes (that is, "tasks") until you have the number you need. If you only schedule one task, Workload Balancing will switch to that mode and never switch back.
- Click **OK**.

To delete or pause a scheduled optimization mode task

1. Select the pool in the **Resources** pane, click on the **WLB** tab, and then click **Settings**.

2. Click **Optimization Mode**.

3. Select a scheduled task and click one of the following:

   - **Delete the task permanently**. Click the **Delete** button.
   - **Stop the task from running temporarily**. Right-click the task and click **Disable**.

To re-enable a task, right click the task in the **Scheduled Mode Changes** list.

To edit a scheduled optimization mode task

1. Select the pool in the **Resources** pane, click on the **WLB** tab, and then click **Settings**.
2. Select a scheduled task.
3. Click **Edit**.
4. In the **Change to** box, select a different mode or make other changes as desired.

Optimizing and Managing Power Automatically

September 30, 2020

You can configure Workload Balancing to accept optimization recommendations automatically (Automation) and turn servers on or off automatically (Power Management).

Accepting Optimization Recommendations automatically

Workload Balancing lets you configure for it to accept optimization recommendations on your behalf and perform the optimization actions it recommends automatically. You can use this feature, which is known as Automation , to apply any recommendations automatically, including ones to improve performance or power down hosts. However, to power down hosts as virtual-machines usage drops, you must configure automation, power management, and Maximum Density mode.

By default, Workload Balancing does not accept optimizations automatically. You must enable Automation if you want Workload Balancing to accept recommendations automatically. If you do not, Workload Balancing still prompts you to accept recommendations manually.

Workload Balancing will not automatically apply recommendations to hosts or virtual machines if the recommendations conflict with High Availability settings. If a pool will become overcommitted by applying Workload Balancing optimization recommendations, XenCenter will prompt you whether or not you want to continue applying the recommendation. When Automation is enabled, Workload Balancing will not apply any power-management recommendations that exceed the number of host failures to tolerate in the High Availability plan.

It is possible to tweak how Workload Balancing applies recommendations in automated mode. For information, see Advanced Settings.

Enabling Power Management

The term power management refers to the ability to the turn the power on or off for physical hosts. In a Workload Balancing context, this term refers to powering hosts in a pool on or off based on the pool's total workload.

Configuring Workload Balancing power management on a host requires that:

- The hardware for the host server has remote power on/off capabilities
- The Host Power On feature is configured for the host
- The host has been explicitly selected as a host to participate in (Workload Balancing) Power Management

In addition, if you want Workload Balancing to power off hosts automatically, you also need to configure:

- Workload Balancing is configured to apply recommendations automatically
- Workload Balancing is configured to apply Power Management recommendations automatically

When enabled and the pool is in Maximum Density mode, if Workload Balancing detects unused resources, it recommends powering off hosts until it eliminates all excess capacity in the pool. If Workload Balancing detects there is not sufficient host capacity in the pool to turn off servers, it recommends leaving the servers on until the pool's workload decreases enough to power off servers. When you configure Workload Balancing to power off extra servers automatically, it applies these recommendations automatically and, consequently, behaves in the same way.

When a host is set to participate in Power Management, Workload Balancing makes power-on/off recommendations as needed. If you turn on the option to apply Power Management recommendations automatically, you do so at the pool level. However, you can specify which hosts from the pool you want to participate in Power Management.

Understanding Power Management Behavior

Before Workload Balancing recommends powering hosts on or off, it selects the hosts to transfer virtual machines to (that is, to "fill"). It does so in the following order:

1. Filling the pool master since it is the host that cannot be powered off.
2. Filling the host with the most virtual machines.
3. Filling subsequent hosts according to which hosts have the most virtual machines running.

When Workload Balancing fills the pool master, it does so assuming artificially low (internal) thresholds for the master. Workload Balancing uses these low thresholds as a buffer to prevent the pool master from being overloaded.

Workload Balancing fills hosts in this order to encourage density.
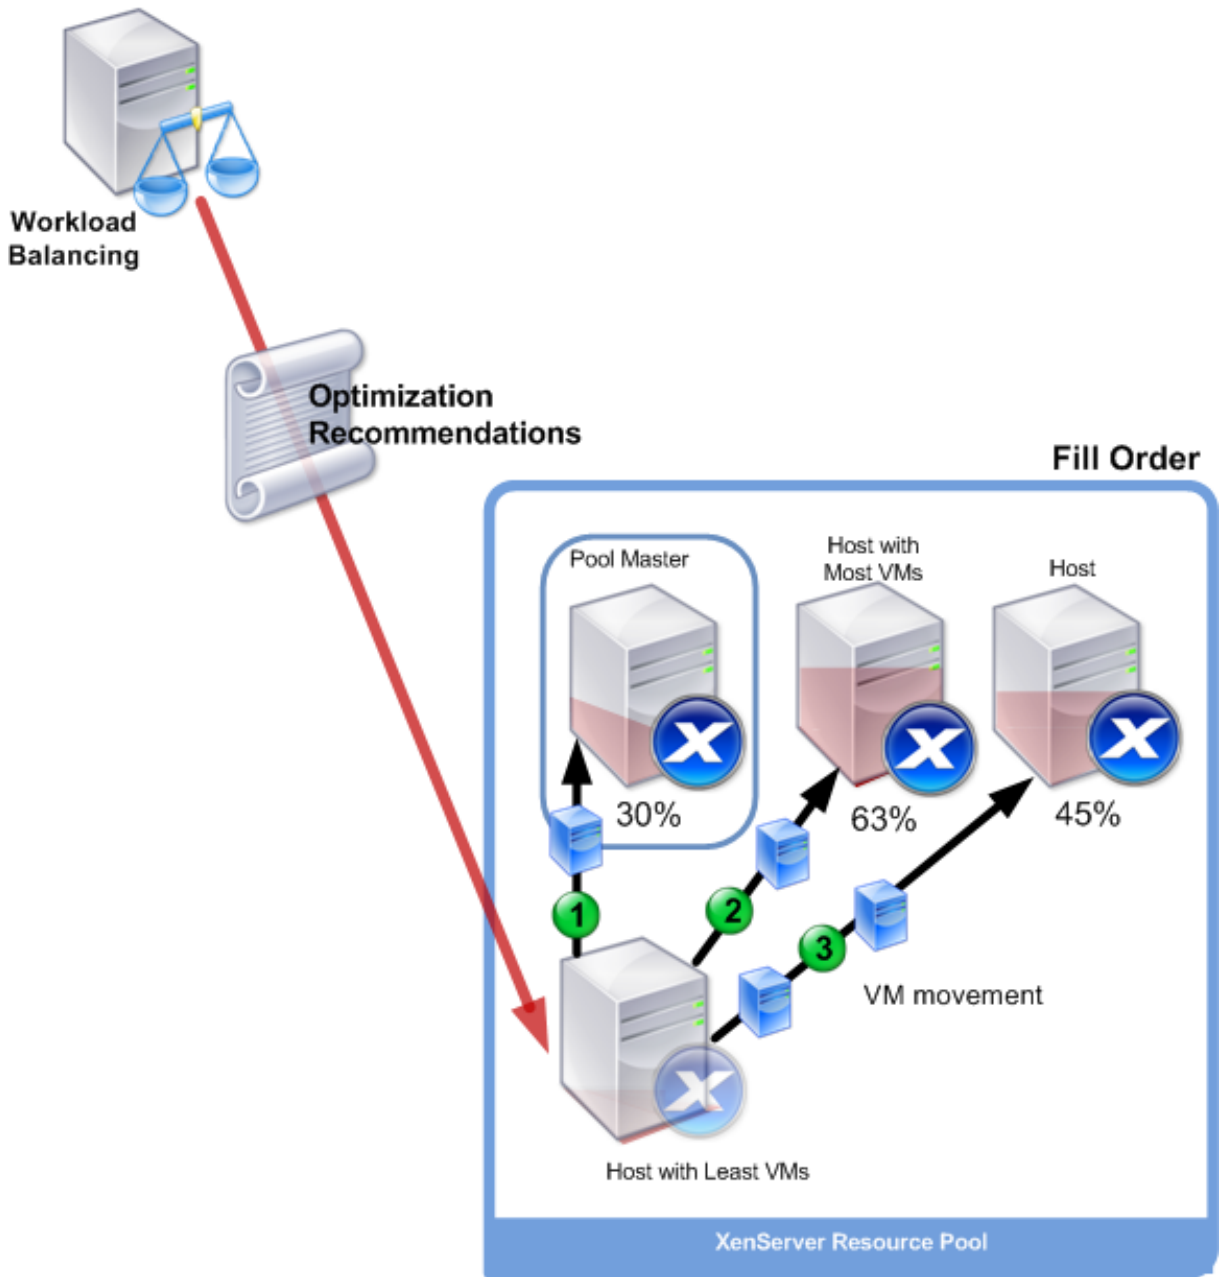
This illustration shows how, when consolidating VMs on hosts in Maximum Density mode, XenServer attempts to fill the pool master first, the most loaded server second, and the least loaded server third.

If Workload Balancing detects a performance issue while the pool is in Maximum Density mode, it attempts to address the issue by recommending migrating workloads among the powered-on hosts. If Workload Balancing cannot resolve the issue using this method, it attempts to power-on a host. (Workload Balancing determines which host(s) to power-on by applying the same criteria it would if the optimization mode was set to Maximum Performance.)

When Workload Balancing is running in Maximum Performance mode, Workload Balancing recommends powering on hosts until the resource utilization on all hosts in the pool falls below the High

threshold.

If, while migrating one or more virtual machines, Workload Balancing determines that increasing capacity would benefit the pool's overall performance, it powers on hosts automatically or recommends doing so.

> **Important:**
>
> Workload Balancing never recommends powering on a host unless Workload Balancing powered it off.

Designing Environments for Power Management and VM Consolidation

When you are planning a XenServer implementation and you intend to configure automatic VM consolidation and power management, consider your workload design. For example, you may want to:

- **Place Different Types of Workloads in Separate Pools.** If you have an environment with distinct types of workloads (for example, user applications versus domain controllers) or types of applications that perform better with certain types of hardware, consider if you need to locate the virtual machines hosting these workloads in different pools.

  Because power management and VM consolidation are managed at the pool level, you should design pools so they contain workloads that you want consolidated at the same rate, factoring in considerations such as those discussed in the Advanced Settings topic.

- **Exclude Hosts from Workload Balancing.** Some hosts may need to be on at all times. See Excluding Hosts from Recommendations for more information.

To apply optimization recommendations automatically

1. In the **Resources** pane of XenCenter, select **XenCenter > your-resource-pool**.
2. In the **Properties** pane, click the **WLB** tab.
3. In the **WLB** tab, click **Configure WLB**.
4. In the left pane, click **Automation**.
5. Select one or more of the following check boxes:
   - **Automatically apply Optimization recommendations.** When you select this option, you do not need to accept optimization recommendations manually. Workload Balancing automatically accepts optimization and placement recommendations it makes.
   - **Automatically apply Power Management recommendations.** The behavior of this option varies according to the pool's optimization mode:
     - **Maximum Performance Mode.** When **Automatically apply Power Management recommendations** is enabled, Workload Balancing automatically powers on hosts when doing so improves host performance.

- **Maximum Density Mode.** When **Automatically apply Power Management recommendations** is enabled, Workload Balancing automatically powers off hosts when resource utilization drops below the Low threshold (that is, Workload Balancing powers hosts off automatically during low usage periods).

6. Do one of the following:
   - If you want to configure power management, click **Automation/Power Management** and proceed to the section below.
   - If you do not want to configure power management and you are finished configuring automation, click **OK**.

