XenServer

XenCenter CR

Contents

What's new in XenCenter

April 25, 2024

XenCenter is updated independently of the version of Citrix Hypervisor or XenServer. To remain supported, ensure that you are using the latest XenCenter version.

The latest version of XenCenter is version 8.2.7. You can download this version of XenCenter from the Citrix Hypervisor downloads page.

What's new in 8.2.7

Released May 11, 2023

This version of XenCenter contains the following behavior changes:

- The Health Check Service has been removed.

  > **Note:**
  >
  > Logs for the Health Check service are retained by Windows for troubleshooting purposes. To remove these logs, delete them manually from `%SystemRoot%\System32\Winevt\Logs` on the Windows machine running XenCenter.

- You can no longer use XenCenter to upload your server status reports (SSRs) to Citrix Insight Services (CIS). Instead, you must generate the report in XenCenter and then go to the Citrix Insight Services website to upload it.

This update includes the following improvements:

- Upgraded the third-party library log4net included in XenCenter to version 2.0.15.

Fixed issues in 8.2.7

This update includes fixes for the following issues:

- XenCenter fails to prompt you to reinsert your credentials during rolling pool upgrades (RPUs) if your credentials were modified since the last connection to the pools in question.
- If you create a new VM in XenCenter and select UEFI Secure Boot for the Boot Mode, XenCenter does not warn you that the selected mode is unavailable.
- During the installation of a driver disk, if the precheck for prerequisite packages fails, XenCenter does not recommend the correct minimum version of the required package.

- When adding a new disk to a VM, or copying or moving a VM or disk within a pool, selecting the SR in which to place your disk takes a long time due to XenCenter scanning all available SRs automatically.
- XenCenter shows a message marking Dynamic Memory Control (DMC) as deprecated. This is no longer the case. DMC is supported in future releases.

Known issues in 8.2.7

This update contains the following known issues:

- After a standalone host is rebooted, including when it's rebooted after applying updates, the host's General Tab does not display the status of the system correctly. We recommend to refresh the host's General Tab by clicking on a different object and back on the host, or by disconnecting and reconnecting.
- [Fixed in XS82ECU1058] In Citrix Hypervisor 8.2 CU 1 pools with hotfix XS82ECU1029 applied that have GFS2 SRs, using XenCenter to generate a server status report (SSR) can fail. To work around this issue, generate your SSRs by running the following command in the host console: `xenserver-status-report`. (CA-375900)
- If a hotfix requires another hotfix to already be installed as a prerequisite, XenCenter does not display the name of the prerequisite hotfix. You can find the prerequisite information in the article on https://support.citrix.com for the hotfix you are trying to install. (CA-383054)
- Changing the font size or dpi on the computer on which XenCenter is running can result in the user interface appearing incorrectly. The default font size is 96 dpi; Windows 8 and Windows 10 refer to this font size as 100%.
- On Windows 10 (1903 and later) VMs, there can be a delay of a few minutes after installing the Citrix VM Tools before the **Switch to Remote Desktop** option is available in XenCenter. You can restart the toolstack to make this option appear immediately.
- It is not advisable to update the same pool from concurrent instances of XenCenter because this action might disrupt the update process.

  If more than one instance of XenCenter is attempting to install multiple hotfixes on a pool, a server might fail to install a hotfix with the error: "The update has already been applied to this server. The server will be skipped."This error causes the whole update process to stop.

  To work around this issue:

  1. Ensure that no other XenCenter instance is in the process of updating the pool
  2. Refresh the update list in the **Notifications > Updates** panel
  3. Start the update from the beginning

- In XenCenter, when you attempt to import an OVF package or a disk image from a folder containing a hash character (#) in its name, the import fails with a null reference exception.
- Use the latest XenCenter to upgrade from Citrix Hypervisor 8.2 CU1 to XenServer 8. Using an older version of XenCenter can result in a loss of connectivity.

  Download the latest XenCenter from the XenServer product downloads page.

Earlier releases

This section lists features in previous releases along with their fixed issues. These earlier releases are superseded by the latest version of XenCenter. Update to the latest version of XenCenter when it is available.

XenCenter 8.2.6

Released Sep 20, 2022

> **Note:**
>
> Only XenCenter 8.2.6 and later can check for and download hotfixes released after Dec 31, 2022.

This update includes the following improvements:

- Updates to third-party packages.
- Usability improvements to the Rolling Pool Upgrade and Install Update wizards:
  - The full name of the VM is displayed, not a shortened version to enable users to correctly map pre-checks problems to the correct VM.
  - The update and upgrade logs now include event timestamps by default. To change this setting, go to **Tools > Options > Display > Log consoles options > Show timestamps** on the update and upgrade log consoles.

Fixed issues This update includes fixes for the following issues:

- If you have the Container Supplemental Pack installed on your XenServer 7.1 CU2 host and attempt to upgrade to Citrix Hypervisor 8.2 CU1 by using XenCenter, you are prevented from upgrading because the supplemental pack is no longer supported.
- When installing updates or performing a Rolling Pool Upgrade in XenCenter, the scrollbar in the output window can become inoperative.
- The Japanese-language edition of XenCenter cannot apply downloaded updates.

- When exporting a VM with a lot of data to a network location, the export fails if the system on which XenCenter is running doesn't have enough storage to accommodate the size of the VM that is being exported.

XenCenter 8.2.5

Released Mar 21, 2022

This update includes the following improvement:

- Improved a misleading confirmation message, which was shown when dismissing update notifications.

Fixed issues This update includes fixes for the following issues:

- In XenCenter 8.2.3, importing an OVF or OVA package can be slower than in earlier versions of XenCenter. This effect is most noticeable for VMs with empty or not very full disks, as these VMs take the same amount of time to import as in the case where the disk is full of data. In XenCenter 8.2.5, this issue is fixed for packages with VHD disks.
- If you attempt to restore update notifications when there are hosts disconnected from XenCenter, XenCenter crashes.
- XenCenter does not display the VM network usage in the **Search** tab.
- When doing an automatic update that includes an update to Citrix Hypervisor 8.2 Cumulative Update 1, XenCenter does not perform all prechecks. As a result, the update can get stuck.

XenCenter 8.2.4

Released Dec 13, 2021

This update includes the following improvements:

- To provide a more secure service for hotfix downloads, XenCenter now requires that you authenticate it with Citrix to automatically download and apply hotfixes. To receive these hotfixes through XenCenter, you must also install the latest version of XenCenter and obtain a client ID JSON file. For more information, see Authenticating your XenCenter to receive updates.
- The version of PuTTY embedded in XenCenter 8.2.4 and later is updated to 0.76.

XenCenter 8.2.3

Released Apr 20, 2021

This update includes the following improvements:

- The mechanism used for OVF/OVA import/export and single disk image import has been simplified and these operations are now performed without using the Transfer VM. This change improves the performance and security of the import and export process.
- XenCenter 8.2.3 now uses the SHA-256 cryptographic algorithm to create a manifest for or digitally sign an exported OVF/OVA package.

  Note: Older versions of XenCenter do not expect this algorithm. If you want to use an older version of XenCenter to import OVF/OVA packages that were exported by the latest version of XenCenter, you must skip the manifest or signature verification step of the import.
- Uses of the MD5 cryptographic algorithm have been removed from XenCenter.

Fixed issues This update addresses the following issues:

- If you have FIPS compliance enabled on the system where XenCenter is installed, you cannot import or export VMs in OVF/OVA format or import Virtual Hard Disk images.
- The Conversion Manager and **New Conversion** wizard are not localized to Simplified Chinese or Japanese.
- When updating a pool to Citrix Hypervisor 8.2, XenCenter can experience a long delay reconnecting to storage after updating each server in the pool. The length of the delay is longer for pools with a larger number of servers.
- XenCenter displays the wrong units on the SR latency graphs for XenServer 7.1 CU2 or Citrix Hypervisor 8.2 servers that are up to date with the latest hotfixes. For example, if the value is 30 milliseconds, XenCenter displays 30 seconds.
- In XenCenter, you cannot delete a custom field.

XenCenter 8.2.2

Released Dec 09, 2020

This update includes the following improvements:

- While pool secret rotation is in progress, XenCenter now prevents you from designating a new pool master or from enabling HA or clustering.
- The **Enter Maintenance Mode** dialog now provides the ability to rerun the pre-checks and refresh the state of VMs on the server. This aids in resolving issues that prevented a Citrix Hypervisor server from entering maintenance mode.

Fixed issues This update addresses the following issues:

- When you are using the Pool Admin role, XenCenter can fail to create an SR and then reports the error "root element missing".

- If XenCenter cannot load the proxy credentials from the user settings, it can crash.
- When attempting to enable clustering in a pool, XenCenter allows you to select a bonded network while its creation is not yet complete. This causes the operation to fail.
- When hovering over a disabled button on the VM's **Storage** tab, XenCenter can display a tool tip that states the wrong reason or no reason for the button being disabled.
- When putting a Citrix Hypervisor server into maintenance mode, the **Enter Maintenance Mode** button can be enabled even if the VMs running on that server are still in the process of shutting down.
- When putting a Citrix Hypervisor server into maintenance mode, the **Scanning for VMs** dialog takes focus while VM shutdown or migration activities are happening in the background. This dialog has now been removed.

XenCenter 8.2.1

Released Sep 15, 2020

This update includes the following improvements:

- Upgrade XenCenter to use .NET Framework 4.8.
- Upgrade the version of PuTTY included in XenCenter to version 0.74.
- For NVIDIA virtual GPUs, the columns **Max resolution** and **Max displays** have been removed from the **GPU** tab of the VM properties and the dialog for the configuration of allowed virtual GPU types on a GPU. These columns are no longer applicable as NVIDIA now support variable resolutions and displays.

Fixed issues This update addresses the following issues:

- If you individually dismiss many alerts from the **Notifications > Updates** tab in quick succession, XenCenter can freeze.
- During storage live migration of a VM, any ISOs in the VM are ejected. These ISOs are not reinserted after migration is complete.
- When launching the **Install Update** wizard from a **Download and Install** action in the **Updates** tab, The **Install Update** wizard can incorrectly show servers as unavailable for update.
- After restarting a Linux VM, the **Open SSH Console** option is not available.
- XenCenter reported an incorrect virtualization state for VMs that were converted from PV to HVM.

XenCenter 8.2.0

This update includes the following improvements:

- Enable and disable read caching from within XenCenter.

  The read caching feature improves performance on NFS, EXT3/EXT4, SMB, or GFS2 SRs that host multiple VMs cloned from the same source. This feature can now be enabled and disabled for each individual SR from the XenCenter console. You might want to disable read caching in the following cases:

  - You have no file-based SRs
  - You do not have any cloned VMs
  - You have insufficient memory available to allocate to dom0 to derive any performance benefits

  For more information, see Changing SR Properties.

Fixed issues This update addresses the following issue:

- When creating an LVM SR from XenCenter and passing CHAP credentials, the operation might fail with an authentication error.

Getting Started with XenCenter

May 25, 2023

With XenCenter, you can manage your Citrix Hypervisor environment and deploy, manage, and monitor virtual machines from your Windows desktop machine. See the topics in the following table to get started.

 Add a New Server	Connecting to Citrix Hypervisor host servers and adding them to the list of managed resources in XenCenter.
 Create a Shared Storage	Creating Citrix Hypervisor 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 Hypervisor and XenCenter, see the system requirements.

Installing XenCenter

April 16, 2024

XenCenter must be installed on a Windows machine that can connect to the Citrix Hypervisor server through your network.

In addition, XenCenter has the following system requirements:

- **Operating System:**
  - Windows 10
  - Windows 11
  - Windows Server 2016
  - Windows Server 2019
  - Windows Server 2022
- **.NET Framework:** Version 4.8
- **CPU Speed:** 750 MHz minimum, 1 GHz or faster recommended
- **RAM:** 1 GB minimum, 2 GB or more recommended
- **Disk Space:** 100 MB minimum
- **Network:** 100 Mbit/s or faster NIC
- **Screen Resolution:** 1024x768 pixels, minimum

XenCenter is compatible with all supported versions of Citrix Hypervisor.

To install XenCenter:

1. Download the installer for the latest version of XenCenter from the Citrix Hypervisor Download page.

2. Launch the installer `.msi` file.
3. Follow the Setup wizard, which allows you to modify the default destination folder and then to install XenCenter.

Connect XenCenter to the Citrix Hypervisor server

To connect XenCenter to the Citrix Hypervisor server:

1. Launch XenCenter. The program opens to the **Home** tab.
2. Click the **Add New Server** icon.
3. Enter the IP address of the Citrix Hypervisor server in the **Server** field. Type the root user name and password that you set during Citrix Hypervisor installation. Click **Add**.
4. The first time you add a host, the **Save and Restore Connection State** dialog box appears. This dialog enables you to set your preferences for storing your host connection information and automatically restoring host connections.

   If you later want to change your preferences, you can do so using XenCenter or the Windows Registry Editor.

   To do so in XenCenter: from the main menu, select **Tools** and then **Options**. The **Options** dialog box opens. Select the **Save and Restore** tab and set your preferences. Click **OK** to save your changes.

   To do so using the Windows Registry Editor, navigate to the key `HKEY_LOCAL_MACHINE\Software\Citrix\XenCenter` and add a key named `AllowCredentialSave` with the string value **`true`** or **`false`**.

Starting or Exiting XenCenter

May 25, 2023

Starting XenCenter

To start your 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 you previously configured XenCenter to restore your server connections on startup and set a master password, XenCenter prompts you to enter this password before continuing. See Store Your Server Connection State to find out more about how to set your server reconnection preferences.

It is possible to run only one XenCenter session per user.

Exiting XenCenter

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

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

If there are any XenCenter tasks in progress, XenCenter warns you when you try to exit. You can choose to exit anyway, in which case unfinished tasks might not complete successfully. Alternatively, you can wait until the unfinished tasks have completed.

Uninstalling XenCenter

May 25, 2023

To uninstall XenCenter:

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

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

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

Uninstalling by using `msiexec`

If you installed XenCenter by using `msiexec`, it might not appear in the **Add or Remove Programs** list. In this case, you can instead use `msiexec` to uninstall the program.

Open a command prompt and run the following command:

```
msiexec /x <xencenter-installation-msi-file-name>
<!--NeedCopy-->
```


Replace <xencenter-installation-msi-file-name> with the name of the XenCenter installer .msi file.