To select servers for power management

1. In the Power Management section, select the hosts that you want Workload Balancing to power on and off automatically.

   > **Note:**
   >
   > Selecting hosts for power management recommendations without selecting the **Automatically apply Power Management recommendations** check box results in Workload Balancing suggesting power management recommendations but not applying them automatically for you.

2. Click **OK**. If none of the physical servers in the resource pool support remote power management, Workload Balancing displays the message, **No hosts support Power Management**

Changing the Critical Thresholds

September 30, 2020

This topic provides guidance about how to modify the default Critical thresholds and how values set for Critical threshold alter High, Medium, and Low thresholds.

This information is only provided for reference while changing thresholds. To understand the concepts discussed in this topic, it is important to read them in the fuller context of the information provided in the *Workload Balancing Administrator's Guide*.

Overview

When evaluating utilization, Workload Balancing compares its daily average to four thresholds: low, medium, high, and critical. After you specify (or accept the default) critical threshold, Workload Balancing sets the other thresholds relative to the critical threshold on a pool. You might want to change Critical thresholds as a way of controlling when optimization recommendations are triggered.

Workload Balancing evaluates CPU, Memory, Network Read, Network Write, Disk Read, and Disk Write utilization for physical hosts in a resource pool.

Workload Balancing determines whether to recommend relocating a workload and whether a physical host is suitable for a virtual-machine workload by evaluating:

- Whether a resource's critical threshold is met on the physical host
- (If the critical threshold is met) the importance assigned to a resource

> **Note:**
>
> To prevent data from appearing artificially high, Workload Balancing evaluates the daily averages for a resource and smooths utilization spikes.

For more information, see the *Workload Balancing Administrator's Guide*.

Workload Balancing determines whether or not to produce recommendations based on if the averaged historical utilization for a resource violates its threshold. As discussed in the *Workload Balancing Administrator's Guide* , **Workload Balancing** recommendations are triggered when the **High threshold** in **Maximum Performance** mode or Low and Critical thresholds for Maximum Density mode are violated. After you specify a new Critical threshold for a resource, **Workload Balancing** resets the resource's other thresholds relative to the new **Critical threshold**. (To simplify the user interface, the **Critical threshold** is the only threshold you can change through XenCenter.)

Default settings for thresholds

The following table shows the default values for the Workload Balancing thresholds:

Setting	Default	High	Medium	Low
CPU Utilization	90%	76.5%	45%	22.5%
Free Memory	51MB	63.75MB	510GB	1020GB
Network Read	25MB/sec	21.25MB/sec	12.5MB/sec	6.25MB/sec
Network Write	25MB/sec	21.25MB/sec	12.5MB/sec	6.25MB/sec
Disk Read	25MB/sec	21.25MB/sec	12.5MB/sec	6.25MB/sec
Disk Write	26MB/sec	21.25MB/sec	12.5MB/sec	6.25MB/sec

To calculate the values for the High, Medium, and Low resource metrics, Workload Balancing multiplies the new value for the Critical threshold with the following factors:

- High Threshold Factor: 0.85

- Medium Threshold Factor: 0.50
- Low Threshold Factor: 0.25

To calculate threshold values for free memory, Workload Balancing multiplies the Critical threshold with these factors:

- High Threshold Factor: 1.25
- Medium Threshold Factor: 10.0
- Low Threshold Factor: 20.0

This means that if you increase, for example, the pool's Critical threshold for CPU Utilization to 95%, Workload Balancing automatically resets the High, Medium, and Low thresholds to 80.75%, 47.5%, and 23.75% respectively.

To perform this calculation for a specific threshold, multiply the factor for the threshold with the value you entered for the critical threshold for that resource:

High, Medium, or Low Threshold = Critical Threshold * Threshold Factor

For example, if you change the Critical threshold for Network Reads to 40MB/sec and you want to know its Low threshold, you multiply 40 by 0.25, which equals 10MB/sec. To obtain the Medium threshold, you multiple 40 by 0.50, and so on.

To prevent the pool master from becoming overloaded, Workload Balancing automatically sets the pool master's Critical Thresholds at lower values.

How other thresholds trigger recommendations

While the Critical threshold triggers many recommendations, other thresholds can also trigger recommendations, as follows:

High threshold.

- **Maximum Performance**.   Exceeding the High threshold triggers optimization recommendations to relocate a virtual machine to a host with lower resource utilization.
- **Maximum Density**. Workload Balancing will not recommend placing a virtual machine on host if moving that virtual machine to the host will cause the utilization of any of the host's resources to exceed the High threshold value.

Low threshold.

- **Maximum Performance**. Workload Balancing does not trigger recommendations from the Low threshold.
- **Maximum Density**. When a metric value drops below the Low threshold, it signals Workload Balancing that hosts are being underutilized and triggers an optimization recommendation to consolidate virtual machines on fewer hosts. Workload Balancing continues to recommend

moving virtual machines onto a host until the metric values for one of the host's resource reaches its High threshold.

However, if after a virtual machine is relocated, a resource's utilization on the virtual machine's new host exceeds its Critical threshold, Workload Balancing will temporarily use an algorithm similar to the Maximum Performance load-balancing algorithm to find a new host for the virtual machines. Workload Balancing continues to use this algorithm to recommend moving virtual machines until resource utilization on hosts across the pool falls below the High threshold.

To change the critical thresholds

1. Select the pool in the **Resources** pane, click on the **WLB** tab, and then click **Settings**.
2. In the left pane, select **Critical Thresholds**.
3. In **Critical Thresholds** page, accept or enter a new value in the **Critical Thresholds** boxes. Workload Balancing uses these thresholds when making virtual-machine placement and pool-optimization recommendations. Workload Balancing strives to keep resource utilization on a host below the critical values set.

Tuning Metric Weightings

September 30, 2020

> **Note:**
>
> Before tuning metric weightings, Citrix recommends reading about the optimization and consolidation process in the *Workload Balancing Administrator's Guide*. The information in this help topic is a subset of the information in that guide and is only meant to be used as a reference when changing the user interface.

Workload Balancing uses metric weightings , a method of assigning importance to resources, to determine what hosts it should optimize first.

The way in which Workload Balancing uses metric weightings when determining which hosts and VMs to process first varies according to the optimization mode, Maximum Density or Maximum Performance.

When **Workload Balancing** is processing optimization recommendations, it creates an optimization order . To determine this, Workload Balancing ranks the hosts to address first according to which hosts have the highest metric values for whatever resource is ranked as the most important in the metric weightings page.

In general, metric weightings are mainly used when a pool is in Maximum Performance mode. However, when Workload Balancing is in Maximum Density mode, it does use metric weightings if a resource exceeds its Critical threshold.

How metric weightings apply in Maximum Performance mode

In **Maximum Performance** mode, **Workload Balancing** uses metric weightings to determine (a) which hosts' performance to address first and (b) which VMs to recommend migrating first.

For example, if you rank Network Writes as the most important resource for Workload Balancing to consider—that is, you move the **Metric Weighting** slider all the way to the right (More Important) and you move all the other sliders to the middle—then Workload Balancing will begin addressing performance issues (that is, making optimization recommendations) on the host with the highest number of Network Writes per second first.

How metric weightings apply in Maximum Density mode

In Maximum Density mode, Workload Balancing only uses metric weightings when a host reaches the Critical threshold—then Workload Balancing applies a Maximum Performance-like algorithm until no Hosts are exceeding the Critical thresholds. When using the Maximum Performance-like algorithm, Workload Balancing uses metric weightings to determine the optimization order in the same way as it does for Maximum Performance mode.

If two or more hosts have resource(s) exceeding their Critical thresholds, Workload Balancing verifies the importance you set for each resource before determining which host to optimize first and which VMs on that host to relocate first.

For example, your pool contains Host A and Host B, which are in the following state:

- The CPU utilization on Host A exceeds the Critical threshold for CPU, and the metric weighting for CPU utilization is set to the far right of the slider **(More Important)**.
- The memory utilization on Host B exceeds the Critical threshold for memory, and the metric weighting for memory utilization is set to the far left of the slider **(Less Important)**.

Workload Balancing will recommend optimizing Host A first because the resource on it that reached the Critical threshold is the resource assigned the highest weight. After Workload Balancing determines that it needs to address the performance on Host A, Workload Balancing then begins recommending placements for VMs on that host beginning with the VM that has the highest CPU utilization, since that CPU utilization is the resource with the highest weight.

After Workload Balancing has recommended optimizing Host A, it makes optimization recommendations for Host B. When it recommends placements for the VMs on Host B, it does so by addressing CPU utilization first, since CPU utilization was assigned the highest weight.

If there are additional hosts that need optimization, Workload Balancing addresses the performance on those hosts according to what host has the third highest CPU utilization.

By default, all metric weightings are set to the farthest point on the slider **(More Important)**.

> **Note:**
>
> The weighting of metrics is relative. This means that if all of the metrics are set to the same level, even if that level is **Less Important** , they will all be weighted the same. The relation of the metrics to each other is more important than the actual weight at which you set each metric.

To edit metric weighting factors

1. Select the pool in the **Resources** pane, click on the **WLB** tab, and then click **Settings**.

2. In the left pane, select **Metric Weighting**.

3. In the **Metric Weighting** page, as desired, adjust the sliders beside the individual resources.

   Moving the slider towards **Less Important** indicates that ensuring virtual machines always have the highest amount of this resource available is not as vital on this resource pool.

Excluding Hosts from Recommendations

June 3, 2019

When configuring Workload Balancing, you can specify that specific physical hosts are excluded from Workload Balancing optimization and placement recommendations, including Start On placement recommendations.

When to exclude hosts

Situations when you may want to exclude hosts from recommendations include when:

- You want to run the pool in Maximum Density mode and consolidate and shut down hosts, but there are specific hosts you want to exclude from this behavior.
- When two virtual-machine workloads always need to run on the same host (for example, if they have complementary applications or workloads).
- You have workloads that you do not want moved (for example, domain controllers or SQL Server).
- You want to perform maintenance on a host and you do not want virtual machines placed on the host.
- The performance of the workload is so critical that the cost of dedicated hardware is irrelevant.

- Specific hosts are running high-priority workloads (virtual machines), and you do not want to use the High Availability feature to prioritize these virtual machines.
- The hardware in the host is not the optimum for the other workloads in the pool.

Regardless of whether you specify a fixed or scheduled optimization mode, hosts excluded remain excluded even when the optimization mode changes. Therefore, if you only want to prevent Workload Balancing from powering down a host automatically, consider not enabling (or deselecting) Power Management for that host instead as described in Optimizing and Managing Power Automatically.

To exclude hosts from placement and optimization recommendations

1. Select the pool in the **Resources** pane, click on the **WLB** tab, and then click **Settings**.
2. In the left pane, select **Excluded Hosts**.
3. In the **Excluded Hosts** page, select the hosts for which you do not want Workload Balancing to recommend alternate placements and optimizations.

Advanced Settings

September 30, 2020

The settings in the **Advanced** dialog primarily fine-tune the way in which Workload Balancing applies recommendations when it is running in automated mode.

Important:

After Workload Balancing is running for a period of time, if you do not receive optimal placement recommendations, Citrix strongly recommends you evaluate your performance thresholds as described in the Workload Balancing Administrator's Guide. It is critical to set Workload Balancing to the correct thresholds for your environment or its recommendations might not be appropriate.

Overview: Setting Conservative or Aggressive Automated Recommendations

When running in automated mode, the frequency of optimization and consolidation recommendations and how soon they are automatically applied is a product of multiple factors, including:

- How long you specify Workload Balancing waits before applying another optimization recommendation
- The number of recommendations Workload Balancing must make before applying a recommendation automatically
- The severity level a recommendation must achieve before the optimization is applied automatically

- The level of consistency in recommendations (recommended virtual machines to move, destination hosts) Workload Balancing requires before applying recommendations automatically

VM Migration Interval

You can specify the number of minutes Workload Balancing must wait after the last time a particular virtual machine was moved, regardless of the cause, before it can generate another optimization recommendation that includes that particular virtual machine.

The recommendation interval is designed to prevent Workload Balancing from generating recommendations for artificial reasons (for example, if there was a temporary utilization spike).

When Automation is configured, it is especially important to be careful when modifying the recommendation interval. If an issue occurs that leads to continuous, recurring spikes, increasing the frequency (that is, setting a lower number) can generate many recommendations and, consequently, relocations.

> **Note:**
>
> Setting a recommendation interval does not affect how long Workload Balancing waits to factor recently rebalanced servers into recommendations for Start-On Placement, Resume, and Maintenance Mode.

Recommendation Count

Every two minutes, Workload Balancing checks to see if it can generate recommendations for the pool it is monitoring. When you enable Automation, you can specify the number of times a consistent recommendation must be made before Workload Balancing can automatically apply the recommendation. To do so, you configure a setting known as the Recommendation Count. The Recommendation Count and the Optimization Aggressiveness setting let you fine-tune the automated application of recommendations in your environment.

As described in the overview section , Workload Balancing uses the similarity of recommendations to a) vet if the recommendation is truly needed and b) determine if the destination host has stable enough performance over a prolonged period of time to accept a relocated virtual machine (without needing to move it off the host again shortly). Workload Balancing uses the Recommendation Count value to determine a recommendation must be repeated before Workload Balancing automatically applies the recommendation.

Workload Balancing uses this setting as follows:

1. Every time Workload Balancing generates a recommendation that meets its consistency requirements, as indicated by the Optimization Aggressiveness setting, Workload Balancing increments the Recommendation Count. If the recommendation does not meet the consistency

requirements, Workload Balancing may reset the Recommendation Count to zero, depending on the factors described in the *Workload Balancing Administrator's Guide*.

2. When Workload Balancing generates enough consistent recommendations to meet the value for the Recommendation Count, as specified in the Recommendations text box, it automatically applies the recommendation.

If you choose to modify this setting, the value you should set varies according to your environment. Consider these scenarios:

- If server loads and activity increase extremely quickly in your environment, you may want to increase value for the Recommendation Count. Workload Balancing generates recommendations every two minutes. For example, if you set this interval to "3", then six minutes later Workload Balancing applies the recommendation automatically.
- If server loads and activity increase gradually in your environment, you may want to decrease the value for the Recommendation Count.

Accepting recommendations uses system resources and affects performance when Workload Balancing is relocating the virtual machines. Increasing the Recommendation Count increases the number of matching recommendations that must occur before Workload Balancing applies the recommendation, which encourages Workload Balancing to apply more conservative, stable recommendations and may decrease the potential for spurious virtual machine moves. It should be noted, however, that the Recommendation Count is set to a conservative value by default.

Because of the potential impact adjusting this setting may have on your environment, Citrix only recommends changing it with extreme caution, preferably by testing and iteratively changing the value or under the guidance of Citrix Technical Support.

Recommendation Severity

All optimization recommendations include a severity rating (Critical, High, Medium, Low) that indicates the importance of the recommendation. Workload Balancing bases this rating on a combination of factors including configuration options you set, such as thresholds and metric tunings; resources available for the workload; and resource-usage history.

When you configure Workload Balancing to apply optimization recommendations automatically, you can set the minimum severity level that should be associated with a recommendation before Workload Balancing automatically applies it.

Optimization Aggressiveness

To provide additional assurance when running in automated mode, Workload Balancing has consistency criteria for accepting optimizations automatically so as to prevent moving virtual machines due

to spikes and anomalies. In automated mode, Workload Balancing does not accept the first recommendation it produces. Instead, Workload Balancing waits to automatically apply a recommendation until a host or virtual machine exhibits consistent behavior over time. The phrase consistent behavior over time refers to factors such as whether a host continues to trigger recommendations and whether the same virtual machines on that host continue to trigger recommendations.

Workload Balancing determines if behavior is consistent by using criteria for consistency and by having criteria for the number of times the same recommendation is made (that is, the Recommendation Count). You can configure how strictly you want Workload Balancing to apply the consistency criteria using a **Optimization Aggressiveness** setting.

While Citrix primarily designed the **Optimization Aggressiveness** setting for demonstration purposes, you can use this setting to control the amount of stability you want in your environment before Workload Balancing applies an optimization recommendation. The most stable setting (Low aggressiveness) is configured by default. In this context, the term stable refers to the similarity of the recommended changes over time, as explained throughout this section.

Workload Balancing uses up to four criteria to ascertain consistency. The number of criteria that must be met varies according to the level you set in the **Optimization Aggressiveness** setting. The lower the level (for example, Low or Medium) the less aggressively Workload Balancing is in accepting a recommendation. In other words, Workload Balancing is stricter about requiring criteria to match (or less cavalier or aggressive) about consistency when aggressiveness is set to Low.

For example, if the aggressiveness level is set to Low, Workload Balancing requires that each criterion for Low is met the number of times specified in the Recommendations box (where you specify the Recommendation Count value) before automatically applying the recommendation.

For example, if you set the Recommendation Count in the **Recommendations** box to "3", you are making Workload Balancing wait until it sees all the criteria listed in the *Workload Balancing Administrator's Guide* for Low are met and repeated in three consecutive recommendations. This helps ensure that the virtual machine actually needs to be moved and that the destination host Workload Balancing is recommending has consistently stable resource utilization over a longer period of time. It reduces the potential for a recently moved virtual machine to be moved off a host due to host performance changes after the move. By default, this setting is set to a conservative setting (Low) to encourage stability.

Citrix does not recommend increasing the **Optimization Aggressiveness** to increase the frequency with which your hosts are being optimized. If you feel that your hosts are not being optimized quickly or frequently enough, try adjusting the Critical thresholds, as described in Changing the Critical Thresholds.

For details about the consistency criteria associated with the different levels of aggressiveness, see the *Workload Balancing Administrator's Guide*.

If you find that Workload Balancing is not automatically applying optimization recommendations fre-

quently enough, you might want to increase the aggressiveness setting. However, Citrix strongly recommends reviewing the information in the *Workload Balancing Administrator's Guide* before doing so.

To configure virtual machine recommendation intervals

1. Select the pool in the **Infrastructure** view, click on the **WLB** tab, and then click **Settings**.

2. In the left pane, click **Advanced**.

3. In the **VM Migration Interval** section, do one or more of the following:

   - In the **Minutes to wait** box, type a value for the number of minutes you want Workload Balancing to wait before making another optimization recommendation on a newly rebalanced server.

   - In the **Recommendation Count** box, type a value for the number of optimization recommendations you want Workload Balancing to make before it applies an optimization recommendation automatically.

   - From the **Recommendation Severity** list, select a minimum severity level before optimizations are applied automatically.

   - From the **Optimization Aggressiveness** list, specify how aggressively Workload Balancing automatically applies optimization recommendations.

Pool Audit Trail Granularity

Workload Balancing enables you to specify the amount of data to be collected in the Pool Audit Trail report. This functionality also allows you to search and filter the audit trail logs by specific users, objects, and by time.

Pool Audit Trail Granularity is set to **Minimum** by default. This option captures limited amount of data for specific users and object types. You can modify the setting at any time based on the level of detail you would require in your report. For example, set the granularity to **Medium** for a user-friendly report of the audit log. If you require a detailed report, set the option to **Maximum**.

> **Important:**
>
> Setting the Pool Audit Trail Granularity to Maximum can cause the Workload Balancing server to use more disk space and memory. If you choose to set the granularity to Maximum, it is recommended that you carefully monitor the WLB server for disk space, memory usage, and CPU usage. If you think the WLB Server is under resource pressure, you should change the granularity setting to Medium or Minimum, or consider expanding your WLB server's memory or the size of the hard disk.

For more information, see Workload Balancing Report Glossary and Audit Log Events.

Administering Workload Balancing

September 30, 2020

Some administrative tasks you may want to perform on Workload Balancing include:

- Disconnecting from Workload Balancing
- Changing the Workload Balancing virtual appliance that a pool uses
- Changing the credentials Workload Balancing or XenServer use to communicate

You can also administer the Workload Balancing virtual appliance using the Workload Balancing service commands. These commands let you determine Workload Balancing virtual appliance status, change user accounts, and increase logging detail. For information about these commands, see the *Workload Balancing Administrator's Guide*.

> **Note:**
>
> For information about configuring Workload Balancing to use a different certificate or configuring XenServer to verify the identity of a certificate, see the *Workload Balancing Administrator's Guide*.

Disconnecting from Workload Balancing

September 30, 2020

If you want to stop Workload Balancing from monitoring your pool, you must disable Workload Balancing on the pool by disconnecting the Workload Balancing server.

When you disconnect a pool from the Workload Balancing virtual appliance, Workload Balancing permanently deletes information about the pool from the Workload Balancing database and stops collecting data for that pool. This means that if you want to use the same Workload Balancing virtual appliance to manage the pool again, you must re-enter the appliance's information in the **Connect to WLB Server** dialog box.

> **Important:**
>
> If you only want to stop Workload Balancing temporarily, click the **WLB** tab and click the **Pause** button.

To disconnect from Workload Balancing

1. In the **Resource** pane of XenCenter, select the resource pool on which you want to stop Workload Balancing.

---

2. From the **Pool** menu, select **Disconnect Workload Balancing Server**. The **Disconnect Workload Balancing** dialog box appears.

3. Click **Disconnect** to stop Workload Balancing from monitoring the pool.

> **Note:**
>
> If you disconnected the pool from the Workload Balancing virtual appliance, to re-enable Workload Balancing on that pool, you must reconnect to the appliance.

Reconfiguring a Pool to Use Another WLB Appliance

September 30, 2020

You can reconfigure a pool to use a different Workload Balancing virtual appliance.

However, to prevent the old Workload Balancing appliance from inadvertently remaining configured and collecting data for the pool, you must disconnect the pool from the old Workload Balancing appliance **before** connecting the pool to the new Workload Balancing appliance.

Once the pool is disconnected from the old Workload Balancing appliance, reconnect the pool by specifying the new Workload Balancing appliance name.

To use a different Workload Balancing appliance

1. On the pool you want to use a different Workload Balancing appliance, from the **Pool** menu, select **Disconnect Workload Balancing Server** and click **Disconnect** when prompted. For instructions, see Disconnecting from Workload Balancing.

2. In the **WLB** tab, click **Connect**. The **Connect to WLB Server** dialog appears.

3. In the **Address** box, type the IP address or host name (FQDN) name of the new Workload Balancing appliance.

   If the new Workload Balancing appliance uses different credentials, you must also enter the new credentials.

   > **Note:**
   >
   > You must enter all the information that you would normally enter when you initially connect a pool to Workload Balancing. For information, see Connecting to Workload Balancing.

Updating Workload Balancing credentials

September 30, 2020

After initial configuration, if you need to update the credentials XenServer and the Workload Balancing appliance use to communicate, you can do so using a three-step process:

1. Disconnect from Workload Balancing, as described below.
2. Change the WLB credentials by editing the WlbConfig file (run the WlbConfig command in the console on the Workload Balancing virtual appliance). See the *Workload Balancing Administrator's Guide* for details.
3. Reenable Workload Balancing and specify the new credentials, as described below.