Exploring the XenCenter workspace

May 25, 2023

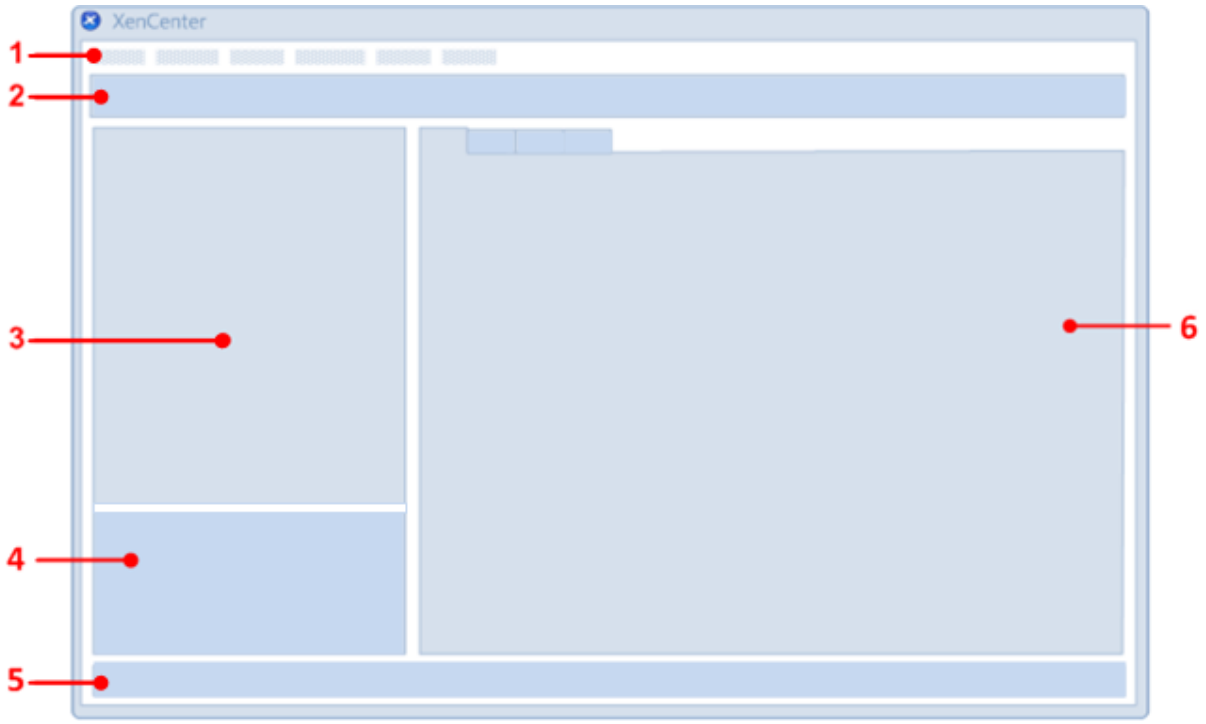

Ref #	Name	Description
1	Menu bar	Includes all the commands you need to manage servers, pools, SRs, VMs, and templates.
2	Toolbar	Provides quick access to a subset of the most frequently used menu commands. See The Toolbar.
3	Resources pane	Lists the servers, pools, VMs, templates, and SRs currently being managed from XenCenter. See The Resources Pane

Ref #	Name	Description
4	Navigation pane	Lists the navigation buttons. Click 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

May 25, 2023

The XenCenter toolbar provides quick access to some of the most common XenCenter tasks, for example, to connect to new servers and create VMs.

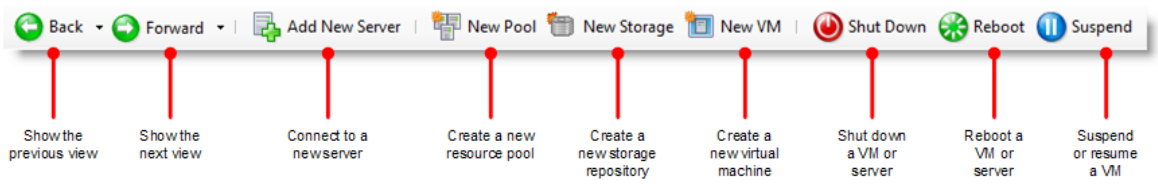

Using the Back and Forward buttons on the Toolbar

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, select **Back**.
- To display the next resource view (if you have used **Back**), select **Forward**.
- To display one of the resource views you used in this session, select the down arrow next to the **Back** or **Forward** buttons, and then select the view from the list.

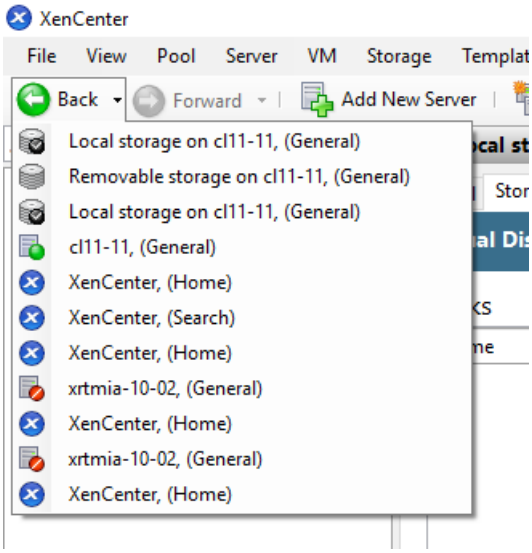

Showing and hiding the Toolbar

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, select to remove the **Show Toolbar** check mark.
- On the **View** menu, select to remove the **Toolbar** check mark.

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

The Resources Pane

May 25, 2023

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. See the table in the following section for information about various buttons in the **Navigation** pane.

To do a simple text search on resource names, type 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. For more information, see Saved searches.

You can also create and add your own custom searches to this list at any time. For more information, see Create a Search Query.

To apply a saved search to the contents of the **Resources** pane, select **Saved Searches** in the **Navigation** pane and select a search query from the list.

The following table lists the 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.
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

November 16, 2023

The XenCenter **Navigation** pane provides various options to view and access managed resources. The following navigation buttons provide a quick way to view and manage your resources:

- **Infrastructure**
- **Objects**
- **Organizations Views**
- **Saved Searches**
- **Notifications**

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

Infrastructure

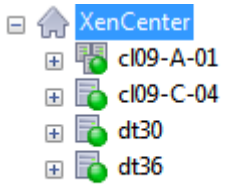

This view is the default view. The **Infrastructure** view displays a tree view of the resources by their physical location. It provides a list of servers, VMs, templates, and storage resources by the pool or the server to which they belong.

Objects

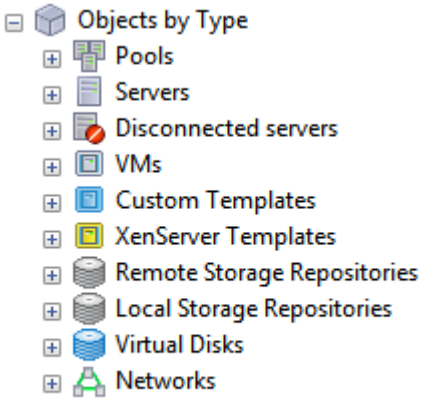

Select **Objects** to see a list of the resources by categories such as pools, servers, VMs. Expand the nodes to view items in each category.

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

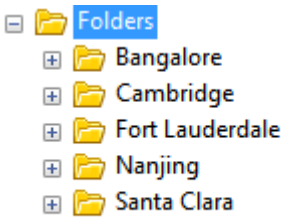

Select this option to view your resources by folders. You can create folders to group your resources by location, function, resource type, and so on

> **Note:**
>
> Organizing resources into a folder is conceptual, and not physical. The resources are note 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

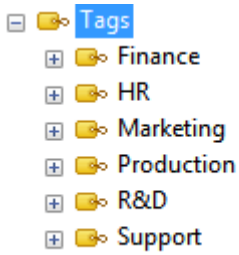

Select this option to view your resources by the tags that you have previously defined. Tags are labels that you specify to view resources based on the criteria that you define. A single resource can contain multiple tags. For example, a server with the tag 'Production'can also be tagged as 'R&D'.

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

Objects by Custom Field

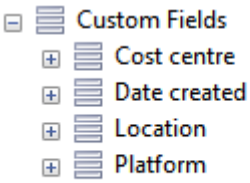

Select this option to view your resources by the customized fields you have previously defined. XenCenter enables you to add custom fields to your resources and provide a value to effectively manage your resources. You simply add a custom field to a server, VM, or any other resource in the pool, then 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

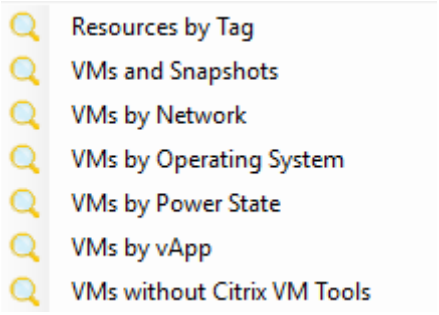

Click this button and select an option from the list to view resources that match the search criteria. By default, XenCenter includes a few saved searches that allow you to search your resources. You can create and add your own query to this list at any time.

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

Notifications

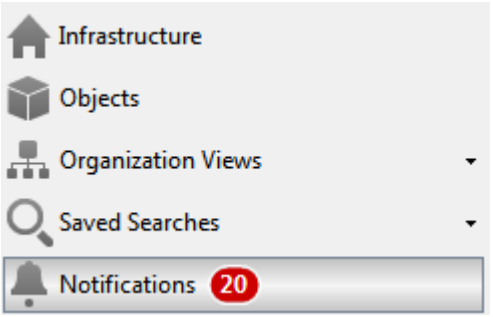

Click this button for the **Notifications** view. The **Notifications** view enables users to see all notifications in a centralized location and perform specific actions to address them. It contains Alerts, Updates, and the Events view.

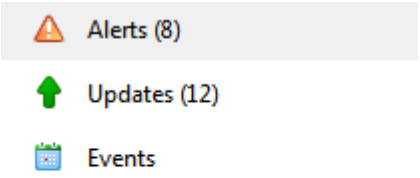

Alerts

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 Citrix Hypervisor 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

April 16, 2024

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. 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 or Switch to Default Desktop

Switches between Windows remote console types

Switch to Graphical Console or Switch to Text Console

Switches between Linux remote console types. You might 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:**
>
> 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.

Most keyboard shortcuts are transmitted to the server or VM when you use a remote console. However, your local system always intercepts the **Ctrl+Alt+Del** key sequence and prevents it from being sent if you 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 Citrix VM Tools, shut down, reboot or suspend a virtual machine from within the floating console window, select the lifecycle icon in the top-left corner of the window and then click 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) or 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, but this behavior 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. The **Connection bar** shows 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 various console settings in the **Options** dialog box. For example, the text clipboard on your local machine is shared with the remote console by default. Items you cut or copy are 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 various 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.

Copy any of the values shown on this pane to the Windows clipboard by right-clicking on the value and 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 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 the following articles:

- Configuring Virtual GPU
- Change Pool Properties.

> **Note:**
>
> - GPU Pass-through and Graphics Virtualization are available for Citrix Hypervisor Premium Edition customers, or those customers who have access to Citrix Hypervisor 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.

USB

The **USB** tab allows you to pass through individual physical USB devices to a VM so the VM's OS can use it as a local USB device. You can enable or disable pass-through by clicking the **Enable Pass-through** or **Disable Pass-through** button on the **USB** tab. To attach a USB, perform the following steps:

1. Shut down the VM.
2. Right-click the VM and select **Properties**.
3. On the left pane, click **USB**.
4. Click **Attach**.
5. In the Attach USB dialog box, click **Attach**.
6. Start the VM. The USB is now attached to the VM.
7. In the same way, click **Detach** to detach the USB from the VM.

USB pass-through is supported only on the following guests:

Windows

- Windows 8.1
- Windows 10
- Windows Server 2016

Linux

- RHEL 7
- Debian 8

> **Note:**
>
> - USB pass-through is supported for the following USB versions: 1.1, 2.0, and 3.0.
> - USB pass-through supports a maximum of 6 USBs to be passed through to a single VM.
> - Snapshot/Suspend/ Pool Migrate/ Storage Migrate operations are not supported when USB is passed through to VM.
> - USB pass-through feature is available for Citrix Hypervisor Premium Edition customers.
> - Plugging in untrustworthy USB devices to your computer might put your computer at risk. Assign USB devices with modifiable behavior only to trustworthy guest VMs.
> - Do not boot BIOS from USB devices.
> - Ensure that the USB device to pass-through is trustworthy and can work stably in normal Linux environment (for example, CentOS 7).
> - USB device pass-through is blocked in a VM if high availability is enabled on the pool and the VM has restart priority as **Restart**. The USB attach button is disabled and the following message is displayed: **The virtual USB cannot be attached because the VM is protected by HA**. When configuring high availability for a pool, if a VM is not agile, the **Restart** option is disabled with the following tooltip: **The VM has one or more virtual USBs. Restart cannot be guaranteed**.

High availability

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

- Enable high availability using the **Configure HA** button.
- Change the pool's high availability configuration using the **Configure HA** button.
- Disable high availability.

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

For more information, see the following articles:

- Configure high availability
- Disable high availability
- Change high availability Settings

Home

The **Home** tab allows you to add a server to the list of managed servers or open a browser window to find out more about Citrix Hypervisor.

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 use DMC. DMC 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.

For more information, see the following articles:

- Changing the Control Domain Memory
- About VM Memory Configuration

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.

For more information, see the following articles:

- About Citrix Hypervisor Networks
- View and Change Network Properties.

NICs

View detailed information about the physical 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 is automatically routed over the second NIC, ensuring server management connectivity. See Configuring NICs.

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

Performance

View performance data for your VMs and managed servers on the **Performance** tab. Full performance data is only available for VMs with Citrix VM Tools installed.

The tab provides real-time monitoring of performance statistics across resource pools and 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, VM, or SR. For more information, see Configuring Performance Alerts.

Search

Select the top-level XenCenter item, pool, or server in the **Resources** pane and then click 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, high availability status or restart priority, and power state.

For more information, see the following articles:

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

Snapshots

Create, delete and export VM snapshots, revert a VM to a selected snapshot, and use existing snapshots to create 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. This information includes 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 disk or **Attach** to attach an existing disk.	Configuring Virtual Storage
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 an 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 Citrix Hypervisor users and groups through AD user account provisioning and Role Based Access Control (RBAC) on the **Users** tab. In this tab you can do the following tasks:

- Join a pool or server to an Active Directory (AD) domain
- Add an AD user or group to a pool
- Assign roles to users and groups.

For more information, see Managing Users.

WLB

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

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

Resource Status Icons

September 18, 2023

The 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.

Icon	Description
	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 Citrix Hypervisor-related problems. This file 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 is restored to normal.
	A server for which updates are available. See Updating Managed Servers.
	A server that is running an older version of Citrix Hypervisor 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.
	Citrix Hypervisor VM template.

Icon	Description
	A custom (user-defined) VM template.
	Citrix Hypervisor vApps. 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.
	A storage repository that is currently unavailable. See Resolving SR Connectivity Problems.
	A virtual disk.
	A virtual disk snapshot. This object is a snapshot of a VM's disks, and is created when a snapshot is made of the VM. See VM Snapshots for information about taking VM snapshots, and see Snapshots to find out more about disk snapshots.

Keyboard Shortcuts

May 25, 2023

You can use the keyboard in addition to 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.
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**.
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 **New VM** on the **VM** menu. Some shortcut keys are shown on the menus and in toolbar tooltips. For numeric keypad keys, ensure that **Num Lock** is off.

Key	Action
F1	Display the online Help

Key	Action
Alt+F4	Exit XenCenter and close the XenCenter window
Ctrl+Enter	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
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

Keyboard shortcuts are also available for working with VM consoles. These shortcuts are configurable. For more information, see Console Settings.

Changing XenCenter Options

December 9, 2023

You can change various settings that affect your XenCenter working environment. On the **Tools** menu, click **Options**, click 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 TLS 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
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.

XenCenter Updates settings

Option	Description	Default
Check for new versions of XenCenter	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:

- Software updates

- Update your hosts
- Updating XenCenter

Display settings

This tab enables you to configure how graphs are displayed in the **Performance** tab.

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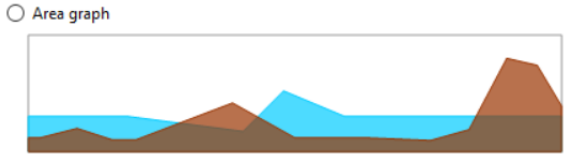
Line graph	Choose this radio button if you want to view the performance data shown on the **Performance** tab as line charts. For example: 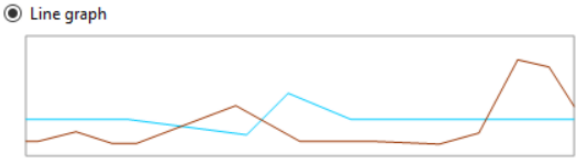

See also: Monitoring system performance.

You can also use this tab to configure whether XenCenter remembers the last selected tab for a resource.

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 are placed on the clipboard and made available for pasting on either your local computer or on the remote console.	On
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 the 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 key 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

Option	Description	Default
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
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 where you run 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. When opening a Remote Desktop console session, XenCenter uses this session instead of creating a new virtual 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. When this option is enabled, XenCenter continues to scan the RDP port even if the **Automatically switch to Remote Desktop** option is turned off. You can switch to RDP when it becomes available.	On

Any changes you make to the Windows Remote Desktop console settings 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. Use HTTP CONNECT to establish a secure TLS tunnel to your servers. Enter the address of the proxy server and the port number to use.

To have **Citrix Hypervisor** connection requests made directly and not through the proxy server, select the **Bypass proxy server for Citrix Hypervisor connections** check box. To have all connection

requests made through the proxy server, clear the check box.

Citrix Hypervisor connection is any connection which provides communication between XenCenter and the Citrix Hypervisor system, such as sending commands and using the console. A non-Citrix Hypervisor connection is something such as checking for updates.

Select the **Provide credentials** check box and enter the **Username** and **Password** that corresponds to a user account that is set up on the specified proxy server.

Select the desired authentication method: Basic or Digest (default).

Authentication method is used to authenticate to the proxy server. Select the same authentication method as the one the proxy server is set for.

For example, if the proxy server requests XenCenter to authenticate using Digest, then XenCenter fails to authenticate if the Basic authentication method is selected.

Default: Don't use a proxy server

Connection timeout

You can specify how long to wait when establishing a connection with a managed server by adjusting the number of seconds to wait for a connection timeout. Do not set this value too low if you don't want to receive many 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 whether to store your login credentials for managed servers. Stored login credentials can be used to reconnect automatically 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 user name and password - for all your managed servers can be stored between XenCenter. These credentials are used to automatically reconnect to them at the start of each new XenCenter session. When this check box is selected, XenCenter remembers the connection state of all your managed servers at the end of each session. XenCenter attempts to restore these servers at the start of your next session.	Off
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 are prompted to enter this master password before connections to your managed servers are automatically restored.	Off
Change Master Password	Select to change the current master password. You are 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 the XenCenter plug-in. For example, you might do this action 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 webpage, and can call out to other services on your network or to your VMs.

Plug-in components

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.