Situations when you may want to use these steps include:

- If you need to change the user account XenServer uses to communicate with Workload Balancing
- If you receive an error message that the Workload Balancing credentials are no longer valid
- If the service is unavailable

If you want to modify settings for thresholds and change the priority given to specific resources, see Editing Workload Balancing Settings.

To disconnect from Workload Balancing

1. In the **Resource** pane of XenCenter, select the resource pool on which you want to stop Workload Balancing.
2. From the **Pool** menu, select **Disconnect Workload Balancing Server**. The **Disconnect Workload Balancing** dialog box appears.
3. Click **Disconnect** to permanently stop Workload Balancing from monitoring the pool.

To reenable Workload Balancing and specify the new credentials

1. After the progress bar completes, click **Connect**. The Connect to WLB Server dialog box appears.
2. Click Update **Credentials**.
3. In the Server Address section, modify the following as desired:
   - In the **Address** box, type the IP address or FQDN of the Workload Balancing appliance.
   - (Optional.) If you changed the port number during Workload Balancing Configuration, enter that port number. The port number you specify in this box and during Workload Balancing Configuration is the port number XenServer uses to connect to Workload Balancing. By default, XenServer connects to Workload Balancing on port 8012.

> **Note:**
>
> Do not edit this port number unless you have changed it during Workload Balancing Setup. The port number value specified during Setup and in the **Workload Balancing Configuration** dialog must match.

4. In the **WLB Server Credentials** section, enter the user name (for example, wlbuser) and password the computers running XenServer will use to connect to the Workload Balancing server.

5. In the **XenServer Credentials** section, enter the user name and password for the pool you are configuring (typically the password for the pool master). Workload Balancing will use these credentials to connect to the computers running XenServer in that pool. To use the credentials with which you are currently logged into XenServer, select the **Use the current XenCenter credentials** check box.

Entering Maintenance Mode with Workload Balancing Enabled

September 30, 2020

When Workload Balancing is enabled, if you take a physical host offline for maintenance (that is, suspend a server by entering Maintenance Mode), XenServer automatically migrates the virtual machines running on that host to their optimal servers when available. XenServer migrates them based on Workload Balancing recommendations (performance data, your placement strategy, and performance thresholds).

If an optimal server is not available, the words **Click here to suspend the VM** appear in the **Enter Maintenance Mode** dialog box. In this case, Workload Balancing does not recommend a placement because no host has sufficient resources to run this virtual machine. You can either suspend this virtual machine or exit Maintenance Mode and suspend a virtual machine on another host in the same pool. Then, if you reenter the **Enter Maintenance Mode** dialog box, Workload Balancing might be able to list a host that is a suitable candidate for migration.

> **Note:**
>
> When you take a server offline for maintenance and Workload Balancing is enabled, the words "Workload Balancing" appear in the upper-right corner of the **Enter Maintenance Mode** dialog box.

To enter maintenance mode with Workload Balancing enabled

1. In the Resources pane, select the server and then do one of the following:
   - Right-click and click **Enter Maintenance Mode** on the shortcut menu.
   - On the **Server** menu, click **Enter Maintenance Mode**.

2. Click **Enter Maintenance Mode**. The virtual machines running on the server are automatically migrated to the optimal host based on Workload Balancing's performance data, your placement strategy, and performance thresholds.

To take the server out of maintenance mode

1. In the **Resources** pane, select the server and then do one of the following:

   - Right-click and click **Exit Maintenance Mode** on the shortcut menu.
   - On the **Server** menu, click **Exit Maintenance Mode**.

2. Click **Exit Maintenance Mode**.

   When you remove a server from maintenance mode, XenServer automatically restores that server's original virtual machines to that server.

Troubleshooting Workload Balancing

June 3, 2019

While Workload Balancing usually runs smoothly, this help system provides a series of topics with guidance in case you encounter issues. Additional troubleshooting topics are provided in the *Workload Balancing Administrator's Guide*.

Here are a few tips for resolving general Workload Balancing issues:

General Troubleshooting Tips

Start troubleshooting by reviewing the Workload Balancing log. You can find the log in Workload Balancing appliance in this location (by default):

```
1  /var/log/wlb
2  <!--NeedCopy-->
```

Also, you can also see the event logs in the XenCenter Navigation pane, click **Notifications** and then **Events** for more information.

Error Messages

Workload Balancing displays error messages in the **Alerts** view in XenCenter and, in some cases, on screen as dialog boxes.

Issues Entering Workload Balancing Credentials

June 3, 2019

If you cannot get Workload Balancing to accept the appliance user account and password when configuring the **Connect to WLB Server** dialog, try the following:

- Make sure that Workload Balancing appliance imported and was configured correctly and all of its services are running by running the service workloadbalancing start command.
- Using Issues Starting Workload Balancing as a guide, check to make sure you are entering the correct credentials.
- Enter the Workload Balancing server's IP address if you are having trouble entering the Workload Balancing FQDN.

You can enter the host name of the Workload Balancing appliance in the **Address** box, but it must be a fully qualified domain name (FQDN). For example, yourcomputername.yourdomain.net .

Issues Starting Workload Balancing

June 3, 2019

If after importing and configuring the Workload Balancing appliance, you receive an error message that XenServer and Workload Balancing cannot connect to each other, you might have entered the incorrect credentials in the **Connect to WLB Server** dialog. To isolate this issue, try:

- Verifying the credentials you entered in the **Connect to WLB Server** dialog match the credentials that you created on the Workload Balancing server and on XenServer
- Verifying the IP address or FQDN of the Workload Balancing appliance you entered in the **Connect to WLB Server** dialog is correct.
- Verifying the account credentials for the Workload Balancing account you created during Workload Balancing Configuration match the credentials you entered in the **Connect to WLB Server** dialog.

Workload Balancing Connection Errors

June 3, 2019

If you receive a connection error in the Workload Balancing Status line on the **WLB** tab, you might need to reconfigure Workload Balancing on that resource pool.

Click the **Connect** button on the **WLB** tab and enter the server credentials again.

Typical causes for this error include changing the credentials of the WLB virtual appliance or pool master or changing the name of the WLB virtual appliance.

Issues changing Workload Balancing servers

June 3, 2019

If you connect a resource pool to a different Workload Balancing virtual appliance without first disconnecting the pool from the original Workload Balancing appliance, both the old and new Workload Balancing appliances will monitor the pool.

To solve this problem, you can either shut down and delete the old Workload Balancing appliance or manually stop the Workload Balancing services (analysis, data collector and Web service) so that the appliance no longer monitors the pool.

Citrix does not recommend using the pool-initialize-wlb xe command to remove or change Workload Balancing server configuration.

Monitoring System Performance

June 3, 2019

Topics

- About Performance Monitoring
- Viewing Performance Data
- Configuring Performance Graphs
- Configuring Performance Alerts

About Performance Monitoring

September 30, 2020

The **Performance** tab in XenCenter provides real time monitoring of performance statistics across resource pools as well as graphical trending of virtual and physical machine performance.

- You can view up to 12 months of performance data and zoom in to take a closer look at activity spikes. To learn more, see Viewing performance data.

- By default, graphs showing CPU, memory, network I/O and disk I/O are displayed on the tab. However, you can add more performance data and change the appearance of the graphs. To learn more, see Configuring performance graphs
- Performance alerts can be generated when CPU, memory usage, network, storage throughput, or VM disk activity go over a specified threshold on a managed server, virtual machine, or storage repository. For more information, see Configuring performance alerts.

Note:

Full performance data is only available for virtual machines with the XenServer paravirtualized drivers (XenServer Tools) installed.

Viewing Performance Data

June 3, 2019

The **Performance** tab shows performance data for the selected server or virtual machine in graph form.

For servers you can view:

- CPU, memory and network I/O usage data, and you can add graphs showing additional resource usage data, if required. For example, you can include the Control Domain Load: this is the average (Linux loadavg) of the number of processes queued inside the XenServer Control Domain (dom0) over the last 5 minutes.
- Lifecycle events for all the VMs hosted on the server are shown in the **VM Lifecycle Events** pane.

For virtual machines, graphs showing CPU, memory, network I/O and disk usage data are shown by default.

At the bottom of the tab, the summary graph gives a quick overview of what has been happening on the machine and also allows you to adjust the time frame shown in the other graphs, either to show data from a longer or shorter period of time, or to show data from an earlier period of time.

To include other types of performance data on the tab or to change the appearance of the graphs, see Configuring performance graphs.

To view data from a longer or shorter time period

By default, data from the last 10 minutes is displayed. To view data from a longer or shorter time period, do one of the following:

- To view available performance data for the last hour, 24 hours, week month or year, click **Zoom** and then select **1 Hour** , **1 Day** , **1 Week** , **1 Month** or **1 Year**.

- To resize the time period displayed in the graphs, in the summary graph, point to the vertical split bar at the edge of the sample area. When the pointer changes to a double-headed arrow ↔ , drag the vertical split bar right or left. For example:
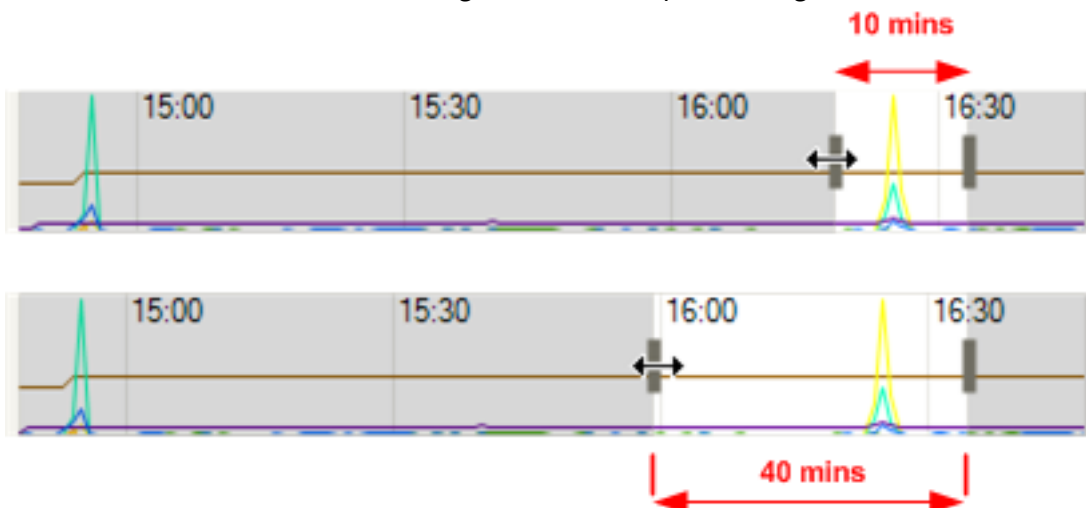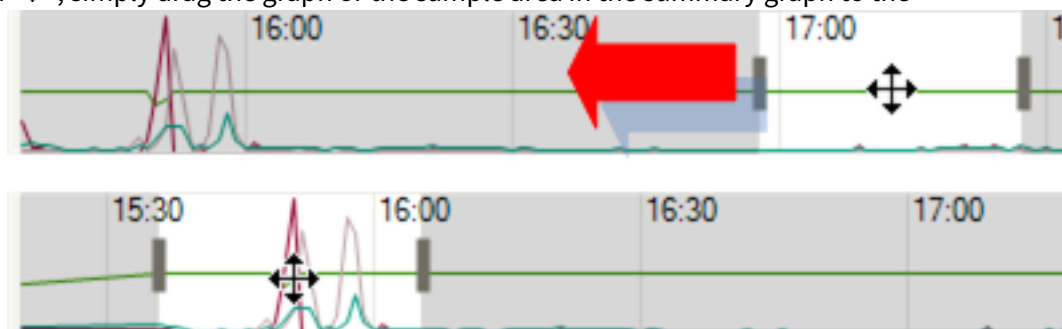
To view data from a different time period

To move the time frame for data displayed in the graphs, point to any graph and when the pointer changes to a move cursor ✛ , simply drag the graph or the sample area in the summary graph to the left or right. For example:

To view VM lifecycle event data on a server

To view lifecycle events for the VMs hosted on a server, use the **VM Lifecycle Events** list.