Place the plug-in components into a plug-ins subfolder in your XenCenter installation folder. The components are loaded when XenCenter starts. For example, in a default installation of XenCenter, your plug-ins would be located here:

```
C:\Program Files (x86)\Citrix\XenCenter\plugins \<
   your_organization_name>\<your_plugin_name>
<!--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, select **Options**. The **Options** dialog box is displayed. From the list of options on the left pane, select **Plugins**.

Default: On

Creating plug-ins

To learn how to create plug-ins for XenCenter, see the samples and accompanying documentation in the XenCenter Plug-in Specification and Examples repository. You can access this webpage anytime from XenCenter by clicking **XenCenter Plugins online** on the **Help** menu.

For more information, see the XenCenter Plug-in Specification Guide](/en-us/citrix-hypervisor/developer/xencente
plugin-specification.html).

Confirmations settings

Use the settings on this tab to configure whether to display a confirmation dialog in the following cases:

- When you dismiss an alert
- When you dismiss an update
- When you dismiss an event

Hidden Objects

May 25, 2023

Hide networks, PIFs, and VMs from XenCenter by adding the key `HideFromXenCenter=true` to the `other_config` parameter for the object in the Citrix Hypervisor Management API. For example, certain VMs can be hidden to prevent them being used directly by general users in your environment. Objects flagged with this key are hidden objects, and by default they do 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 developer documentation to find out more about flagging objects using the `HideFromXenCenter` key.

Organizing Resources

May 25, 2023

XenCenter provides various 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

May 25, 2023

A folder is a container that groups 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 Citrix Hypervisor 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"or "not contained in"relation and a list hierarchy of folders. For more information, see Create a Search Query.

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

To create a folder

The simplest way to create a folder is through Resources pane. Click **Organization Views** in the **Navigation** pane, and then select **Objects by Folder**. In the **Resources** pane, click 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 is stored, and then click **Create**.

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

1. In the **Resources** pane, select a pool, server, SR, virtual machine, or template, then click 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 is 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

Move a resource into a folder by dragging it from the **Resources** pane or **Search** tab and to the folder. Resources can only be in one folder. If the resource is already in another folder, it is moved when you drag 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.
2. Select the **General** tab and click **Properties**.
3. On the **General** tab of the **Properties** dialog box, click **Change** in the **Folder** box.
    - To remove the resource from its current folder, select **Not in any folder**.
    - To move the resource into a different folder, select **In this folder** and choose a folder or a subfolder from the list.
    - 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 is stored, and then click **Create**.
4. 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 shortcut menu.
2. Type the new name.

To delete a folder

You do not actually delete the resources in the folder when you delete the folder. The resources return to the general collection.

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

All the contents are moved out of the folder and then the folder is deleted.

Using Tags

May 25, 2023

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

You make up tags when you need them and you can use as many as you like. You can also build searches based on your tags, for example, "all Windows 10 test machines located in Cambridge".

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

To create a tag

1. In the **Resources** pane, select a resource.
2. Select its **General** tab and then select **Properties**.
3. On the **General** tab of the **Properties** dialog box, select **Edit tags**.
4. Type a word or phrase in the **New Tag** box and then select **Create**.
5. The new tag is automatically assigned to the selected resource. To remove the tag, clear the check box.
6. Select **OK**.
7. On the **Properties** dialog box, select **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 is displayed on the **Resources** pane.
2. Select the tag you want to delete, right-click, and then select **Delete Tag**. The tag is 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 it onto the tag in the **Resources** pane. You can drag resources from within the **Resources** pane on the **Tags** group or 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.
2. Select the **General** tab and then select **Properties**.
3. On the **General** tab of the **Properties** dialog box, select **Edit tags**.

- To create a tag and add it to the selected resource, type a word or phrase in the **Edit Tags** box. Click **Create**.
- To add an existing tag, select to select the tag's check box in the **Tags** list. Click **OK**.

4. On the **Properties** dialog box, select **OK** to apply your changes.

To untag a resource

To untag (remove a tag from) a resource, in the **Navigation** pane, select **Organization Views** and then select **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.
2. Select its **General** tab and then select **Properties**.
3. On the **General** tab of the **Properties** dialog box, select **Edit tags**.
4. Clear the check box for the tag in the **Tags** list and select **OK**.
5. On the **Properties** dialog box, select **OK** to apply your changes.

Using Custom Fields

May 25, 2023

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

For example, you might identify all hosts with their physical location. Alternatively, you might label the cost center and applications running on all of your VMs. 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. If you set a custom field on any resource in a pool, this field is available to all resources in the pool. You can access custom fields 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 custom field

1. In the **Resources** pane, select any resource.
2. Select the **General** tab and then click **Properties**.
3. On the **Custom Fields** tab, click **Edit Custom Fields**.
4. Click **Add**, type a name for the custom field and select the field type.
5. 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.
2. Select the **General** tab and then click **Properties**.
3. On the **Custom Fields** tab in the **Properties** dialog box, enter a value for the custom fields.
4. 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.
2. Select the **General** tab and then click **Properties**.
3. On the **Custom Fields** tab in the **Properties** dialog box, click **Edit Custom Fields**.
4. Select the custom field in the list and then click **Delete**.
5. Click **Yes** to confirm.

Searching Resources

May 25, 2023

XenCenter enables you to perform complex searches of your managed resources. By default, XenCenter includes several searches. These searches allow you to search resources by tag. They also enable you to search VMs by network, operating system, power state, vApp, and Citrix VM 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 select 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. The **Search** box is located above the **Resources** pane. Matching resources are displayed as you

type. To remove the query and view all your resources again, select the **x** button at the right of the **Search** box. For more information, see the following topics:

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

Create a Search Query

May 25, 2023

Using the **Search** tab, you can construct queries based on object types, folders, and attributes. Attributes can include name, description, tags, high availability status, restart priority, and power state.

To create a 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 resource types such as servers, VMs, and SRs, this list also contains some common combinations of resource types. It also provides 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 when you select an option under **Search for**. 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 the **Location** list to choose the server where the search query metadata is saved.
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 when you select a filter option, and the results are 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. Resources can only be in one folder, so if the resource is already in another folder it is 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

May 25, 2023

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 you search for servers, you can filter the results by server name, server IP address, its resource pool, and the VMs on it. When you search for virtual disks, you can filter the results using criteria such as whether they are shared or the storage repository they are on.

To filter your search results

1. Click the filter button and choose a filter category from the list:

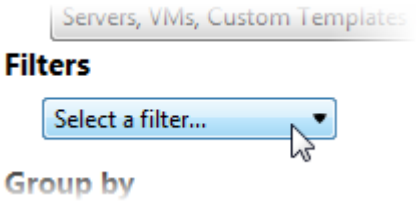

   You can only select filters that are applicable to the resources you are searching for.

2. If applicable, select an operator, such as **Is**, **Contains**, or **Start Time** and then enter a value. The filter applies to the search result when you select an option here. The results are updated immediately.
3. To add more filters, click the filter button and choose a multi-filter category from the list:

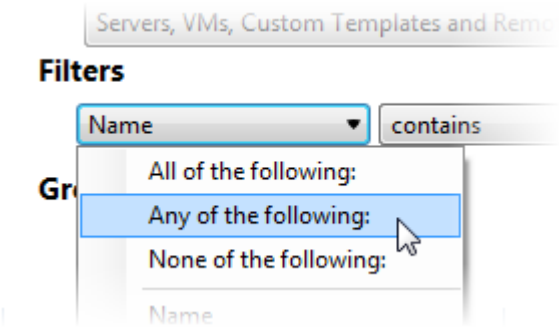

4. To remove a filter, click the **Remove** button on the right of the **Search** tab.

To group search results

1. Under **Group by**, select the grouping options you want to apply from the list. The grouping is applied when you select an option here. 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 the group button and select **Remove Grouping**.

Saved Searches

May 25, 2023

XenCenter includes several useful saved searches. These searches allow you to search resources by tag. They also enable you to search VMs by network, operating system, power state, vApp, and Citrix VM Tools status. Modify these search queries by clicking the **Edit Search** button on the **Search** tab.

You can create and add your own custom searches to this list at any time. For more information, see Create a Search Query. Custom searches appear at the top of the **Saved Searches** list in the **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 **Navigation** pane. Search results are 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

May 25, 2023

XenCenter enables you to save search queries as `.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, select **Export**.
2. Enter a file name and location.
3. Click **Save** to confirm.

To import a search

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

To save the imported search as a new custom search

1. Click **Save**.
2. Specify a **Name** for the search query.
3. Choose the **Location** where to store the search query metadata.
4. Click **Save** to confirm.

About Citrix Hypervisor Licensing

August 30, 2023

Citrix Hypervisor requires a License Server to run with a Premium Edition or Standard Edition license. For information about how to install and run Citrix Licensing, see Citrix Product Documentation. To use Citrix Hypervisor Express Edition, you do not require a license or a License Server. For more information, see Citrix Hypervisor Licensing.

After purchasing support for Citrix Hypervisor you are provided with the `.LIC` license access code. Install this license access code on a Windows server running the Citrix License Server software.

> **Note:**
>
> Previously, a Linux-based Citrix License Server virtual appliance was available. This virtual appliance is no longer supported.

Citrix Hypervisor Licensing depends on the version of the Citrix Hypervisor product you have installed on your server. For more information, see:

- Licensing Overview
- Managing Citrix Hypervisor Licenses

Licensing Overview

January 25, 2024

Citrix Hypervisor is available in two commercial editions:

- Standard Edition
- Premium Edition

The **Standard Edition** is our entry-level commercial offering. It includes a range of features that deliver a robust and high performing virtualization platform, but not the premium features offered by the Premium Edition. With Standard Edition, you can still benefit from the assurance of comprehensive Citrix Support and Maintenance.

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

- Automated Windows VM driver updates
- Automatic updating of the Management Agent
- Support for SMB storage
- Direct Inspect APIs
- Dynamic Workload Balancing
- GPU virtualization with NVIDIA vGPU, AMD MxGPU, and Intel GVT-g
- VMware vSphere to Citrix Hypervisor conversion utilities

- Export pool resource data
- In-memory read caching
- PVS-Accelerator
- Citrix Hypervisor live patching
- Enablement for Citrix Virtual Desktops tablet mode
- Changed block tracking
- IGMP snooping
- USB pass-through
- SR-IOV network support
- Thin provisioning for shared block storage devices

> **Notes:**
>
> If you have purchased Citrix Virtual Apps and Desktops, you continue to have an entitlement to Citrix Hypervisor that includes all features in the preceding list.
>
> Automated Updates were previously restricted to Citrix Hypervisor Premium Edition customers or Citrix Virtual Apps and Desktops customers. However, in pools with hotfix XS82ECU1053 applied, this feature is available to all users.

Citrix Licensing

Citrix Hypervisor 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 Citrix Hypervisor, you will receive a `.LIC` license access code. Install this license access code on a Windows server running the Citrix License Server software.

> **Important:**
>
> Citrix Hypervisor 8.1 requires Citrix License Server 11.14 or higher.

When you assign a license to your Citrix Hypervisor host, Citrix Hypervisor contacts the specified Citrix License Server and requests a license for the specified servers. If successful, a license is checked out. The **License Manager** displays information about the license the hosts are licensed under.

When you request or apply licenses, information about the Citrix Hypervisor version and license type might be transmitted to Citrix. No other information relating to users, VMs, or the Citrix Hypervisor environment is collected or transmitted to Citrix. The limited information transmitted to Citrix during the licensing process is handled in accordance with our privacy policy. For more information, see our privacy policy.

Licensing Citrix Hypervisor servers

Citrix Hypervisor does not support partial licensing, all servers in a pool must be licensed. If your Citrix Hypervisor pool contains servers that are licensed with different license types, the lowest license applies to the whole pool. Mixed pools of licensed and unlicensed hosts behave as if all hosts were unlicensed. For information on managing licenses in Citrix Hypervisor, see Managing Citrix Hypervisor Licenses.

Upgrades to the Premium Edition are available from the Standard edition. For detailed information about Citrix Hypervisor licensing, see Licensing. To upgrade or to buy a Citrix Hypervisor license, visit the Citrix website.

License expiry

XenCenter notifies you when your license is due to expire. Purchase a license before it expires. When your Citrix Hypervisor license expires:

- XenCenter License Manager displays the status as **Unlicensed**.
- You cannot access licensed features or receive Citrix Support for any server within the pool until you purchase another license.

License grace period

Citrix licensing has built-in timeout technology. After a Citrix Hypervisor server checks out a startup license, the Citrix Hypervisor server and the License Server exchange "heartbeat"messages every five minutes. These heartbeat messages indicate to each server that the other is still up and running. If your Citrix Hypervisor host cannot contact the License Server, the server lapses into a 30-day licensing grace period. During the grace period, Citrix Hypervisor licenses itself through cached information. The servers are allowed to continue operations as if they are still in communication with the License Server. The grace period is 30 days and when the grace period runs out, Citrix Hypervisor reverts to an unlicensed state. After communication is re-established between Citrix Hypervisor and the License Server, the grace period is reset.

Managing Citrix Hypervisor Licenses

May 25, 2023

This topic contains information about managing licenses in Citrix Hypervisor.

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

- **Assign** Citrix Hypervisor license to managed servers. When you assign a license, Citrix Hypervisor 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** Citrix Hypervisor licenses. When you release a license, Citrix Hypervisor contacts the Citrix License Server and checks the license back in.

> **Important:**
>
> Citrix Hypervisor requires Citrix License Server v11.14 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 Citrix Hypervisor license

1. From the **Tools** menu, select **License Manager**.
2. Select one or more hosts or pools that you want to assign a license and then select **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 Citrix Hypervisor licenses, see Licensing Overview.
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, see the licensing topics on the Citrix Product Documentation website.

XenCenter contacts the specified Citrix License Server and requests a license for the specified servers. If successful, a license is checked out and the information displayed in the XenCenter License Manager is updated.

To release a Citrix Hypervisor license

1. On the **Tools** menu, select **License Manager**.
2. Select the servers or pools and then select **Release License**.

Getting Help

May 25, 2023

There are several different places you can find the information you need about using Citrix Hypervisor and XenCenter:

- **XenCenter product documentation**. Comprehensive reference documentation aimed at XenCenter users.
- **Citrix Hypervisor product documentation**. Comprehensive reference documentation aimed at Citrix Hypervisor administrators and developers.
- **Citrix Knowledge Center**. Browse or search for knowledge base articles and technical notes.

XenCenter help

In XenCenter 8.0 and later, the information that was previously in the XenCenter in-product help is provided as an HTML documentation set.

- Use the table of contents on the left to navigate to the information that you need
- Use the search box in the top-right to search for specific information
- See an outline of the information in each article in the 'In this article' box
- Print individual articles by using the 'Print' button
- Download all content as a PDF for offline viewing by using the 'View PDF' button

Managing Servers

May 25, 2023

Connecting and disconnecting

- Add a 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

Configuring networking

- Citrix Hypervisor 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

May 25, 2023

- 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
- Install a TLS Certificate
- Store Your Server Connection State
- Back up and Restore a Server
- Remove a Server From XenCenter

Add a server

May 25, 2023

To monitor and manage activities on a server from XenCenter, first identify the server as a managed resource. When you first connect to a server, the server appears in 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 can also appear here. A managed server can then be disconnected, reconnected, shut down or put into maintenance mode. 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. By using this dialog, you can set your preferences for storing connection information and restoring server connections at the start of each XenCenter session. For more information, 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: `203.0.113.28` or `server.example.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 user name and the password set up during Citrix Hypervisor installation. If Active Directory (AD) authorization has been enabled in your Citrix Hypervisor environment, you can enter your AD credentials here. For more information, eee RBAC overview.
4. Click **Add**. A connection progress monitor is displayed: to cancel the connection, click **Cancel**.

Security Certificates

You can configure XenCenter to display a warning message whenever it finds a new or modified TLS security certificate while connecting to a managed server. Click **View Certificate** to view the security certificate. To prevent TLS certificate warnings from being generated, use the **Security Settings** tab in the XenCenter **Options** dialog box.

Disconnect a Server

May 25, 2023

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

To see which of your servers are currently disconnected, switch to the **Objects** view in the **Navigation** pane and click **Disconnected servers**.

To disconnect a server:

1. Select the server in the **Resources** pane.
2. 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

May 25, 2023

After you have added a server to XenCenter, it remains accessible in the **Resources** pane throughout the current XenCenter session. It is accessible regardless of the server status: connected or disconnected, running normally or in maintenance mode.

To reconnect to a disconnected server, select it in the **Resources** pane, or right-click and then select **Connect** on the shortcut menu. Connection information for the server is remembered for the current XenCenter session. You do 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, select **Connect/Disconnect** and then **Reconnect As**.
3. Enter the new user name and password. If Active Directory authorization has been enabled in your Citrix Hypervisor environment, you can enter your AD credentials here. See RBAC overview.
4. Click **OK**.

Reboot a Server

May 25, 2023

When you reboot a server in XenCenter, the server shuts down any VMs running on it. After the VMs shut down, the server is disconnected and rebooted. If the server is a member of a pool, the loss of connectivity on shutdown is handled and the pool recovers when the server returns. If you shut down another pool member (not the master), the other pool members and the master continue to function. If you shut down the master, the pool is out of action until the master is rebooted and back on line. When the master restarts, the other members reconnect and synchronize with the master. Alternatively, you can make one of the other members into the master by using the xe CLI.

VMs with Citrix VM Tools installed are shut down gracefully when you reboot the host server. However, VMs without Citrix VM Tools installed are shut down using a forced shutdown. To avoid forced shutdowns, install the Citrix VM Tools on your VMs. For more information, see Installing Citrix VM Tools.

After a server reboot, XenCenter attempts to reconnect to the server automatically. After the server reconnects, restart any VMs that were running on it unless they are configured to start automatically on server reboot. For more information, 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

May 25, 2023

When you shut down a server in XenCenter, the server shuts down any VMs running on it, and then the server is disconnected and powered off. If the server is a member of a pool, the loss of connectivity on shutdown is handled and the pool recovers when the server returns. If you shut down another pool member (not the master), the other pool members and the master continue to function. If you shut down the master, the pool is out of action until the master is rebooted and back on line. At which point the other members reconnects and synchronizes with the master. Alternatively, you can make one of the other members into the master, which you can do by using the xe CLI.

VMs with Citrix VM Tools installed are shut down gracefully. However, VMs without Citrix VM Tools installed are shut down using a forced shutdown. To avoid forced shutdown, install the Citrix VM Tools on all VMs. For more information, see Install Citrix VM tools.

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

To shut down a server

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

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

Restart Toolstack

November 16, 2023

The **Restart Toolstack** option allows you to restart the Citrix Hypervisor management toolstack. This toolstack controls VM lifecycle operations, host and VM networking, VM storage, and user authentication. It allows the management of Citrix Hypervisor resource pools. The toolstack provides the publicly documented Management API, which is used by all tools that manage VMs and resource pools.

> **Note:**
>
> Sometimes, the **Restart Toolstack** option can be used for troubleshooting Citrix Hypervisor issues. However, be cautious when using this option, as incorrect usage can cause unexpected results.
>
> Do not restart the toolstack while HA is enabled. If possible, temporarily disable HA before restarting the toolstack.

To restart the toolstack

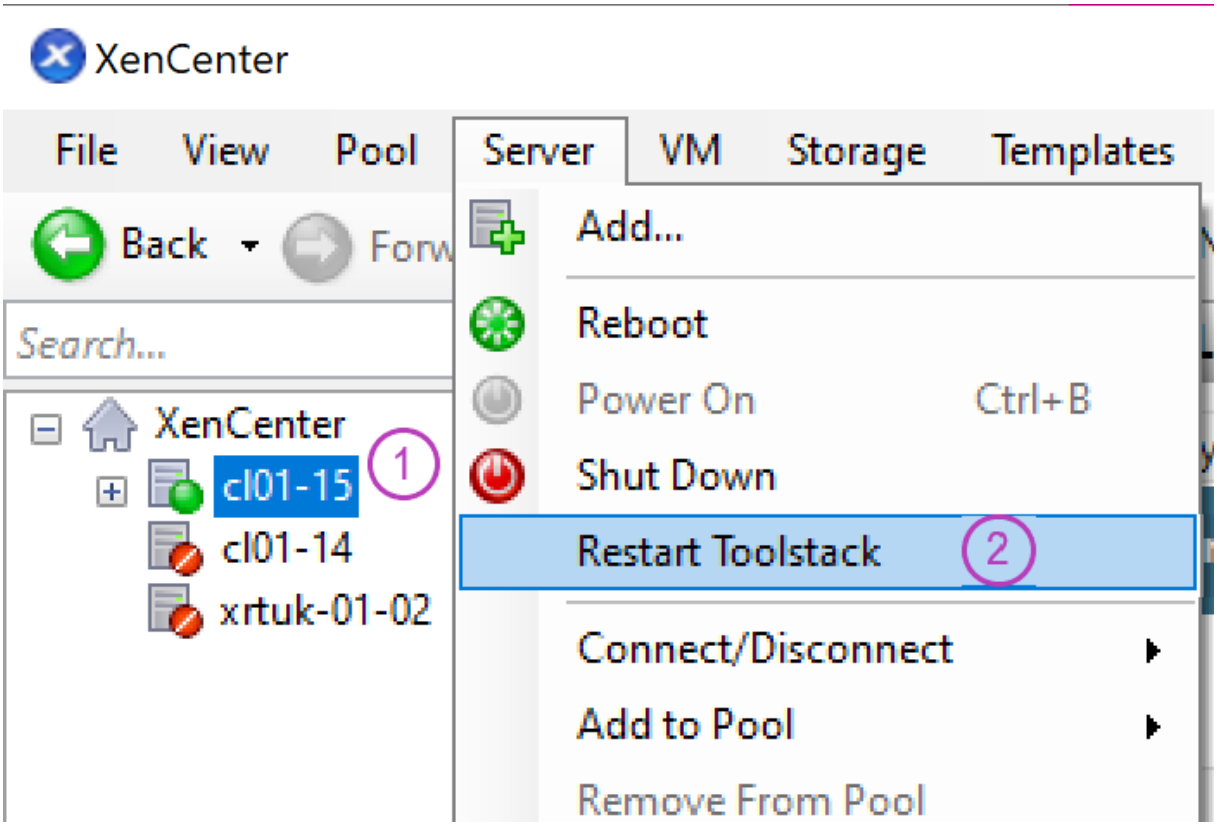

1. Select the server in the **Resources** pane.
2. On the **Server** menu, click **Restart Toolstack**.
3. Click **Yes** to confirm.

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

Configuring Host Power On

May 25, 2023

The Citrix Hypervisor 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 server supports remote power control. That is, the server has Wake on LAN functionality, a DRAC card, or it is using a custom script.
2. Enable the Host Power On functionality. To perform this procedure for DRAC processors, you need the credentials for the processor, which are set in its firmware.

After Host Power On has been configured on a server, you can power the server on from XenCenter. Select the server and then, on the **Server** menu, click **Power On**.

If you have installed and configured Workload Balancing, you can also configure Citrix Hypervisor to turn hosts on and off 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 Citrix Hypervisor with DRAC, 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.
- A custom power-on script based on the Management API that enables you to turn the power on and off through Citrix Hypervisor. For DRAC, 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 Hosts and Resource Pools.

To enable or disable Host Power On

You can enable Host Power On for an individual server by using the server **Properties** window, or on multiple servers by using the pool **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 disable the Host Power On feature.
   - **Wake on LAN (WOL)** To use this option, the host must have a Wake on LAN-enabled network card.
   - **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 the Citrix Hypervisor host from a remote location. For information about creating the script, including a list of supported key/value pairs, see the Hosts and resource pools.
3. If you selected 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 DRAC is configured.
   - **User name** This is the DRAC user name that is associated with the management processor. You might have changed this value 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 file name and path to the custom script you created. 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 file name 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

May 25, 2023

The Host Power On feature allows you to remotely power on managed servers from XenCenter. For servers to use this feature, they must fulfill the following criteria:

- The server has remote power control support: Wake-on-LAN functionality, a DRAC card, or a custom power-on script.
- You have enabled Host Power On in the server **Properties** settings. This feature can be enabled once for multiple servers at pool-level. For more information, see Configuring Host Power On.

After Host Power On has been configured, select the server or servers and then do one of the following:

- On the **Server** menu, select **Power On**.
- Right-click and select **Power On**.

Run in maintenance mode

May 25, 2023

You might want to take a managed server offline for various reasons. For example:

- To do a rolling upgrade of virtualization software
- To add or test connectivity to a new network
- To diagnose an underlying hardware issue
- To add connectivity to a new storage system.

Use XenCenter to take a server offline temporarily by placing it into *maintenance mode*. When you place a server in a resource pool into maintenance mode, all running VMs on it are automatically migrated to another server in the same pool. If the server is the pool master, a new master is also selected for the pool.

When Workload Balancing is enabled, it migrates the running VMs on that server to their optimal servers when available. These migrations are 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 choose **Enter Maintenance Mode** from the shortcut menu.

- On the **Server** menu, select **Enter Maintenance Mode**.

After all running VMs successfully migrate off the server, the server's status in the **Resources** pane changes to show the server maintenance mode icon. 

To take a server out of 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 choose **Exit Maintenance Mode** on the shortcut menu.
- On the **Server** menu, select **Exit Maintenance Mode**.

Install a TLS certificate on your server

November 16, 2023

The Citrix Hypervisor server comes installed with a default TLS certificate. However, to use HTTPS to secure communication between Citrix Hypervisor and Citrix Virtual Apps and Desktops, install a certificate provided by a trusted certificate authority.

> **Note:**
>
> This feature is supported only for Citrix Hypervisor 8.2 and later. If your Citrix Hypervisor server is an earlier version, XenCenter does not provide the option to install a new certificate on it.

This article contains information about how to use certificates in XenCenter. For information about working with certificates by using the xe CLI, see Hosts and resource pools.

Requirements

Ensure that your TLS certificate and its private key meet the following requirements:

- The certificate and key pair are an RSA key
- The key matches the certificate
- The key is provided in a separate file to the certificate
- The certificate is provided in a separate file to any intermediate certificates
- The key file must be one of the following types: `.pem` or `.key`
- Any certificate files must be one of the following types: `.pem`, `.cer`, or `.crt`
- The key is greater than or equal to 2,048 bits and less than or equal to 4,096 bits in length

- The key is an unencrypted PKCS #8 key and does not have a passkey
- The key and certificate are in base-64 encoded 'PEM' format
- The certificate is valid and has not expired
- The signature algorithm is SHA-2 (SHA256)

XenCenter warns you when the certificate and key you choose do not meet these requirements.

Install a certificate

You can use XenCenter to install a certificate that is on the XenCenter system into a Citrix Hypervisor server.

To install a certificate on a Citrix Hypervisor server, you must have the Pool Admin role and the Citrix Hypervisor server must not have HA enabled.

1. Go to the **Install Certificates** dialog. You can get to this dialog in one of the following ways:
   - In the **Server** menu, select **Install Certificates**.
   - Right-click on the server in the resources pane and choose **Install Certificates** from the context menu.
   - In the **General** tab of the server, right-click on the **Certificates** section and choose **Install Certificates** from the context menu.
2. In the **Install Certificates** dialog, browse to the location of the private key file and select it.
3. Browse to the location of the server certificate file and select it.
4. You can choose to add any number of intermediate certificates from the certificate chain.
   a) Click **Add**
   b) Browse to the location of one or more intermediate certificates and select them.
5. Click **Install**.

   XenCenter validates and installs the certificates.

   - If there is a problem with a certificate, XenCenter shows an error message. Attempt to correct the problem and click **Install** again.
   - If the certificate is installed successfully, XenCenter shows a success message. You can now click **Close** to close the dialog.

When the certificate on a Citrix Hypervisor server is changed, the server closes any open connections. XenCenter expects this behavior and reopens the connection with the Citrix Hypervisor server. However, you might have to manually reopen any other connections that were previously open to the server - for example, from another API client or the remote xe CLI.

View certificate information

In the **General** tab for a Citrix Hypervisor server, a section called **Certificates** displays the following information for the server:

- The certificate validity period. This text appears red when the certificate is approaching its expiry date.
- The certificate thumbprint

Certificate alerts

When your certificates are nearing their expiry date, XenCenter shows alerts in the **Alerts** section of the **Notifications** tab. You can choose to open the **Install Certificates** dialog from the action menu of these alerts.

For more information about alerts, see XenCenter Alerts.

Store Your Server Connection State

November 16, 2023

Login credentials - your user name 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 you enable this feature, XenCenter remembers the connection state of all your managed servers at the end of each session. XenCenter attempts to restore the servers at the start of your next session. If a server was connected at the end of your previous session, it is reconnected automatically without prompting you for your server login details. If a server was disconnected at the end of your previous session, it is not reconnected automatically.

If you disable the automatic reconnection feature, you must reconnect to all your managed servers each time you open XenCenter. You then enter your user name and password for each server.

> **Note:**
>
> Your system administrator can disable the saving of server login credentials, so this feature might 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 are 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 the **Save and Restore** tab.
3. Select or clear the **Save and restore server connection state on startup** check box.

Using a master password

When you choose to store login credentials in XenCenter, you can also set a master password. You must enter this master password 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.

If you lose or forget the master password, it cannot be recovered. You must 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 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 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 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

November 16, 2023

You can back up a managed server to one Citrix Hypervisor backup file (`.xbk`). This backup file can be used to restore the server if there is a hardware failure.

> **Note:**
>
> This file backs 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 or software failure. When backing up servers in XenCenter, note the following points:

- Do not create the backup on Citrix Hypervisor control domain (dom0). For more information about Citrix Hypervisor control domains, see the product documentation.
- Citrix Hypervisor backup files can be large.

To restore a server, you can select and restore the backup file within XenCenter. Reboot the server from Citrix Hypervisor installation ISO to complete the restore.

To back up your server configuration and software

1. Select the server in the **Resources** pane.
2. On the **Server** menu, click **Back Up**.
3. Browse to locate the folder where you want to create the backup file and enter the file name.
4. Click **Save** to begin the backup.

The backup might take some time. You can select 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.
2. On the **Server** menu, click **Restore From Backup**.
3. Browse to locate the backup file.
4. Click **Open** to begin the restore.
5. On the server, reboot to the installation CD and select **Restore from backup**.

Remove a Server From XenCenter

May 25, 2023

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 breaks the connection between XenCenter and the server and its VMs. The server is no longer displayed in XenCenter.

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

To return a server that you removed to the list of managed resources, add it again to XenCenter in the same way as you first connected to it. For more information, see Add a server.

Configuring Networking

July 11, 2023

Each managed server has one or more networks. Citrix Hypervisor networks are virtual Ethernet switches that can be connected to an external interface or can be entirely virtual, internal to an individual server or pool. The external interface can be with or without a VLAN tag.

When the Citrix Hypervisor product 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 VM (VIF) and a physical network interface (PIF) associated with a NIC on the server.

When you move a managed server into a pool, these default networks are merged and physical NICs with the same device name are attached to the same network. Typically, you add a network in the following cases:

- to create an internal network
- to set up a new VLAN using an existing NIC
- 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 traffic on storage networks and VM networks. 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. The possible MTU value range is 1500–9216.

Network types

There are three different physical network types to choose from when creating a network within XenCenter.

Single-Server private network

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

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. The bridge enables VMs to connect to external resources through the server's physical NIC.

Bonded network

This type of network bonds 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 network traffic for the host automatically routes over the second NIC. This mode provides load balancing of virtual machine traffic across the physical NICs in the bond.

- **Active-passive**

  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

- Configure vSwitch as the network stack to be able to view the LACP bonding options in XenCenter and to create a 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. However, you can only bond two NICs when a Linux bridge is the network stack.

For more information about the support for NIC bonds in Citrix Hypervisor, see Networking.

Add a network

November 16, 2023

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, select the **Networking** tab and then click **Add Network**.

To add an external network

An external network has an association with a physical NIC and provides a bridge between virtual machines and your external network. This bridge enables 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 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 1500–9216.
    d) To create a VLAN on an SR-IOV network, choose the NIC on which SR-IOV is enabled (Step 4a). Check the **Create the VLAN on the SR-IOV network** check box.
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 single-server private network

A single-server private network is an internal network that has no association with a physical network interface. It provides connectivity only between the VMs on a given server. This network has 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. This selection ensures that the new network is 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 bonded network

This type of network bonds 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 when you initially create your resource pool and before joining more servers to the pool or creating VMs. The bond configuration is automatically replicated to servers as they joined the pool. This action 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 can 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. With this bond, traffic is balanced between the bonded NICs. If one NIC within the bond fails, the 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 only becomes 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. With this bond, 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 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. This bond uses the source IP address, source port number, destination IP address, and destination port number 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:**
>
> - To be able to view the LACP bonding options in XenCenter and to create a LACP bond, configure vSwitch as the network stack. 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. However, you can only bond two NICs when the Linux bridge is the network stack.

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

To add an SR-IOV Network

Single Root I/O Virtualization (SR-IOV) is a PCI device virtualization technology that allows a single PCI device to appear as multiple PCI devices on the physical PCI bus. The physical device is known as a Physical Function (PF). The others are known as Virtual Functions (VF). SR-IOV enables the hypervisor to directly assign one or more of these VFs to a Virtual Machine (VM) using SR-IOV technology. The guest can then use the VF as any other directly assigned PCI device.

1. Open the **New Network** wizard.
2. On the first page of the wizard, choose **SR-IOV Network** and then click **Next**.
3. Enter a name and an optional description for the new network, and then click **Next**.
4. Choose a NIC from the list. NIC0 is not available in the list.

5. 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.
6. Click **Finish**.

   Creating an SR-IOV network affects network connection status. XenCenter connections to the pool can be temporarily disturbed.
7. Click **Create SR-IOV anyway** to create the network and close the wizard. The network created appears in the **NICs** tab indicating the number of VFs remaining or if it is disabled.

Remove a Network

May 25, 2023

1. Select the server or pool in the **Resources** pane.
2. Select the **Networking** tab.
3. On the **Networking** tab, select the network in the list.
4. Click **Remove Network**.

View and Change Network Properties

May 25, 2023

To view a server's current networking configuration

Select 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 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 is 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 address of the network adapter (NIC). This value is a unique identifier for a particular network adapter.
MTU	A Maximum Transmission Unit value between 1500–9216 allows the use of jumbo frames.

To change a server's networking configuration

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

Bond mode

This configuration option appears on bonded networks only.

- Select **Active-active** to configure an active-active bond. With this bond, traffic is balanced between the bonded NICs. 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 only becomes 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. With this bond, 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. This bond uses the source IP address, source port number, destination IP address,

and destination port number 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:**
>
> - To be able to view the LACP bonding options in XenCenter and to create a LACP bond, configure vSwitch as the network stack. 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 1500–9216.

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

November 16, 2023

Citrix Hypervisor 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 can 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, the server's network traffic automatically routes over the second NIC.

Active-passive: 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:**
>
> Configure vSwitch as the network stack to be able to view the LACP bonding options in XenCenter and to create a 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 Networking.

When you bond separate NICs using XenCenter, a new NIC is created. This NIC is the bond master, and the other NICs are referred to as the bonded NICs. The NIC bond can then be connected to Citrix Hypervisor 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. Use the 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 NIC 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 physical interface on your server, for example, a new Ethernet controller, it might not appear in the list on the **NICs** tab. If this situation happens, click **Rescan** on the **NICs** tab to force the server to scan for new cards.

To create a NIC bond

1. Ensure that the NICs you want to bind together are not in use. Shut down any VMs with virtual network interfaces using the bonded NICs before creating the bond. After you have created the bond, reconnect the virtual network interfaces to an appropriate network.
2. Select the server in the **Resources** pane on the left, then click 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 can 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. However, 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. With this bond, traffic is balanced between the bonded NICs. 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 only becomes 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. With this bond, 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. This bond uses the source IP address, source port, destination IP address, and destination port 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 Citrix Hypervisor, see the Networking.

5. To use jumbo frames, set the Maximum Transmission Unit (**MTU**) to a value between 1500–9216.
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 automatically moves management and secondary interfaces from bonded NICs to the bond master when the new bond is created.

A server with its management interface on a bond is not permitted to join a pool. 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 before 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) are dropped.

To delete a bond

1. Select the server in the **Resources** pane on the left, then click 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

November 16, 2023

The NIC used as the management interface on a managed server is initially specified during Citrix Hypervisor installation. XenCenter, xe CLI, and any other management software that runs on separate machine use the IP address of the management interface to connect to the server.

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 and dedicate NICs to a specific function, such as storage or other types of network traffic.

When a new server joins a pool, the new server inherits the pool master's networking configuration, including network and bond information. However, the joining server's management interface is not changed to match the master. Reconfigure it after joining to use the same bond as the pool master server.

> **Note:**
>
> A server with its management interface on a bond is not permitted to join a pool. Reconfigure the server's management interface and move it back on to a physical NIC before it can join a pool.

In XenCenter, use the **Configure IP Addresses** dialog to assign an IP address to a NIC and change the management interface for a server or pool. The following sections provide instructions for completing these actions.

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 are 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, establish separate networks for server management, application production traffic, and storage traffic. In the default Citrix Hypervisor networking configuration, all network traffic to IP-based storage devices occurs over the NIC used for the management interface. It is important to note that the secondary interfaces inherit the DNS server settings from the management interface.

To assign an IP address to a NIC, to carry out a specific function, ensure the appropriate network configuration is in place to ensure the NIC is used for the desired traffic. For example, to dedicate a NIC to storage traffic, assign the newly created interface an IP address that fulfills the following criteria:

- The IP address is on the same subnet as the storage controller, if applicable.
- The IP address is on a different subnet than the management interface.
- The IP address is not on the same subnet as any other secondary interfaces.

Ensure that you configure the NIC, storage target, switch, and VLAN so that the target is only accessible over the assigned NIC. This action allows the 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**, select **Configure**.
2. Click **Add IP address**.

3. Enter a name for the new secondary interface.
4. Choose your network from the **Network** list.
5. Configure the networking settings for the new interface:
   - To use the automated 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, select **Use these settings** and enter the required values. Enter an IP address and a subnet mask. You can optionally enter the gateway settings.
6. To configure extra interfaces, select **Add IP address** again and repeat the preceding configuration steps.
7. When you have finished, click **OK** to save your configuration choices.

> **Note:**
>
> If you choose to configure the network settings manually, you are 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**, select **Configure**.
2. From the list of configured interfaces, 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**, select **Configure**.
2. On the **Primary** tab, choose your network from the **Network** list.

> **Note:**
>
> The tagged VLAN networks are also displayed in this Network list.

3. Configure the networking settings for the management interface:
   - To use automated DHCP to assign network settings automatically, including the IP address, subnet mask, gateway and DNS server, select **Automatically obtain network settings using DHCP**.

- To configure network settings manually, select **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 are prompted to confirm your settings. Click **Reconfigure anyway** to confirm.

Changing Server Properties

May 25, 2023

Select any connected server in the **Resources** pane and select the **General** tab to see its properties and 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 **Properties** dialog.

- 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, select **Change** in the **Folder** box and select a folder. For more information, see 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). Citrix Hypervisor 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 Citrix Hypervisor support for iSCSI storage, see Storage.

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

To change the iSCSI IQN value for a managed server

> **Note:**
>
> Before you change a server's iSCSI IQN value, all existing SRs must be detached. Changing the server IQN might 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, select 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

Custom fields allow you to add information to managed resources to make it easier to search and organize them. For more information, see Using custom fields.

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 the email notification for system alerts generated on a standalone server. This feature is configured at pool level for servers in a pool. For more information, see XenCenter Alerts.

Multipathing

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

For more information, see Storage Multipathing.

Power on (standalone servers)

Use this tab to configure your Citrix Hypervisor 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

Citrix Hypervisor system log messages are stored locally on the server itself. You can also choose to forward these logs to a remote server.

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

To specify a remote Citrix Hypervisor log destination

1. Select the server in the **Resources** pane, select the **General** tab, and click **Properties**.
2. Select the **Log Destination** tab in the **Properties** dialog box.
3. Select **Also store the system logs on a remote server**.
4. In the **Server** field, enter an IP address or the host name of a server running the syslogd daemon.
5. Click **OK** to save your changes and close the dialog box.

GPU] 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.

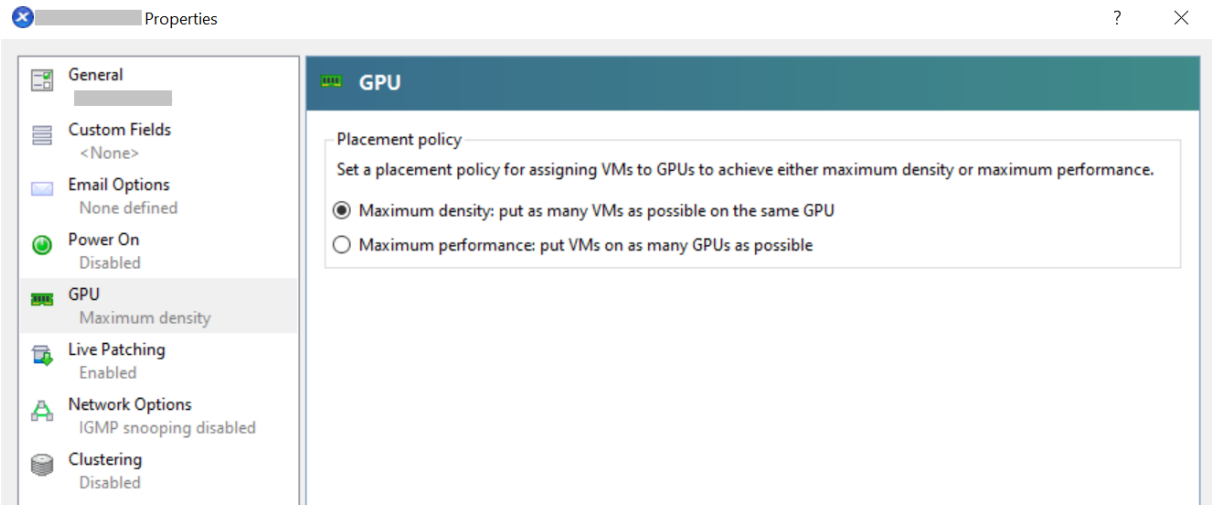

The **GPU** tab displays **Mixed** setting only when different settings are used for different GPU groups. For a **Mixed** setting, certain GPU groups are configured to achieve maximum density, and the rest are configured to achieve maximum performance.

It is **not** possible to set or edit the **Mixed** setting using XenCenter. Use the xe CLI if you want to use different settings for different GPU groups.

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

Integrated GPU pass-through

When your Citrix Hypervisor server is connected to an Intel GPU on an Intel server, the control domain of the server is connected to the integrated GPU device. In such cases, the GPU is not 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 Citrix Hypervisor server's external console output (for example, VGA, HDMI, DP) is not avail-

> able after disabling the connection between dom0 and the GPU.

Changing the Control Domain Memory

November 16, 2023

The control domain, also known as 'dom0', is a secure, privileged Linux Virtual Machine (VM) that runs the Citrix Hypervisor management toolstack (XAPI). The control domain provides the Citrix Hypervisor management function. It 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 Citrix Hypervisor installation. The amount is based on the amount of physical memory on the server. For more information, see Memory usage.