- Each event has a tooltip with the full message for that lifecycle event ("Virtual Machine 'Sierra' has been started").
- You can use the cursor keys to navigate the items in list.
- Double clicking or pressing **Enter** will zoom the graphs to the point when the selected lifecycle event occurred.
- Selecting (single click or highlight with cursor keys) one of the events causes the lifecycle event on the graph itself to be highlighted.

Configuring Performance Graphs

June 3, 2019

To Add a New Graph

1. On the **Performance** tab, click **Actions** and then **New Graph**. The **New Graph** dialog box will be displayed.
2. Enter a name for the graph in the **Name** field.
3. From the list of datasources, select the check boxes for the datasources you want to include in the graph.
4. Click **Save**.

To Edit a Graph

1. Navigate to the **Performance** tab and select the graph that you would like to edit.
2. Click **Actions** and then **Edit Graph**.
3. On the graph details window, make the necessary changes and click **OK**.

To Delete a Graph

1. Select the graph that you would like to remove from the list of graphs displayed on the **Performance** tab.
2. Click **Actions** and then **Delete Graph**.
3. Click **Yes** to confirm the deletion.

To Reorder a Graph

1. Navigate to the **Performance** tab and select the graph that you would like to reorder.
2. Click the **Move Up** or **Move Down** tab to move the graph from its current location.

To Change Datasource Color in Graphs

1. Navigate to the **Performance** tab.
2. Double click on the graph for which you would like to change the color of the datasource. The graph details dialog box will be displayed.
3. Click the colored check box located against the required datasource and select a new color from the color picker.
4. Click **OK** to confirm.

To Change the Graph Type

Data on the performance graphs can be displayed as lines or as areas:

Line Graph:
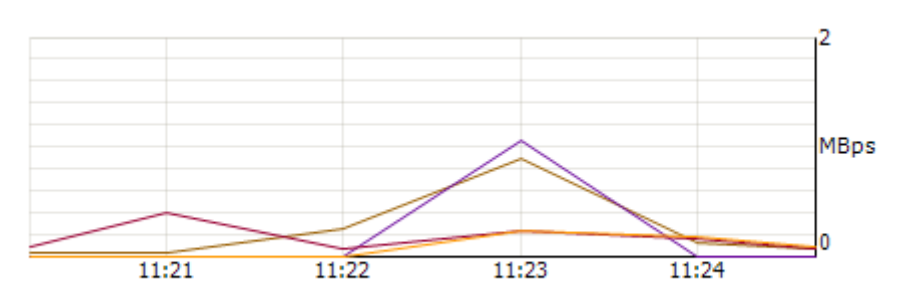
Area Graph:
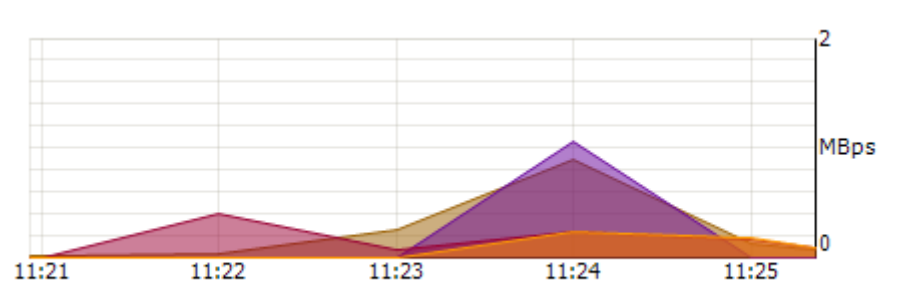
To change the graph type:

1. On the **Tools** menu, click **Options** and then click on the **Graphs** tab.
2. To view performance data as a line graph, click on the **Line graph** radio button.
3. To view performance data as an area graph, click on the **Area graph** radio button.
4. Click **OK** to save your changes.

Configuring Performance Alerts

September 30, 2020

Performance alerts can be generated when CPU, memory usage, network, storage throughput, or VM disk activity exceeds a specified threshold on a managed server, virtual machine, or storage repository. By default, the alert repeat interval is set to 60 minutes, and it can be modified if required.

Performance alerts will appear in the **Alerts** view (accessed by clicking on the **Notifications** button on the left hand pane). In addition, you can have performance alerts emailed to you along with other alerts: see XenCenter Alerts.

To configure performance alerts:

1. Select the server, virtual machine, or storage repository in the **Resources** pane, click on the **General** tab and then click **Properties**.
2. Click on the **Alerts** tab and then:
   - To request CPU performance alerts for a server or virtual machine, select the **Generate CPU usage alerts** check box, then set the CPU usage and time threshold that will trigger the alert.
   - To request network performance alerts for a server or virtual machine, select the **Generate network usage alerts** check box, then set the network usage and time threshold that will trigger the alert.
   - To request memory performance alerts for a server, select the **Generate memory usage alerts** check box, and then set the memory usage and time threshold that will trigger the alert.
   - To request disk usage performance alerts for a virtual machine, select the **Generate disk usage alerts** check box, then set the disk usage and time threshold that will trigger the alert.
   - To request storage throughput alerts for a storage repository, select the **Generate storage throughput alerts** check box, then set the storage throughput and time threshold that will trigger the alert.

     Note:

     This option generates alerts when the total read/write storage throughput activity on a Physical Block Device (PBD) exceeds the specified limit. PBD represents the interface between a specific XenServer host and an attached SR. When the total read/write SR throughput activity on a PBD exceeds the threshold you have specified, alerts will be generated on the host connected to the PBD. Unlike other host alerts, this must be configured on the relevant SR.
   - To change the alert repeat interval, enter the number of minutes in the **Alert repeat interval** box. Once an alert threshold has been reached and an alert generated, another alert will not be generated until after the alert repeat interval has elapsed.
3. Click **OK** to save your changes.

Updates and Upgrades

June 3, 2019

Topics

- About Software Updates and Upgrades
- Upgrading Managed Servers - Rolling Pool Upgrade Wizard

- Updating Managed Servers - Install Update Wizard
- Live Patching
- Applying Automated Updates
- Installing Supplemental Packs
- Updating XenCenter
- Update Notification

About Software Updates and Upgrades

June 3, 2019

XenCenter issues notifications about available XenServer and XenCenter updates and upgrades on the **Updates** tab in the **Notifications** view. A software upgrade is a major, standalone version of XenServer. Upgrades offer new features or technologies, or replace a release that is no longer supported. A software update updates a major version of XenServer, but does not upgrade it to the next major version.

XenCenter is configured by default to automatically check for new XenServer and XenCenter updates and upgrades at regular intervals, and you will be notified when a new update or product version is available. It is recommended that you install all published updates. You can check for available updates manually at any time and verify that you are running the latest version of both XenServer and XenCenter. To do this, select **Notifications** > **Updates** and then click **Refresh**.

Applying updates to XenServer hosts

Updates to a version of XenServer can be delivered as a Hotfix or a Cumulative Update. Hotfixes generally supply bug fixes to one or more specific issues. Cumulative Updates contain accumulated bug fixes, and occasionally, feature improvements and enhancements. Updates can be quickly applied to your managed servers. For more information, see Updating managed servers.

New Current Releases of XenServer are also delivered as updates. You can apply a Current Release as an update to some previous Current Releases of XenServer. This update moves you to a newer version of XenServer.

Updating your version of XenCenter

Updates for XenCenter are supplied as XenCenterSetup.exe files which can be downloaded and installed using a Download Manager; see Updating XenCenter.

Upgrading XenServer hosts

To upgrade XenServer hosts, use the **Rolling Pool Upgrade** wizard. You can use this wizard to upgrade multiple servers in a pool with minimal service interruption for running VMs. VMs are automatically migrated onto other available servers as the upgrade is applied to each server in turn. The wizard can also be used to upgrade standalone servers. See Upgrading managed servers.

Upgrading Managed Servers

September 30, 2020

You can use the **Rolling Pool Upgrade** wizard to upgrade XenServer - standalone servers or a pool of servers to a newer version.

> **Note:**
>
> The Rolling Pool Upgrade wizard is available for licensed XenServer customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. For more information, see About XenServer Licensing.

The Rolling Pool Upgrade wizard guides you through the upgrade procedure and organizes the upgrade path automatically. The Rolling Pool Upgrade wizard allows you to upgrade multiple servers and pools simultaneously. Each of the servers in the pool is upgraded in turn, starting with the pool master. Before starting an upgrade, the wizard conducts a series of prechecks to ensure that certain pool-wide features, such as HA and WLB, are temporarily disabled and that each host in the pool is prepared for upgrade (for example, that the CD/DVD drive of each host is empty). Only one server is offline at a time, and any running VMs are automatically migrated off each server before the upgrade is installed on that server.

The wizard can operate in manual or automatic mode:

- In manual mode, you must manually run the XenServer installer on each server in turn and follow the on-screen instructions on the serial console of the server. When the upgrade begins, XenCenter prompts you to insert XenServer installation media or specify a PXE boot server for each server that you upgrade.
- In automatic mode, the wizard uses network installation files located on an HTTP, NFS or FTP server to upgrade each server in turn, without requiring you to insert install media, manually reboot, or step through the installer on each server. If you choose to perform a rolling pool upgrade in this manner, unpack the XenServer installation media onto your HTTP, NFS or FTP server before starting the upgrade.

You can also use the **Rolling Pool Upgrade** wizard to upgrade standalone server s, that is, server s which do not belong to any resource pool.

Important: Before you upgrade

Upgrading a pool of server s requires careful planning. As you plan your upgrade, it is very important to be aware of the following:

- Download and install the latest version of XenCenter. For example, when you are upgrading your hosts to XenServer 7.1 CU2, you must use XenCenter issued with XenServer 7.1 CU2. Using earlier versions of XenCenter to upgrade to a newer version of XenServer is not supported.
- VMs can only be migrated from a server running an older version of XenServer to one running the same version or higher. You cannot migrate VMs from an upgraded server to one running an older version of XenServer. Ensure to allow for space on your servers accordingly.

Citrix strongly advises against running a mixed-mode pool (one with multiple versions of XenServer coexisting) for longer than necessary, as the pool operates in a degraded state during upgrade.

Do not attempt to perform any key control operations as these are not available during upgrade. Though VMs continue to function as normal, VM actions other than migrate may not be available (for example, shut down, copy, and export). In particular, it is not safe to perform storage-related operations such as adding, removing or resizing virtual disks.

The wizard will always upgrade the pool master first. Do not place the pool master into maintenance mode using XenCenter before performing the upgrade as this will cause a new master to be designated.

Take a backup of the state of your existing pool using the pool-dumpdatabase xe CLI command (see the *XenServer Administrator's Guide* for details). This allows you to revert a partially complete rolling upgrade back to its original state without losing any VM data. As it is not possible to migrate a VM from an upgraded server to a server running an older version of XenServer, it may be necessary to shut down VMs if you need to revert the rolling upgrade for any reason.