You might want to increase the memory that is allocated to dom0 in the following cases: Storage Read Caching scenarios, PVS-Accelerator scenarios, or when running more than 50 VMs per Citrix Hypervisor server. On servers with a smaller amount of memory, you might want to reduce the memory that is 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, see Command line interface.

> **Note:**
>
> - Citrix recommends that you do not reduce the dom0 memory below 1 GiB.
> - Increasing the amount of dom0 memory causes less memory being available to VMs.
> - Customers cannot use XenCenter to reduce dom0 memory below the value that was initially set during Citrix Hypervisor installation.

To update the dom0 memory

> **Note:**
>
> 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. This information includes available memory, dom0 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. Any change to the dom0 memory causes the server to reboot.
4. Click **OK** to confirm the changes and reboot the server.

Exporting and Importing a List of Managed Servers

May 25, 2023

You can export your list of managed servers from XenCenter to a configuration file. You can import this configuration file into a XenCenter session on another computer. This feature can be useful, for example, to copy your list of managed servers from your desktop computer to a laptop. You can avoid having to manually adding a long list of servers on the new machine.

XenCenter saves the IP address or DNS name, port, and display name of each managed VM in XML format in a file with a `.config` file name extension. Your login credentials are not stored.

To export your list of managed servers

1. On the **File** menu, select **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, select **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

May 25, 2023

Citrix Hypervisor pools allow you to view multiple servers and their connected shared storage as a single unified resource. Use this view to deploy VMs based on their resource needs and business priorities. A pool can contain up to 64 servers running the same version of Citrix Hypervisor software,

at the same patch level, and with broadly compatible hardware. For more information, see Pool Requirements.

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

If the pool master shuts down, the pool is unavailable until the master is back online or until you nominate another pool member as the new pool master. Every member of a resource pool contains all the information necessary to take over the role of master, if necessary. On an HA-enabled pool, a new pool master is automatically nominated if the master is shut down.

Pool Requirements

November 16, 2023

A resource pool is a homogeneous or heterogeneous aggregate of one or more servers, up to a maximum of 64. Before you create a pool or join a server to an existing pool, ensure the following requirements are satisfied for all servers in the pool.

Hardware requirements

All servers in Citrix Hypervisor resource pools 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, there are several other prerequisites for a server that joins a pool:

- It must have a consistent IP address (a static IP address on the server or a static DHCP lease). This requirement 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 cannot be a member of an existing resource pool.
- It cannot have any running or suspended VMs or any active operations in progress on its VMs. All VMs must be shut down before a server can join a pool.
- It cannot have any shared storage already configured.

- It cannot have a bonded management interface. Reconfigure the joining server's management interface and move it back on to a physical NIC before joining the pool. After the server has successfully joined the pool, you can reconfigure it. For more information, see Configuring IP Addresses.
- It must be running the same version of Citrix Hypervisor software, at the same patch level, as the servers already in the pool.
- It must be configured with the same supplemental packs as the servers already in the pool. Supplemental packs are used to install add-on software into dom0 (Citrix Hypervisor control domain). To prevent inconsistencies in the user experience across a pool, ensure you install the same supplemental packs at the same revision on all the servers in the pool.
- It must have the same Citrix Hypervisor license as the servers already in the pool. For example, you cannot add a server with Citrix Hypervisor Standard Edition license to an existing resource pool that contains servers with Citrix Hypervisor Premium Edition. 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 Citrix Hypervisor Licensing.

Homogeneous pool

A homogeneous resource pool is an aggregate of servers with identical CPUs. In addition to the requirements in the preceding sections, a server that joins a homogeneous pool must have the same CPUs as those servers already in the pool. CPUs are considered the same if they have the same vendor, model, and features.

Heterogeneous pool

Citrix Hypervisor 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 applying 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 capability enables you to create pools of hosts with disparate CPUs but still safely support live migrations. Servers joining heterogeneous pools must meet the following requirements:

- The CPUs of the server that joins the pool must be of the same vendor (AMD, Intel) as the CPUs on servers already in the pool. However, the specific type of CPU (family, model, and stepping numbers) is note required to be the same.
- The CPUs of the server joining the pool must support either Intel FlexMigration or AMD Enhanced Migration.

Citrix Hypervisor simplifies the support for heterogeneous pools. You can add servers 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 continues 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. If the pool level drops when a less-capable server joins the pool, a running VM can migrate to any server in the pool, except the newly added server. When you move or migrate a VM to a different server within or across pools, Citrix Hypervisor compares the VM feature set to that of the destination server. If the feature sets are found to be compatible, the VM is allowed to migrate. This capability enables the VM to move freely within and across pools, regardless of the CPU features the VM is using. If you use Workload Balancing to choose an optimal destination server to migrate your VM, a server with an incompatible feature set is not recommended as the destination server.

> **Note:**
>
> To update a running VM to use the pool's new feature set, power off the VM and start it again. 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 requirement for creating a resource pool, the advantages of pools are only available if the pool has one or more shared storage repositories (SRs). These advantages include running a VM on the most appropriate server and VM migration between servers.

We recommend that you do not attempt to create a pool until shared storage is available. After you add shared storage, 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 have to create a new shared SR for the pool: see Creating a New SR.

Create a New Pool

May 25, 2023

Before you attempt to create a pool, ensure that the requirements identified in Pool requirements are satisfied for all the servers that are in the new pool.

To create a 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 is 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 can add it to the list by clicking **Add New Server**. If a managed server is not listed, it might be because it does not satisfy one or more of the pool join requirements listed in Pool requirements.
5. Select **Create Pool** to create the 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 have to create a new shared SR for the pool. For more information, see Creating a New SR.

More 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

November 16, 2023

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

> **Important:**
>
> 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, select **Add to Pool** and then select the target pool.
- Right-click and select **Add to Pool** on the shortcut menu. Select 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:

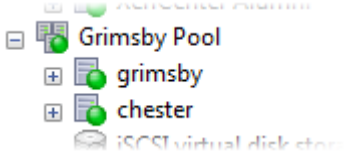

When you add a server to a pool, XenCenter attempts to resolve any pool configuration issues if possible:

- 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, if you add a server with Standard Edition license to a pool whose master is licensed with Premium Edition, you are prompted to upgrade the joining server license to match the master 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 Citrix Hypervisor 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.

There might be other hardware or configuration issues that prevent a server from successfully joining a pool. For more information, see Pool requirements.

When a new server joins a pool, that server automatically inherits the pool master's networking configuration, including network and bond information. However, the joining server's management interface does not change to match the master. Reconfigure it after joining to use the same bond as the pool master. For more information, see To change the management interface.

To place a server in a new pool

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

1. In the **Resources** pane, select the server.

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

Remove a Server From a Pool

November 16, 2023

1. Move any data stored on local disks to a shared storage repository in the same resource pool. For more information, see Move virtual disks.

   > **Important:**
   >
   > When you remove a server from a resource pool, all VM data stored on local disks is erased. Ensure you complete this step to retain any important data.

2. Shut down any VMs running on the server. For more information, see Shut Down a VM.
3. In the **Resources** pane, select the server and do one of the following:

   - Right-click and select **Remove Server from Pool** in the **Resources** pane shortcut menu.
   - In the **Pool** menu, select **Remove Server**.

Destroy a server from a pool

May 25, 2023

> **Important:**
>
> Destroying a server from a resource pool forgets the specified Citrix Hypervisor server without contacting it explicitly. It permanently removes the server 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, you cannot undo the destroy server operation. Reinstall the server before it can be used again.

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, select **Destroy**.

2. Click **Yes, Destroy** to confirm.

Export Resource Data

November 16, 2023

Export Resource Data enables you to generate a resource data report for your pool and export the report into an `.xls` or `.csv` file. This report provides detailed information about various resources in the pool such as:

- servers
- networks
- storage
- virtual machines
- VDIs
- GPUs

Use this feature to track, plan, and assign resources based on various workloads such as CPU, storage, and Network.

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

To export resource data:

1. In the XenCenter **Navigation** pane, select **Infrastructure** and then click the pool.
2. From the XenCenter menu, select **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 is available only if there are GPUs attached to your Citrix Hypervisor host.

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

Change Pool Properties

January 2, 2024

Select any resource pool in the **Resources** pane and select the **General** tab to see its properties and 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 add the pool to a folder or to move it to a different folder, select **Change** in the **Folder** box. Choose a folder. For more information, see 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 the email notification for system alerts that are generated on any servers or VMs in the pool. For more information, see XenCenter Alerts. Users who want to receive performance alert emails can choose the preferred language from the list. The languages available are English, Chinese, and Japanese.

The default language for configuring performance alert email language for XenCenter is English.

Power on

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

- In the list of servers at the top of the tab, select the servers for which you want to configure power management.
- Under **Power On mode**, specify the **Power On** settings (Disabled, Wake-on-LAN, 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. The options you must specify depend on the **Power On mode** option you chose.

For more information about prerequisites and configuration options for the Host Power On feature, see Configuring Host Power On.

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 you configure certain GPU groups within a pool for maximum density, and configure the other GPU groups for maximum performance. It is **not** possible to set or edit the **Mixed** setting using XenCenter. To use different settings for different GPU groups, use the xe CLI.

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

Security

The **Security** tab enables you to specify a security protocol to use for communication with the pool.

- **TLS 1.2 only**: This option accepts communication from Management API clients and appliances (including third-party appliances) that can communicate with the Citrix Hypervisor pool using the TLS 1.2 protocol. The **TLS 1.2 only** option uses the following cipher suites:
- `TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384` (Citrix Hypervisor 8.2 and later)
- `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384` (Citrix Hypervisor 8.2 and later)
- `TLS_RSA_WITH_AES_256_CBC_SHA256` (Citrix Hypervisor 8.2 and later)
- `TLS_RSA_WITH_AES_128_CBC_SHA256`

  > **Important:**
  >
  > Ensure that all Management API clients and appliances that communicate with the Citrix

> Hypervisor pool are compatible with TLS 1.2 before choosing this option.
>
> In Citrix Hypervisor 8.2 and later, this option is the only option provided.

- **Backwards compatibility mode (TLS 1.2 and earlier protocols)**: (Citrix Hypervisor 8.1 and earlier) Choose this option to allow both TLS and SSL protocols for pool-wide communication. For example, you might need both protocols for backward compatibility reasons. This option uses the following cipher suites 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:**
>
> Citrix Hypervisor Live Patching is available for Citrix Hypervisor Premium Edition customers, or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement. For more information about licensing, see About Citrix Hypervisor Licensing.

Network options

This tab allows you to enable or disable IGMP snooping. Citrix Hypervisor sends multicast traffic to all guest VMs. This behavior leads to unnecessary load on host devices by requiring them to process packets they have not solicited. If you enable IGMP snooping, it prevents hosts on a local network from receiving traffic for a multicast group that they have not explicitly joined. This action improves the performance of multicast. IGMP snooping is especially useful for bandwidth-intensive IP multicast applications such as IPTV. This option is disabled by default.

> **Note:**
>
> - IGMP snooping is available only when the network back-end uses Open vSwitch.
> - When enabling this feature on a pool, it might also be necessary to enable IGMP querier on one of the physical switches. Or else, multicast in the sub network falls back to broadcast and can decrease Citrix Hypervisor performance.
> - When enabling this feature on a pool running IGMP v3, VM migration or network bond failover results in IGMP version switching to v2.
> - Citrix Hypervisor IGMP snooping is available for Citrix Hypervisor Premium Edition customers, or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement. For more information about licensing, see About Citrix Hypervisor Licensing.

Clustering

This tab allows you to enable or disable clustering. Enable clustering on a pool to use thin provisioned storage repositories with GFS2.

> **Note:**
>
> Citrix recommends that you apply clustering only on pools that contain three or more servers and a GFS2 SR.
>
> Do not enable clustering on pools that don't include a GFS2 SR.

When enabling this feature on a pool, specify a network. The clustering mechanism uses this network to communicate with all servers in the pool. If a server cannot communicate with most other servers in the clustered pool, after a timeout that server self-fences. To decrease the chance of a host self-fencing unnecessarily, ensure the network you use for clustering is reliable. Citrix recommends you use a physically separate bonded network. For more information, see Add a new network.

Pool security

May 25, 2023

Change the root password

You can change the root password for a pool –that is, for all servers in a pool –by completing the following steps:

1. In the **Resources** pane, select the pool or any server in the pool
2. On the **Pool** menu or on the **Server** menu, select **Change Server Password**

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 is remembered. For more information, see Store your server connection state.

When you change the root password you are also advised to rotate the pool secret.

Rotate the pool secret

The pool secret is a secret shared among the servers in a pool that enables the server to prove its membership to a pool. Users with the Pool Admin role can view this secret when connecting to the server over SSH. Rotate the pool secret if one of these users leaves your organization or loses their Pool Admin role.

You can rotate the pool secret for a pool, by completing the following steps:

1. In the **Resources** pane, select the pool or any server in the pool
2. On the **Pool** menu, select **Rotate Pool Secret**

When you rotate the pool secret you are also prompted to change the root password.

Delete a Pool

May 25, 2023

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

May 25, 2023

Citrix Hypervisor storage repositories (SR) are storage containers on which virtual disks are stored. Both storage repositories and virtual disks are persistent, on-disk objects that exist independently of

Citrix Hypervisor. SRs can be shared between servers in a resource pool and can exist on different types of physical storage device, both internal and external. These devices include local disk devices and shared network storage. Various different types of storage are available when you create a storage repository using the **New Storage Repository** wizard. Depending on the type of storage selected, several advanced storage features can be configured in XenCenter. These features include:

- **Dynamic multipathing**. On Fibre Channel and iSCSI SRs, you can configure storage multipathing using round robin mode load balancing. For more information, see Storage Multipathing.
- **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. This behavior guarantees that virtual disks never run out of space and therefore 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. This behavior allows for over provisioning of the available storage and maximum utilization of your storage assets.
- **Reclaiming Freed Space**. On a thinly provisioned block-based SR, you can free up unused space (for example, deleted VDIs in a LUN). The storage repository can then reuse the reclaimed space. 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. This SR is 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. This behavior provides 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 automatically becomes the default SR for that server.

It is possible to use different SRs for VMs, crash dump data and suspended VM using Citrix Hypervisor xe CLI. For more information, see Command line interface.

Creating a New SR

May 25, 2023

To create a storage repository, select **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**.

Select the physical storage type on the first page of the **New Storage Repository** wizard. 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 following links to find out more about creating different types of SR.

SR creation steps

The **New Storage Repository** wizard takes you through the process of creating an 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 migrate between servers in the same resource pool.
   - iSCSI:
     Software 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. Complete the configuration required to expose the LUN before running the **New Storage Repository** wizard: the wizard automatically probes for available LUNs and displays a list of all the LUNs found.
   - SMB Storage:
     SMB servers are a common form of Windows shared filesystem infrastructure. These servers 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.
   - Software FCoE (deprecated):
     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 feature eliminates the need for using expensive HBAs. Before you use the New Storage Repository wizard to create a Software FCoE storage, manually complete 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, the wizard automatically generates a description of the SR, 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. If you select iSCSI or Hardware HBA as your storage type, the wizard displays the **Provisioning** page. Select the type of provisioning to use for this SR. The options available are

   - Thin provisioning (GFS2). This type of provisioning is only available on clustered pools. For more information about clustering, see Change pool properties
   - Full provisioning (LVM)

4. 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 depend on the type of storage you selected on the first page of the wizard.
5. Click **Finish** to create the SR and close the wizard.

NFS Storage

May 25, 2023

In an NFS storage repository (SR), VM images are stored as thin-provisioned VHD format files on a shared NFS target. Existing NFS servers (that support any version of NFSv3 or NFSv4 over TCP/IP) can be used immediately as an SR for virtual disks.

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

Since virtual disks on NFS SRs are created as sparse, 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, the wizard generates a description of the SR. This description includes a summary of the configuration options that 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 is 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

November 16, 2023

Software iSCSI is supported using the open-iSCSI software iSCSI initiator or by using a supported iSCSI Host Bus Adapter (HBA).

Dynamic multipathing support is available for iSCSI storage repositories. By default, multipathing uses round robin mode load balancing, so both routes have active traffic on them during normal operation. You enable and disable storage multipathing in XenCenter by using the **Multipathing** tab on the server **Properties** dialog. For more information, 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, the wizard generates a description of the SR. This description includes 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 **Provisioning** page, select the type of provisioning to use. The options available are
   - Thin provisioning (GFS2). This type of provisioning is only available on clustered pools. For more information about clustering, see Change pool properties
   - Full provisioning (LVM)

   Click **Next** to continue.
5. On the **Location** page, specify the iSCSI target details:
   - **Target Host**: The IP address or DNS name of the iSCSI target. This can also be a comma-separated list of values.
   - **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 user name credential to apply when connecting to the target.
     - CHAP Secret: the CHAP authentication password credential to apply 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. Every iSCSI target and initiator must have a unique IQN. If a non-unique IQN identifier is

> used, data corruption can occur or access to the target can be denied or both.

- **Target LUN**: To specify the LUN on which to create the storage repository, click the **Discover LUNs** button. Choose a LUN from the **Target LUN** list.

  Each individual iSCSI storage repository must be contained entirely on a single LUN. The SR cannot span more than one LUN. If the LUN already contains an SR, choose either to use the existing SR or to replace the existing SR with a new one. Replacing the existing SR destroys any data present on the disk.

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

Hardware HBA Storage

November 16, 2023

Hardware HBA SRs connect to Fibre Channel (FC), Fibre Channel over Ethernet (FCoE) or shared Serial Attached SCSI (SAS) LUNs through an HBA. Do the configuration required to expose the LUN before you run the **New Storage Repository** wizard. The wizard automatically probes for available LUNs and displays 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. To open the **New Storage Repository** wizard, you can do any of the following actions:
   - On the toolbar, select **New Storage**.
   - On the **Storage** tab for the selected pool or server, select **New SR**.
   - On the **Storage** menu, select **New SR**.
   - In the **Resources** pane, select a server or pool then right-click and select **New SR** on the shortcut menu.
2. Select **Hardware HBA** as the physical storage type and then select **Next**.
3. On the **Name** page, enter the name of the new SR. By default, the wizard generates a description of the SR. This description includes 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 **Provisioning** page.

4. On the **Provisioning** page, select the provisioning type. The options available are
    - Thin provisioning (GFS2). This type of provisioning is only available on clustered pools. For more information about clustering, see Change pool properties
    - Full provisioning (LVM)

Click **Next** to continue to the **Location** page.

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 is displayed if there are existing SRs on the LUN you have selected. Review the details and choose one fo the following options.
>
> - To use the existing, click **Reattach**.
> - To delete the existing SR and to create an SR, click **Format**.
> - If you prefer to select a different LUN, click **Cancel** and select a LUN from the list.

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

SMB Storage

May 25, 2023

SMB servers are a common form of Windows shared filesystem infrastructure that you can use as a storage repository substrate for virtual disks. VM images in SMB servers are stored as thinly provisioned VHD files on an SMB target. As SMB servers are shared SRs, you can start VMs that have virtual disks in SMB servers on any server in a pool. These VMs readily migrate between the servers.

> **Note:**
>
> - SMB storage is available for Citrix Hypervisor Premium Edition customers, or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement. To learn more about Citrix Hypervisor licensing, see About Citrix Hypervisor Licensing.
> - When using SMB storage, do not remove the share from the storage before detaching the SMB SR.

To configure an SMB SR