Ensure that your servers are not over-provisioned - that is, they have sufficient memory to carry out the upgrade. It is best to suspend any VMs that are not critical during the upgrade process.

While the **Rolling Pool Upgrade** wizard checks that the following actions have been taken, you may choose to perform them before you begin the upgrade:

- Empty the CD/DVD drives of the VMs in the pool(s). For details and instructions, see the *XenServer Installation Guide*.
- Disable HA.
- Disable WLB

To upgrade XenServer using the Rolling Pool Upgrade wizard

1. Open the Rolling Pool Upgrade wizard: on the **Tools** menu, select **Rolling Pool Upgrade**.

2. Read the Before You Start information, and then click **Next** to continue.

3. Select the pool or standalone servers that you would like to upgrade and then click **Next**.

4. Choose the Upgrade Mode. You can either select **Automatic Mode** or **Manual Mode** depending on whether you are planning an automated upgrade from network installation files located on an HTTP, NFS, or FTP server, or a manual upgrade from either a CD/DVD or a server through PXE boot.

   > **Note:**
   >
   > If you choose **Manual Mode** , you must run XenServer installer on each server in turn and follow the on-screen instructions on the serial console of the server. Once the upgrade begins, XenCenter prompts you to insert the XenServer installation media or specify a PXE boot server for each server that you upgrade.

5. Choose whether you want XenCenter to automatically download and install the minimal set of updates (hotfixes) after upgrading the servers to a newer version. The apply updates option is selected by default. However, you must have internet connection to download and install the updates.

6. After selecting the Upgrade Mode, click **Run Prechecks**.

7. Follow the on-screen recommendations to resolve any upgrade prechecks that have failed. If you want XenCenter to automatically resolve all failed prechecks, click **Resolve All**. When the prechecks have been resolved, click **Next** to continue.

Prepare the XenServer installation media:

- If you chose **Automatic Mode** , enter the installation media details. Choose **HTTP** , **NFS** or **FTP** and then specify the path, username and password, as appropriate.
- If you chose **Manual Mode** , note the upgrade plan and instructions.

Click **Start Upgrade** to begin the upgrade.

When the upgrade begins, the wizard guides you through any actions you need to take to upgrade each server. Follow the instructions until you have upgraded and updated all servers in the pool.

> **Note:**
>
> If the upgrade or the update process fails for any reason, the Rolling Pool Upgrade wizard halts the process. This allows you to fix the issue and resume the upgrade process by clicking the **Retry** button.

The Rolling Pool Upgrade wizard prints a summary when the upgrade is complete. Click **Finish** to close the wizard.

Updating Managed Servers

September 30, 2020

XenCenter issues notifications about available XenServer updates in the **Updates** tab on the **Notifications** view. Updates to XenServer can be delivered as one of the following types of update:

- **Hotfixes** , which contain bug fixes for one or more specific issues. Hotfixes are provided for XenServer releases in the Long Term Service Release (LTSR) and Current Release (CR) streams and for earlier supported releases that are not part of either stream

  Hotfixes on the latest CR is available to all XenServer customers. However, hotfixes on previous CRs that are still in support are only available for customers with an active Citrix Customer Success Services (CSS) account.

  Hotfixes on the LTSR stream are available to customers with an active CSS account. For more information, see the *XenServer Licensing FAQ*

- **Cumulative Updates** , which contain previously released hotfixes and may contain support for new guests and hardware. Cumulative Updates are provided for XenServer releases in the LTSR stream and are available only to customers with an active Citrix Customer Success Services account.

- **Current Releases** , which are full versions of XenServer from the Curent Release (CR) stream. Current Releases are available as updates that can be applied to a previous Current Release of XenServer.

This topic contains information about applying XenServer updates to your managed servers. You must pay careful attention to the release note that is published with each update. Each update might have unique installation instructions, particularly with regard to preparatory and post-update operations. Some updates might only be available to licensed users or to Customer Success Services customers.

In addition to XenServer updates, the Updates tab also notifies users about the availability of new versions of XenServer and new versions of XenCenter. Some new XenServer releases from the CR stream can be applied as updates to earlier versions of XenServer from the CR stream. However, many new releases can only be reached by the upgrade process. For information about upgrading XenServer, see Upgrading Managed Servers. To update your XenCenter to a newer version, see Updating XenCenter. For information about installing supplemental packs, see Installing Supplemental Packs.

XenCenter also enables you to dismiss updates listed on the **Updates** tab. Dismissing an update hides the update entry from the list. Select any unwanted updates from the list and click **Dismiss** and then **Dismiss Selected**. To dismiss all the updates, select **Dismiss All**. If you would like to see updates which were previously dismissed, click **Restore Dismissed Updates**.

Before you Update

Before you apply an update to your servers, pay careful attention to the following:

1. Citrix strongly recommends that you read the release notes published with each update.
2. Back up your data before applying an update, just as you would with any other maintenance operations. For backup procedures, see the *XenServer 7.1 Administrator's Guide.*
3. We recommend that you reboot all your servers before installing an update and then verify their configuration, for example, to check that the VMs start and that storage is accessible. This is because some configuration changes will only take effect when a server is rebooted, so the reboot may uncover configuration problems that could cause the update to fail.
4. When you are upgrading a pool of servers to a newer version, you must upgrade each server in a pool starting with the pool master and ensure that the pool is up and running **before** applying any updates.
5. Update ALL servers in a pool within a short period: running a mixed-mode pool (a pool that includes updated and non-updated servers) is not a supported configuration. This should be scheduled to minimize the amount of time that a pool runs in a mixed state.
6. Update all servers within a pool sequentially, always starting with the pool master.
7. After applying an update to all servers in a pool, update any required driver disks before rebooting the servers.

Viewing Available Updates

The **Updates** section of the **Notifications** view lists the updates that are available for all connected servers and pools.

Notes:

- By default, XenCenter periodically checks for XenServer and XenCenter updates. Click **Refresh** to manually check for available updates.
- If the **Updates** tab cannot find any updates because you have disabled automatic check for updates, a message appears on the Updates tab. Click **Check for Updates Now** to manually check for updates.

You can select from the **View** dropdown whether to view the list of updates **By Update** or **By Server**.

When you view the list of updates **By Update** , XenCenter displays the list of updates, which can be ordered by **Server / Pool** or by **Date**.

- Cumulative Updates and new releases are displayed at the top of this list. Not all new releases can be applied as an update.
- To export this information as a .csv file, click **Export All**. The .csv file lists the update name, a description of the update, the servers that this update can be applied to, the timestamp of the update, and a reference to the web page that the update is downloaded from.

- To apply an update to a server, from the **Actions** drop-down list for that update select **Download and Install**. This extracts the update and opens the **Install Update** wizard on the **Select Servers** page with the relevant servers selected. For more information, see the following section Updating a Pool Automatically.
- To open the release note of an update in your browser, click the **Actions** drop-down list and select **Go to Web Page**.

When you view the list of updates **By Server**, XenCenter displays the list of servers connected to XenCenter. This list shows both the updates that can be applied to the servers and the updates that are installed on the servers.

- To export this information as a .csv file, click **Export All**. The .csv file lists the **Pool** that the server belongs to, the **Server** name, the **Status** of the installed XenServer, the update **Status** of the server, the **Required Updates** for this server, and the **Installed Updates** for this server.
- To apply the updates, click **Install Updates**. This opens the **Install Update** wizard on the **Select Update** page. For more information, see the following section Updating a Pool Automatically.

Updating a Pool Automatically

XenCenter allows you to apply automated updates that are needed to bring your servers up-to-date. You can apply these updates to one or more pools. When you choose to apply automated updates, XenCenter applies the minimum set of updates that are required to bring the selected pool or the standalone server up-to-date. XenCenter minimizes the number of reboots required to bring the pool or the standalone server up-to-date, and where possible, limits it to a single reboot at the end. For more information, see Applying Automated Updates.

Applying an Update to your Managed Servers

The update installation mechanism in XenCenter allows you to download and extract the selected update from the Citrix Support website, and enables you to apply an update to multiple servers and pools simultaneously using the **Install Update** wizard. During the process, the **Install Update** wizard automatically migrates VMs off each server, places the server in Maintenance mode, applies the update, reboots the server if required, and then, migrate the VMs back to the updated server. Any actions that were taken at the pre-check stage to enable the updates to be applied, such as turning off HA, will be reverted.

When you choose to install a Current Release or an LTSR, the Install Update mechanism offers to apply the minimum set of updates (hotfixes) on the new version to bring the servers up to date.

The following section provides step-by-step instructions on extracting and applying an update using the **Install Update** wizard. If you are planning to apply an update that you have already downloaded from the Citrix Support website, see the section *Installing previously downloaded updates*.

1. From the XenCenter menu, select **Tools** and then **Install Update**.

2. Review the information on the **Before You Start** page and click **Next** to continue.

3. Select the update to install and click **Next** to continue.

4. Select the servers to install the update on and click **Next** to continue.

   Note:

   - If you are installing a Current Release or an LTSR, XenCenter also offers to apply the minimum set of updates (hotfixes) after the installation process.
   - If you are installing an update (hotfix), XenCenter downloads and extracts the update, and uploads it to the servers you have specified. The Upload page displays the status of the upload.

5. The **Install Update** wizard performs a number of update prechecks, including the space available on the servers, to verify that the update can be applied on the selected servers and displays the result. The wizard also checks whether the servers need to be rebooted after the update is applied and displays the result. In addition, the Install Update wizard checks whether a live patch is available for the update and whether the live patch can be successfully applied to the servers. For information about Live Patching, see [Live Patching in XenServer](#).

6. Follow the on-screen recommendations to resolve any update prechecks that have failed. If you prefer XenCenter to automatically resolve all failed prechecks, click **Resolve All**. When the prechecks have been resolved, click **Next** to continue.

7. If you are installing a Current Release or an LTSR, XenCenter downloads the updates, uploads them to the default SR of the pool, and installs the updates. The **Upload and Install** page displays the progress.

   Notes:

   - If the default SR in a pool is not shared, or does not have enough space, XenCenter uploads the update to another shared SR with sufficient space. If none of the shared SRs have sufficient space, the update is uploaded to local storage of the pool master.
   - If the update process cannot complete for any reason, XenCenter halts the process. This allows you to fix the issue and resume the update process by clicking the **Retry** option.

   See Step 10 to complete the Current Release or the LTSR installation process.

8. If you are installing an update (hotfix), choose an **Update Mode**. Review the information displayed on the screen and select an update mode. If the update contains a live patch that can be successfully applied to the servers, it displays **No action required** on the **Update Mode** page.

   > **Note:**
   >
   > If you click **Cancel** at this stage, the **Install Update** wizard reverts the changes and removes

> the update file from the server.

9. Click **Install update** to proceed with the installation. The Install Update wizard shows the progress of the update, displaying the major operations that XenCenter performs while updating each server in the pool.

10. Click **Finish** to close the Install Update wizard. If you chose to carry out the post-update tasks, do so now.

Installing previously downloaded updates

XenCenter enables you to install updates that you have already downloaded. Update files are delivered as zip files on the Citrix Support website.

> **Note:**
>
> Updates to XenServer are published as .zip files. Download the zip file to a known location on your computer and make sure to extract the .iso file before installing the update.

1. From the XenCenter menu, select **Tools** and then **Install Update**.
2. Read the information displayed on the **Before You Start** page and then click **Next**.
3. On the **Select Update** page, click **Browse** to locate the update file and then click **Open**. Click **Next** to continue.
4. Select the pool and servers that you would like to update. Any servers or pools that cannot be updated will be grayed out. Click **Next** to continue.
5. Follow the instructions from step 3 to 9 in the section *Applying an Update to your Managed Servers* above to complete the update installation process.
6. Click **Finish** to exit the wizard.

Live Patching in XenServer

September 30, 2020

XenServer customers who deploy XenServer hosts may often need to reboot their servers after applying hotfixes. This rebooting results in unwanted downtime for the servers while customers have to wait till the system is restarted. This also reduces the uptime of the servers and impacts business. Live patching enables customers to install some Linux kernel and Xen hypervisor updates without having to reboot the servers. This reduces maintenance costs and downtime. Such hotfixes will consist of both a live patch, which will be applied to the memory of the server, as well as a hotfix that updates the files on disk.

Live Patching is enabled by default. For more information about enabling and disabling Live Patching, see Change Pool Properties.

When applying an update using the **Install Update** wizard, the **Prechecks** page displays information about the post-update tasks that customers are required to carry out for the update to take effect. In addition, the wizard also checks whether the servers need to be rebooted after the update is applied and displays the result. This enables customers to know the post-update tasks well in advance and schedule the application of updates accordingly.

> **Note:**
>
> XenServer Live Patching is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. For more information about licensing, see About XenServer Licensing.

Live Patching Scenarios

Hotfixes can be live patched across pools, servers, or on a standalone server. Some updates may need reboot, some need the xapi toolstack to be restarted, and a few updates do not have any post-update tasks.

The following scenarios describe the behavior when a Live Patch is and is not available for an update:

- **Updates with live patch** Hotfixes that update Linux kernel and the Xen hypervisor usually do not need a reboot after applying the update. However, in some rare cases, when the live patch cannot be applied, a reboot may be required.
- **Updates without live patch** -No change in the behavior here. It works as usual.

> **Note:**
>
> If a server does not require a reboot, or if the update contains live patches that can be successfully applied to the servers, XenCenter displays `No action required` on the **Update Mode** page.

Applying Automated Updates

September 30, 2020

XenCenter allows you to apply automated updates that are needed to bring your servers up-to-date. You can apply these updates to one or more pools. When you choose to apply automated updates, XenCenter applies the minimum set of updates that are required to bring the selected pool or the standalone server up-to-date. XenCenter minimizes the number of reboots required to bring the pool or the standalone server up-to-date, and where possible, limits it to a single reboot at the end.

As a prerequisite, XenCenter requires Internet access to fetch the required updates. When you choose to apply automated updates, all the required updates are applied. The automated updates feature

---

applies any Cumulative Updates that are available for a host. However, if a new Current Release version is available as an update, the automated updates feature does not apply this update. You must manually select to update your servers to the new Current Release.

To view the list of required updates, perform the following steps:

1. Select the server from the Resources pane.
2. Navigate to the **General** tab.
3. Expand the **Updates** section. You can see:
   - **Applied** - lists already applied updates.
   - **Required Updates** - lists the set of updates required to bring the server up-to-date.
     > **Note:**
     >
     > If there are no updates required, the Required Updates section will not be displayed.
   - **Installed supplemental packs** - lists supplemental pack that are installed on the server (if any).
     > **Note:**
     >
     > If you select a pool instead of a server, the **Updates** section on the **General** tab lists the updates that are already applied as **Fully applied**.

If you would like to install specific updates to a pool or a managed server, see Applying updates to your managed servers.

> **Note:**
>
> The Automated Updates feature is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement. For more information about licensing, see About XenServer Licensing.

The following section provides step-by-step instructions on how to apply automated updates using the **Install Update** wizard.

1. From the XenCenter menu, select **Tools** and then select **Install Update**.

2. Read the information displayed on the **Before You Start** page and then click **Next**.

3. Select **Automated Updates**. his option is visible only if XenCenter is connected to at least one licensed pool or a licensed standalone server.

4. Click **Next**.

5. Select one or more pools or standalone servers that you would like to update and click **Next**. Any server or pool that cannot be updated appears grayed out.

6. The **Install Update** wizard performs a number of update prechecks, including the free space check on the servers.

   Follow the on-screen recommendations to resolve any prechecks that are failed. If you prefer XenCenter to automatically resolve all failed prechecks, click **Resolve All**.

7. When the prechecks have been resolved, click **Next** to continue.

   The **Install Update** wizard automatically downloads and installs the recommended updates. The wizard also shows the overall progress of the update, displaying the major operations that XenCenter performs while updating each server in the pool.

   Notes:

   - The updates will be uploaded to the default SR of the pool. If the default SR is not shared, or does not have enough space, XenCenter tries to upload the update to another shared SR with sufficient space. If none of the shared SRs have sufficient space, the update will be uploaded to local storage on the pool master.
   - If the update process cannot complete for any reason, XenCenter halts the process. This allows you to fix the issue and resume the update process by clicking the **Retry** button.

8. When the updates have been applied, click **Finish** to close the **Install Update** wizard.

Installing Supplemental Packs

September 30, 2020

Supplemental packs are used to modify and extend the functionality of XenServer, by installing software into the control domain (Dom0). For example, an OEM partner might wish to ship XenServer with a suite of management tools that require SNMP agents to be installed, or provide a driver that supports the latest hardware. Users can add supplemental packs either during initial XenServer installation, or at any time afterwards. Facilities also exist for OEM partners to add their supplemental packs to the XenServer installation repositories to allow automated factory installations. For more information, refer to the Developer Documentation.

To install a Supplemental Pack using XenCenter

1. Download the supplemental pack (`filename.iso`) to a known location on your computer. Supplemental packs are available to download from the XenServer Downloads page.

2. From the XenCenter menu, select **Tools** and then **Install Update**.

3. Read the information on the **Before You Start** page and then select **Next** to continue.

4. On the **Select Update** page, click **Browse** to add the supplemental pack and then click **Next** to continue.

5. On the **Select Servers** page, select the pool or server to which you would like to apply the supplemental pack and then click **Next**. This uploads the supplemental pack to the default SR of the pool or the server.

> **Note:**
>
> If the default SR in a pool is not shared, or does not have enough space, XenCenter tries to upload the supplemental pack to another shared SR with sufficient space. If none of the shared SRs have sufficient space, the supplemental pack will be uploaded to local storage on each server.

6. The **Upload** page displays the status of the upload. If there is not enough space on the SR, an error will be displayed. Click **More Info** for details and take necessary action to free up the space required for the upload.

7. After the file is successfully uploaded, XenCenter performs a series of prechecks to determine whether the supplemental pack can be applied onto the selected servers and displays the result. Follow the on-screen recommendations to resolve any update Prechecks that have failed. If you would like XenCenter to automatically resolve all failed prechecks, click **Resolve All**.

8. Choose the **Update Mode**. Review the information displayed on the screen and select an appropriate mode. If you click **Cancel** at this stage, the Install Update wizard will revert the changes and removes the supplemental pack from the SR.

9. Click **Install update** to proceed with the installation. The Install Update wizard shows the progress of the update, displaying the major operations that XenCenter performs while updating each server in the pool.

10. When the supplemental pack installation is complete, click **Finish** to close the wizard. The newly installed supplemental pack will be displayed in the **Updates** section on the **General** tab of the host or the pool. .

For information on installing supplemental packs using the CLI, see the Developer Documentation.

Updating XenCenter

June 3, 2019

If automatic update notification is configured, you may occasionally be notified that a new update is available for XenCenter. XenCenter updates are supplied as XenCenterSetup.exe files and can be downloaded from the Citrix Support website.

To check for new XenCenter versions manually at any time, select **Notifications** , **Updates** and then click **Refresh**.

To download and install a new version of XenCenter:

1. From the XenCenter navigation pane, select **Notifications** and then **Updates**. This displays a list of available updates.

2. Select the required XenCenter update from the list and select **Go to Web Page** from the **Actions** drop-down list. This opens the XenCenter update page in your web browser.

3. Click **Download**. If this is the first time you have downloaded XenServer updates, you will need to install the Akamai Download Manager. When prompted, accept the Download Manager security certificate, and then click **Install** to install the Download Manager.

4. Click **Download Now** to download the XenCenterSetup.exe file.

5. Browse to locate the folder where you want to save the new file and then click **Save** to open the Download Manager and begin the download.

6. Exit your current XenCenter session.

7. When the download is complete (progress shown as 100% in the Download Manager), click **Launch** to begin installing the new version of XenCenter.

8. When XenCenter installation is complete, close the Download Manager.

Update Notifications

June 3, 2019

You can configure XenCenter to periodically check for available XenServer and XenCenter updates and new versions.

To configure updates notification:

1. On the **Tools** menu, click **Options** and then click the **Updates** tab.

2. Select **Check for new versions of XenServer** to have XenCenter periodically check and notify you when a new XenServer version is available.

3. Select **Check for XenServer updates** to have XenCenter periodically check and notify you when updates for XenServer are available.

4. Select **Check for new XenCenter versions** to have XenCenter periodically check and notify you when a new XenCenter version is available.

5. Click **OK** to apply your changes and close the Options dialog box.

Troubleshooting

June 3, 2019

Topics

- XenCenter Alerts
- XenCenter Event Log
- Health Check

- Creating a Server Status Report
- Resolving SR Connectivity Problems
- VM Recovery Mode

XenCenter Alerts

September 30, 2020

You can view different types of system alerts in XenCenter by clicking **Notifications** and then **Alerts**.

The **Alerts** view displays various types of alerts, for example:

- **Performance alerts**. Performance alerts may be generated when CPU, memory usage, network, storage throughput, or VM disk activity exceeds a specified threshold on a managed server, virtual machine, or storage repository. For information on configuring performance alerts, see Configuring performance alerts.
- **HA (High Availability) status alerts**. Alerts may be generated for changes to a pool's high availability status, such as when a pool becomes over committed.
- **License expiry alerts**. Alerts will be generated when the XenServer licenses on your managed servers are approaching their expiry dates or have expired.

Working with alerts

XenCenter is equipped with powerful filtering capabilities. It enables you to filter alerts displayed on the Alerts tab so you can view alerts only from specific pools or servers, or only those generated during a specific time period. For some alerts, it may be possible to quickly address the issue that caused the alert to be generated. The following table lists various options available in the Alerts view.

Filter by Severity

Filters alerts by their severity

By default, alerts of all severity levels will be displayed on the Alerts tab. To view alerts of a particular severity, click **Filter by Severity** and then cancel the selection on other severity levels from the list. Click **Show All** to view all the alerts.

Filter by Location

Filters alerts by the source from which they originate

By default, alerts from all hosts connected to XenCenter will be displayed. To stop displaying alerts from a specific host, click the drop-down list and cancel the selection on the host. Clicking again on the host toggles your selection.

Filter by Date

Filters alerts based on the time of occurrence

By default, all alerts for the current XenCenter session will be displayed. Click the drop-down list and select a date range from the list. Alternatively, click **Custom** to define your own date range by specifying the start and end date/time. Click **Show All** to view all the alerts.

Refresh

If new alerts are generated when the Alerts tab is open, they may not appear in the list. Click **Refresh** to view an updated the list.

Export All

Exports alerts as a comma delimited (.csv) file for viewing and analysis in external applications.

Dismiss All

Removes alerts from the view

To dismiss or remove all the alerts, click **Dismiss All**. To dismiss a specific set of alerts, select the required alerts from the list, and click **Dismiss Selected**.

Actions

Enables you to perform specific actions on the alerts displayed. The Actions drop-down list displays all actions available for the selected alert.

Select an alert from the list and then click the action item relevant to the alert to address it. For example, click:

- **Alarm Settings** to manage alerts for your host's CPU, memory usage, network activity, and storage throughput. This action opens the **Host Properties** dialog box.
- **Copy** to copy information about the alert to the clipboard.
- **Dismiss** to dismiss the alert.
- **Go to Web Page** to open the update page in a web browser.
- **HA settings** to manage High Availability alerts. This action opens the **Configure HA** dialog box.
- **Help** to open the Help topic related to the alert.

- **License Manager** to manage your licenses. This action opens the **License Manager** dialog box.
- **View Log Files** to open the directory where logs are stored.

Receiving alert notifications by email

You can configure XenCenter to send email notifications when alerts are generated for any of the servers and virtual machines in a pool, or for a standalone server and its virtual machines.

When you turn on the email notification feature, you will be sent an email notification when alerts with a priority of 3 or higher are generated. (You can assign a priority for different types of alerts through the XenServer xe CLI; see the *XenServer Administrator's Guide* for details.)

To turn on email notification

1. Select a pool or standalone server in the **Infrastructure** view.

2. click on the **General** tab and then **Properties**.

3. Click on the **Email Options** tab in the **Properties** dialog box.

4. Select the **Send email alert notifications** check box and then type the delivery address details.

   > **Note:**
   >
   > You should enter the details of an SMTP server which does not require authentication. Emails sent through SMTP servers which require authentication will not be delivered. For instructions on using authenticated SMTP servers to receive email notifications, see the *XenServer Administrator's Guide*.

5. Click **OK** to save your changes and close the dialog box.

XenCenter Event Log

June 3, 2019

XenCenter maintains an event log which can be helpful with troubleshooting. You can view a summary of events in the current XenCenter session by clicking **Notifications** and then **Events**. A much more detailed, permanent record of XenCenter events is stored in a log file in your profile folder, and you can use this to troubleshoot any problems that may arise during a XenCenter session.

Viewing events in the current session

To view the events summary for your current XenCenter session, click **Notifications** and then **Events**.

Viewing the XenCenter event log file

A permanent XenCenter log file (syslog) is generated when you use XenCenter. This file includes a complete description of all operations and errors that occur when using XenCenter. It also contains informational logging of events that provide an audit trail of various actions that have occurred in XenCenter and on your managed resources.

The XenCenter log file is stored in %appdata%\Citrix\XenCenter .

The log output from XenCenter is invaluable when diagnosing problems in your XenServer environment. To quickly locate the XenCenter log file, for example, when you need to email it to your support organization, from the XenCenter menu, click **Help** and then **View XenCenter Log Files**.

Working with events in the current session

XenCenter enables you to filter events in the current session and perform a specific action to address them. The following table lists various options available in the **Events** view.

Filter by Status

Filters events by their progress

By default, all events for the current XenCenter session will be displayed. Select a specific status from the drop-down list to toggle the selection.

Filter by Server

Filters events by the source from which they originate

By default, events from all hosts connected to XenCenter will be displayed. To stop displaying events from a specific host, click the drop-down list and cancel the selection on the host. Clicking again on the host toggles the selection.

Filter by Date

Filters events based on the time of occurrence

By default, all events for the current XenCenter session will be displayed. Click the drop-down list and select a date range from the list. Alternatively, click **Custom** to define your own date range by specifying the start and end date/time.

Dismiss All

Removes events from the Events view

To dismiss or remove all the current events click **Dismiss All**. To dismiss a specific set of events, select the required events from the list, and click **Dismiss Selected**.

Actions

Enables you to perform specific actions on the events displayed

Select an event from the list and then click:

- **Dismiss** to dismiss the event
- **Go To** to navigate to the host from which the event originated. Selecting this action item takes to you the Infrastructure view
- **Copy** to copy information about the event to the clipboard

Health Check

September 30, 2020

In previous versions of XenServer, customers had to manually generate and upload the server status reports to Citrix Insight Services (CIS). The Health Check feature enables you to automate the process of generating and uploading the server status report to CIS. After enrolling a pool to Health Check, you will receive notifications in XenCenter regarding the health of the pool. This allows you to proactively monitor the health of the XenServer systems based on the report that CIS generates.

In order to use the Health check feature, you should meet the following requirements:

- All hosts in the pool should be running XenServer 7.1 CU2 or higher
- You should connect to the XenServer pool using XenCenter shipped with XenServer 7.1 CU2 or higher
- XenCenter should have access to the internet
- The Health Check Service should be installed and running on the XenCenter machine. For more information, see Health Check Service later in this topic.
- If using Active Directory (AD), you should have Pool Operator or a higher role

Enrolling a pool to Health Check

When you connect any eligible pool (or standalone hosts) to XenCenter, you will be prompted to enable Health Check. Alternatively, you can select **Tools** > **Health Check** from the XenCenter menu at any time to enable Health Check.

To enroll a pool to Health Check, on the **Health Check Overview** dialog, click **Enroll now**. This opens the **Health Check Enrollment** dialog which allows you to configure Health Check settings based on your requirements. The Health Check feature operates at a pool-level. After enrolling a pool to Health Check, all pool members inherit the same configuration. Note that it is not possible to have different settings between hosts in the same pool.

The following sections guide you through the configuration settings you need to provide to enroll your pool to Health Check.

Health check upload schedule

Specify the schedule at which you would like to upload the server status report to CIS. You can specify the frequency, time, and day of the week you would like to automatically upload the server status report to CIS. The Health Check Enrollment dialog displays some values by default. You can modify these values based on your preference. When an upload is due, the Health Check Service generates a server status report and uploads it to CIS. For information about the report content, see Health Check Server Status Report.

> **Note:**
>
> You can also generate a health check report on demand. For more information, see the section *Uploading server status report on demand*.

XenServer Credentials

When the pool is enrolled to Health Check, the Health Check Service periodically connects to the pool and checks whether an upload is due. You must provide XenServer credentials the Health Check Service can use to establish a connection to the pool. You can use existing credentials or specify new credentials in the **XenServer Credentials** section.

Click **Test Credentials** to verify the credentials.

Authentication with Citrix Insight Services

To successfully upload the server status report and to retrieve CIS analysis report, you should authenticate your uploads with CIS. To do this, you should provide your MyCitrix credentials. If you do not have a Citrix account, visit www.citrix.com to create a new account. Your Citrix credentials will be used to get an authentication token. The token will be stored on the pool master. MyCitrix credentials will not be stored by XenServer or XenCenter. You can use the same credentials to authenticate multiple pools. If you have previously authenticated with CIS, click on the **Use existing authentication** button. Click **OK** to confirm the settings.

Re-authenticating with CIS could solve the issues of repeated failures to upload and failures to retrieve the Health Check analysis reports. To do this, perform the following steps.

1. From the Tools menu, click on Health Check, and then click Edit Health Check settings.
2. In the Authentication with Citrix Insight Services section, select Authenticate using MyCitrix credentials and provide your MyCitrix credentials.

These credentials will be used to get a new authentication token which will be used in the future to upload the server status report and to retrieve analysis reports.

Viewing CIS analysis reports

After you have successfully enrolled your pool to Health Check, the Health Check Service automatically uploads the server status report based on the settings you have specified during enrollment. A summary of the status will be displayed on the Health Check Overview page. It contains information about the last successful upload and the schedule configured during enrollment.

When the CIS analysis report is ready, you will receive an email notification from CIS. The Health Check Overview dialog also displays any issues detected by CIS and provides a link to the analysis report. Click **View report analysis** to see the detailed report on the CIS web site.

In addition, the **Health Check Overview** dialog allows you to perform the following operations:

- **Edit Health Check settings** to review and edit your current enrollment settings
- **Disable Health check** to disable Health Check for the selected pool
- **Request an additional upload now** to upload a SSR on demand. For more information, see *Uploading Server Status Report on demand*.

Uploading server status report on demand

XenCenter generates a server status report and uploads it to CIS based on the schedule configured during Health Check enrollment. If you prefer to perform an on-demand upload for a pool you have enrolled to Health Check:

1. From the XenCenter menu, select **Tools** and then **Health Check**.

2. On the Health Check Overview dialog, select the pool and then click **Request an additional upload now.**

   > **Note:**
   >
   > There must be a delay of at least 30 minutes between upload requests.

3. The status of the upload will be displayed on the Health Check Overview page. Click **View report analysis** to review your report.

Health Check Service

The Health Check Service is installed by default during XenCenter installation. It is responsible for establishing a connection to the XenServer pool using credentials provided during Health Check enrollment. The Health Check Service periodically checks whether an upload is due for a host and initiates report generation and the upload process.

If you chose not to install the Health Check service during XenCenter installation, and you now wish to use the Health Check feature, re-run the XenCenter installer and follow the on-screen instructions. Once the Health Check Service is installed on your XenCenter machine, the service starts up automatically.

Creating a Server Status Report

June 3, 2019

The **Server Status Report** wizard provides a convenient way to collect and package a comprehensive snapshot of a specific XenServer installation for troubleshooting purposes. Options let you include or exclude a range of different configuration files and log files for selected servers.

The Server Status Report gets packaged as a single zip file that can be stored and/or emailed. The size of the report you generate varies, depending on which items you choose to include. The zip file includes:

- A folder for each server, containing the report types you select in the wizard
- XenCenter log files

To generate a server status report

On the **Tools** menu, click **Server Status Report** and follow the steps in the **Server Status Report** wizard:

1. **Select Servers**. Select the servers for which you want to collect report data.

   All available managed servers are listed. If a server not listed, you may be able to add it to the list by clicking **Add New Server**.

2. **Select Report Contents**. Select the data to include in the report and then click **Next**.

3. **Compile Report**. This page shows the progress of the report compilation and reports any problems with the data collection. Click **Next** when the report compilation is complete.

4. **Report Destination**. Browse to locate the folder where the report will be saved and then click **Finish** to save the report files to the specified folder and close the wizard.

Resolving SR Connectivity Problems

September 30, 2020

In order for a storage repository to be available to a server, a connection must exist between the server and the SR. This connection is provided in software by a Physical Block Device (PBD). A PBD stores information that allows a given SR to be mapped to a server. A PBD must be attached or plugged in to the server in order for the SR to be available. If a PBD is unplugged for some reason, the SR will no longer be available to the server and will appear with a broken storage icon ⬛ in the **Resources** pane.

You may be able to diagnose and resolve some common SR connection problems using the **Repair Storage Repository** tool. In the **Resources** pane, select the storage resource, right-click and click **Repair Storage Repository** on the shortcut menu.

Alternatively, on the **Storage** menu, click **Repair Storage Repository**.

The available storage repositories are listed, and you can see their status.

- **Connected**. The connection between the SR and the server is working normally and the storage provided by the SR is currently available.
- **Unplugged**. The storage is unavailable because the PBD is currently unplugged.
- **Connection missing**. The storage is unavailable because the PBD cannot be found.

Click **Repair** to have XenCenter attempt to repair the storage. Progress and results are displayed within the **Repair Storage Repository** dialog box.

VM Recovery Mode

September 30, 2020

If you experience serious problems with a paravirtualized Linux VM, you can try starting it up in Recovery Mode. This turns HVM mode on temporarily and sets the CD drive as the first boot, allowing you to boot a rescue CD or rescue PXE and then investigate the cause of the problem.

To start a VM in Recovery Mode:

- Select the VM that you wish to start in recovery mode.

- From the main menu, choose **VM** > **Start/Shut down** > **Start in Recovery Mode**.

  > **Note:**
  >
  > You should attach your usual operating system rescue CD, boot the VM from this, and then fix the VM from the rescue CD.

See About VMs and Templates to find out more about HVM and paravirtualized modes.