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

2. Select **SMB** as the physical storage type, then click **Next**.
3. On the **Name** page, enter the name of the new SR. By default, the wizard generates a description of the SR. This description includes 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. Configure the SMB server to export the specified path to all servers in the pool.
   - **User name and Password** (Optional) 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.

Software FCoE Storage (deprecated)

November 16, 2023

Software FCoE provides a standard framework that hardware vendors can plug in their FCoE offload-capable NIC into. By using this framework, they get the same benefits of a hardware-based FCoE. This feature eliminates the need for using expensive HBAs. Software FCoE can be used with Open vSwitch and Linux bridge as the network back end.

Before creating a software FCoE storage, complete the configuration required to expose a LUN to the host. This process includes configuring the FCoE fabric and allocating LUNs to your SAN's public world wide name (PWWN). After you complete this configuration, the available LUN is mounted to the host's CNA as a SCSI device. You can then use the SCSI device 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, see the documentation provided by the vendor. For detailed information about Software FCoE, see the Storage

> **Note:**
>
> - Software FCoE is deprecated and will be removed in a future release.

> - When you configure the FCoE fabric, do not use VLAN 0. The Citrix Hypervisor host cannot find traffic that is on VLAN 0.
> - Software FCoE can be used when using Open vSwitch and Linux bridge as the network back end.

To create a Software FCoE SR

1. To open the **New Storage Repository wizard**, do one of the following actions:
   - On the toolbar, click **New Storage**.
   - 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, the wizard generates a description of the SR. This description includes 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. Choose the LUN (s) that you want to allocate to the SR and click **Next**.

   > **Note:**
   >
   > If the host cannot find any LUNs, an error message is 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

November 16, 2023

This type of SR can be useful for creating shared ISO libraries. For example, use it to create a library of VM installation images. The following 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, the wizard generates a description of the SR. This description includes 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 the SR uses.

     > **Note:**
     >
     > If the underlying storage array does not support NFSv4, NFSv3 is used to mount the share.
5. Click **Finish** to complete the new SR configuration and close the wizard.

Storage properties

April 10, 2024

Viewing storage properties

You can view the details for all storage repositories (SRs) in your Citrix Hypervisor pool from the **Storage** tab of the pool. Select a server or pool in the **Resources** pane and then click the **Storage** tab. This tab displays information about the local and shared storage repositories. This information includes the name, description, storage type, usage, the size of the SR, and the virtual allocation.

To view detailed information about an individual storage repository, select the SR repository in the **Resources** pane. In the main pane, the following tabs are available: **General**, **Storage**, and **Search**.

General

The **General** section of the **General** tab displays the information about the storage repository. This information can include the name, description, any tags applied to the SR, the folder the SR is in, storage type, the size of the SR, the SCSI ID, and the UUID of the SR.

> **Note**
>
> For GFS2 SRs, the size is displayed in the form "*a* GB used of *b* GB total (*c* GB allocated)". The values in this statement have the following meanings:
>
> - *c* is the space that is allocated and used in the SR. (If you run the `du` command, this is the value shown.)
> - *b* is the actual total size of the LUN.
> - *a* is the 'used' space. This value is the sum of the file sizes for the files on the SR and does not take into account blocks that have been freed from sparse files. These freed blocks can now be used by other VMs. (If you run the `df` command, this is the value shown.)
>
>   On GFS2 SRs, the VM disks are stored in the QCOW2 format, which is both thinly provisioned and sparse. When data blocks are deleted from the VM the associated data blocks in the VM disk file are freed back to the SR filesystem and marked as unallocated, but the file size is unchanged. This behavior can cause discrepencies between the value of *a* and the value of *c*.

The **Status** section lists the state of the SR and shows whether it is connected to servers in the pool.

The **Multipathing** section shows whether multipathing is active between the SR and servers in the pool.

Storage

The **Storage** tab lists the virtual disks located on the storage repository. For each disk, the table shows the disk name, description, size, virtual machine name, and whether changed block tracking is en-

abled.

Search

Using the **Search** tab, you can construct queries based on object types, folders, and attributes. For more information, see Searching Resources.

Changing SR properties

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

You can access the **Properties** dialog box for an SR in one of the following ways:

- Select a server or pool in the **Resources** pane and then click the **Storage** tab. This lists the local and shared storage in your pool. Select an SR from the list and click **Properties**.
- Select the storage repository in the **Resources** pane. In the **General** tab for the SR, 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. For more information, see 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. For more information, see Using custom fields.

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. Check **Generate storage throughput alerts** and set the storage throughput and time threshold that triggers the alert.

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

Read Caching

On the **Read Caching** tab, you can choose to enable or disable read caching.

Read caching improves performance on NFS, EXT3/EXT4, SMB, or GFS2 SRs that host multiple VMs cloned from the same source. You might want to disable read caching in the following cases:

- You have no file-based SRs
- You do not have any cloned VMs
- It is not providing any performance benefits

For more information, see Storage read caching.

Removing an SR

May 25, 2023

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 must temporarily take a storage repository offline, for example, for maintenance. A detached SR can then be reattached. For more information, see Reattaching an SR.
- **Forget**. When you forget an SR, the contents of the virtual disks on the SR are preserved. However, the information used to connect 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 about how to destroy an SR, refer to the Storage.

> **Note:**
>
> - You cannot remove a storage repository if it contains virtual disks of a currently running virtual machine.
> - When using SMB storage, do not remove the share from the storage before detaching the SMB SR.

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 is 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

May 25, 2023

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

To reattach a detached SR

1. Select the detached SR in the **Resources** pane and then do one of the following:
   - Right-click and select **Reattach Storage Repository** on the **Resources** pane shortcut menu.
   - On the **Storage** menu, select **Reattach Storage Repository**.
2. Enter the required configuration information in the same way as when you add an SR. See:
   - NFS VHD Storage
   - Software iSCSI Storage
   - Hardware HBA Storage
   - ISO Storage
3. Click **Finish** to complete the SR configuration and close the wizard.

Storage Multipathing

October 4, 2023

Dynamic multipathing support is available for Fibre Channel and iSCSI storage repositories. By default, multipathing uses round robin mode load balancing, so both routes 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 action 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 NICs must be configured to use the same subnet. For more information, see Configuring IP Addresses.)

For more in-depth multipathing information, see Multipathing.

You can use up to 16 paths to a single LUN.

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 is disconnected and might disappear from the **Resources** pane temporarily while a new pool master is assigned. When the server reappears in the **Resources** pane with the Server maintenance mode icon, continue to the next step.

   

2. On the **General** tab, select **Properties** and then select the **Multipathing** tab.
3. To enable multipathing, check 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 the server in the **Resources** pane, right-click, and select **Exit Maintenance Mode**.

Storage Read Caching

April 10, 2024

> **Note:**
>
> Storage Read Caching is available for Citrix Hypervisor Premium Edition customers, or those cus-

> tomers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement.

Read Caching improves a VM's disk performance as, after the initial read from an 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 as it drastically reduces the number of blocks read from disk. For example, Read Caching improves performance in Citrix Virtual Desktops Machine Creation Services (MCS) environments.

This performance improvement can be seen whenever data must be read from disk more than once, as it gets cached in memory. This performance difference 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 narrow time frame (boot storm)
- When a significant number of VMs are scheduled to run malware scans at the same time (antivirus storm)

> **Note:**
>
> The amount of memory assigned to the Citrix Hypervisor control domain (dom0) might 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.

Read Caching is enabled by default, provided:

- The Citrix Hypervisor host is licensed with Citrix Hypervisor Premium Edition or a Citrix Virtual Apps and Desktops license.
- The VM is attached to a VDI on a file-based SR such as NFS, EXT3/EXT4, or GFS2. 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, see Storage read caching.

To disable read caching

You can disable read caching for an SR in its **Properties** dialog.

1. In the **Resources** pane, select the SR that you want to disable read caching on.
2. In the **General** tab, click **Properties**.
3. In the **Properties** dialog, deselect **Enable Read Caching**.
4. Click **OK**.

To enable read caching

You can enable read caching for an SR in its **Properties** dialog.

1. In the **Resources** pane, select the SR that you want to enable read caching on.
2. In the **General** tab, click **Properties**.
3. In the **Properties** dialog, select **Enable Read Caching**.
4. Click **OK**.
5. Restart any VMs that you want to benefit from the changed setting.

PVS-Accelerator

May 25, 2023

The Citrix Hypervisor PVS-Accelerator feature offers extra capabilities for customers using Citrix Hypervisor and Citrix Provisioning (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 Citrix Hypervisor host. To benefit from the PVS-Accelerator feature, use Citrix Hypervisor with Citrix Provisioning 7.12 or higher. For detailed information about PVS-Accelerator, see the product documentation.

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

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

Enabling PVS-Accelerator

To enable the PVS-Accelerator feature, complete the following configuration settings in Citrix Hypervisor and in PVS:

1. Install the PVS-Accelerator Supplemental Pack on each Citrix Hypervisor host in the pool. The supplemental pack is available to download from the Citrix Hypervisor Product Downloads page. For instructions on how to install the supplemental pack, see Installing Supplemental Packs.
2. Configure PVS-Accelerator in Citrix Hypervisor. This configuration can be done using XenCenter or the xe CLI.

After installing the PVS-Accelerator Supplemental Pack, add the PVS-Accelerator configuration details in the Citrix Hypervisor server. This process entails adding a PVS site and specifying the PVS Cache storage.

The following section contains XenCenter instructions. For information about configuring the PVS-Accelerator using the xe CLI, see the Citrix Hypervisor product documentation.

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) Select the pool or the standalone host and then select the **PVS** tab.
b) Select **Configure PVS-Accelerator**.
c) On the **PVS-Accelerator configuration** dialog, select **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 to use:
     - When you select **Memory only**, the feature uses up to the specified cache size in the Control Domain memory. This option is only available after extra 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 uses up to the specified cache size on the SR. It also implicitly uses available control domain memory as a best effort cache tier.

   Important:

   - If neither Memory only nor an SR is specified, the read cache is not activated.
   - PVS-Accelerator has been designed to utilize either memory only, or a combination of disk and memory. Irrespective of the configuration choice, 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 4 GB of Control Domain memory per host to avoid frequent disk accesses that cause higher read-latency and therefore degrade performance. For more information, see Changing the Control Domain Memory.
   - We recommend that you allocate at least 5 GB of cache space per vDisk version that you actively use.
d) Click **OK**. The new PVS site and the chosen cache storage configuration are added in the Citrix Hypervisor server.

3. After configuring the PVS-Accelerator in the Citrix Hypervisor server, complete the cache configuration for the newly created site using the Citrix Provisioning Console or the PowerShell Snap-In CLI. For more information, see Citrix Provisioning Documentation. When this step is complete, you can view a list of PVS Servers configured for the new site by clicking **View PVS Servers** on the **PVS-Accelerator configuration** dialog.

Cache operation

After you start a VM with PVS-Accelerator, the caching status for the VM is 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 a vDisk but not writes or reads from a write cache
- **Based on image versions**. Multiple VMs share cached blocks if they use the same image version
- Devices with **any write cache type**

- A vDisk with the Access mode set to **Standard Image**. Caching is not compatible with any vDisk 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 Citrix Hypervisor Premium Edition customers, or customers who have access to Citrix Hypervisor 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 uses the capabilities of OVS and is therefore not available on hosts that use Linux Bridge as the network back end.
- PVS-Accelerator works on the first virtual network interface (VIF) of a cached VM. Therefore, use the first VIF to connect the PVS storage network for the caching to work.
- After upgrading the PVS-Accelerator supplemental pack, XenCenter might list multiple versions of the PVS-Accelerator. However, only the latest version is active. There is no need to uninstall PVS-Accelerator, as old versions of this feature are always superseded by the newest version.

Reclaiming Freed Space

May 25, 2023

Use the **Reclaim freed space** option in XenCenter you to free up unused blocks on a LUN that is thinly provisioned by the storage array. Once released, the array can then reuse the reclaimed space. The **Reclaim freed space** operation is only available for LVM-based SRs that are thinly provisioned on the array. These SR types are iSCSI, Fibre Channel, or Local LVM. This feature is not enabled on file-based SRs such as NFS and EXT3/EXT4 as these SR types do not require a manual space reclamation operation.

To reclaim freed space:

1. Select the **Infrastructure** view and then select the host or the pool that contains the SR.
2. Select the **Storage** tab.
3. Select the SR from the list and then select **Reclaim freed space**.

> **Note:**
>
> Reclaiming freed space is an intensive operation and can affect the performance of the storage array. 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, select **Notifications** and then **Events**.

Live LUN Expansion

May 25, 2023

To fulfill capacity requirements, you might need to add capacity to the storage array to increase the size of the LUN provisioned to your Citrix Hypervisor host. Use live LUN Expansion to increase the size of the LUN and use the newly gained space without detaching the SR or taking the hosts or VMs offline.

> **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 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.

Creating VMs

September 14, 2023

A virtual machine (VM) is a software container that runs on a host physical computer. The VM 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. This metadata is the setup information needed to create a VM with a specific guest operating system, and with the optimum storage, CPU, memory, and virtual network configuration.

You can create VMs in XenCenter in several different ways:

- The New VM wizard takes you step by step through the process of creating a VM from a template or a snapshot. This wizard enables you to configure the 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 the required VM configuration parameters. You simply select your preconfigured template in XenCenter then right-click and select 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.

Citrix VM Tools

VMs in a Citrix Hypervisor environment might be fully virtualized (HVM) or paravirtualized:

- In HVM (hardware-assisted virtualization or Hardware Virtual Machine) mode, the VM is fully virtualized. An HVM mode VM 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 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 *Update Linux Kernels and Guest Utilities* section in Linux VMs.

- 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.

  > **Note:**
  >
  > Paravirtualized (PV mode) VMs are only supported in Citrix Hypervisor 8.0 and earlier.

  For detailed information about supported guest operating systems, see Guest operating system support.

I/O drivers (also known as Paravirtualized drivers or PV drivers) are available for Windows and Linux VMs to enhance disk and network performance. Install these drivers on all new VMs and update them through the Windows Update mechanism. The I/O drivers and the Management Agent are combined and issued as **Citrix VM Tools** for ease of installation. For more information, see Installing Citrix VM Tools. Citrix Hypervisor features such as VM migration and historical performance data tracking are only available on VMs that have Citrix VM Tools installed.

Using templates

Several different templates are supplied with the Citrix Hypervisor server. These templates 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 VMs. See Guest OS support 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 Citrix Hypervisor templates supplied with the product and any custom templates that you create in the **Resources** pane.

-  Citrix Hypervisor template
-  Custom template

You can control whether to display the Citrix Hypervisor and Custom templates in the **Resources** pane:

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

  This panel displays a tree view of your managed resources in the **Resources** pane.
- To display standard Citrix Hypervisor VM templates: on the **View** menu, select **Citrix Hypervisor Templates**. To hide Citrix Hypervisor templates, select again to remove the check mark.
- To show custom VM templates: on the **View** menu, select **Custom Templates**. To hide custom templates, select again to remove the check mark.

Creating a New VM

May 25, 2023

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 see slightly different VM configuration options presented

on subsequent pages. 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 Citrix Hypervisor environments where Role-Based Access Control (RBAC) is implemented, the **New VM** wizard checks that you have a role with sufficient permissions to create VMs. If your RBAC role does not have sufficient permissions, for example, a VM Operator or Read-only role, you cannot continue with VM creation. For more information, see RBAC overview.

Overview of VM creation steps

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

1. Select a template.

   The first step is to choose a VM template. Templates contain the setup information to create a VM with a specific guest operating system, and with the optimum storage, CPU, memory, and virtual network configuration. Various different templates are supplied, and you can add custom templates of your own. For more information, see Template and BIOS options.

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 makes it easier for you to manage different VMs if you give them meaningful, memorable names. For more information, see VM name and description.

3. Specify the operating system installation media and boot mode.

   The third step in the process is to specify the type and location of the OS installation media and to choose a boot mode for the VM.

   Windows operating systems can be installed from an ISO library, from install media in a physical DVD/CD drive or from network boot. Linux operating systems can be installed from a network install repository, an ISO library, or from install media in a physical DVD/CD drive.

   You can now allow recent versions of Windows guest operating systems to boot in UEFI mode. For more information, see OS installation media.

   > **Note:**
   >
   > Guest UEFI boot is an experimental feature. You can create UEFI-enabled VMs on hosts that are in a production environment. However, UEFI-enabled VMs must not be used for production purposes. You might have to re-create the VMs when you upgrade the host to a newer version.

4. Choose a home server.

This step is optional, but you can choose a home server for the new VM. Citrix Hypervisor always attempts to start the VM on the nominated home server if it can. For more information, see Home server.

5. Configure CPU and memory.

   - **For Windows VMs**: You can assign virtual CPUs (vCPUs) to the new VM, specify the number of cores per socket to present to the vCPUs, and allocate memory. These settings can be adjusted later, after the new VM has been created. For more information, see CPU and memory configuration.
   - **For Linux VMs**: You can assign a maximum number of vCPUs to the VM, specify the number of cores per socket 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. For more information, see CPU and memory configuration.

6. Assign a graphics processing unit (GPU).

   The **New VM** wizard prompts you to assign a dedicated GPU or one or more virtual GPUs to the VM. This feature enables the VM to use the processing power of the GPU, empowering high-end 3D professional graphics applications such as CAD/CAM, GIS, and Medical Imaging applications. For more information, see GPU.

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

7. Configure storage.

   The next step is to configure virtual disks for the new VM. The wizard automatically configures a minimum of one virtual disk and the template you select might include more. For more information, see Virtual disk configuration.

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

   If you are creating a CoreOS VM, you are prompted to specify the cloud-configuration parameters for the VM. For more information, see Cloud-Config Parameters.

9. Configure networking.

   The last step in the process of provisioning a new VM is to configure networking. You can configure up to four virtual network interfaces on each VM. For more information, see Configure virtual network interfaces.

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 when it is created.

VM Template and BIOS Options

May 25, 2023

Various different VM templates are supplied with the Citrix Hypervisor server. These templates can be used in different ways to create VMs. Each template contains installation metadata: the information used to create a VM with a specific guest OS with the optimum storage, CPU, memory, and virtual network configuration.

For a full list of guest operating systems that Citrix Hypervisor supports, see Guest operating system requirements.

You can also create your own customized templates configured with the appropriate guest operating system, memory, CPU, storage, and network settings. For more information, see Creating new templates.

Select a BIOS option

Citrix Hypervisor VMs can be BIOS-generic or BIOS-customized:

BIOS-generic: the VM has generic Citrix Hypervisor BIOS strings;

BIOS-customized: HVM VMs support customization of the BIOS in two ways, namely: Copy-Host BIOS strings and User-Defined BIOS strings.

- Copy-Host BIOS Strings: The VM has a copy of the BIOS strings of a particular server in the pool.
- User-Defined BIOS Strings: The user can set custom values in selected BIOS strings using CLI/API.

> **Note:**
>
> If a VM does not have BIOS strings set when it starts, the standard Citrix Hypervisor BIOS strings are inserted into it, and the VM becomes BIOS-generic.

For more information, refer to Advanced VM information.

When you create a 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 action enables 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 is 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

May 25, 2023

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

You can more easily manage different VMs if you give them meaningful names. For example, include one of the following pieces of information in the VM name:

- The operating system of the VM (Windows 10 64-bit)
- The application software on the VM (Citrix Hypervisor Web Self-Service v1.0 (Build 9057))
- The role of the VM (`db-server`, Outlook Server, Test).

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

May 25, 2023

The options for OS installation media and boot mode available on the **Installation Media** page of the **New VM** wizard depend on the OS or template that you selected on the first page of the wizard.

OS installation media options

Install from ISO library or DVD drive

Templates: Windows and Linux PV and HVM guests

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

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

If there are no ISO images listed here, you need to make the ISOs available to the server by creating an external NFS or SMB/CIFS share directory.

Boot from network

Templates: Windows and Linux HVM guests

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

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

Install from URL

Templates: Linux PV guests

You can install PV versions of CentOS, SUSE Linux Enterprise Server, and Red Hat Linux operating systems 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:

```
nfs://server/path
ftp://server/path
http://server/path
<!--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 more operating system boot parameters, if necessary.

Boot mode

Choose a boot mode for the VM. Specify the boot mode when you create a VM. It is not possible to change the boot mode after booting the VM for the first time.

- Select **BIOS Boot** to boot the VM in legacy BIOS mode.
- Select **UEFI Boot** to boot the VM in UEFI mode.
- Select **UEFI Secure Boot** to boot the VM in UEFI Secure Boot mode.

The most secure boot mode is selected by default. You can only select those boot options that are available to your new VM.

UEFI boot and UEFI Secure Boot are supported only on newly created Windows 10 (64-bit), Windows Server 2016 (64-bit), Windows Server 2019 (64-bit), and Windows Server 2022 (64-bit) VMs. For more information, see Windows VMs.

Home Server

May 25, 2023

A home server is the server which provides the resources for a VM in a pool. When you nominate a home server for a VM, Citrix Hypervisor always attempts to start up the VM on that server if it can. If Citrix Hypervisor cannot start the VM on that server, an alternate server within the same pool is selected automatically.

- To nominate a home server for the VM in the **New VM** wizard, select **Place the VM on this server** and choose 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 uses 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 is automatically 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 does not take effect:

- In pools with Workload Balancing (WLB) enabled, the nominated home server is not used for starting, restarting, resuming, or migrating the VM. Instead, WLB nominates the best server for the VM by analyzing Citrix Hypervisor resource pool metrics and by recommending optimizations.
- If a VM has one or more virtual GPUs assigned to it, the home server nomination does not take effect. Instead, the server nomination is based on the virtual GPU placement policy set by the user. For more information, see GPU Placement Policy.

VM CPU and Memory Allocation

May 25, 2023

When you create a VM, you can allocate virtual CPUs, specify how many cores-per-socket to present to the vCPUs, and set initial memory resources for the VM. You can change the settings at any time 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 get the best performance out of the VM, the number of vCPUs assigned to the VM mustn't 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, see the Citrix Hypervisor Configuration Limits.

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 menu.

To get the best performance out of the VM, the number of vCPUs assigned to the VM mustn't 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, Citrix Hypervisor allocates one core per socket for each vCPU. For example, allocating 4 vCPUs appear as 4 sockets with 1 core per socket. Click the **Topology** menu to change this setting and choose 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. 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 choose from the 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 Citrix Hypervisor templates provide typical VM configurations and set reasonable defaults for the memory, based on the type of guest operating system. The following considerations can affect how much memory you decide to initially allocate to a new VM:

- The kinds of applications that run on the VM.
- Other virtual machines that use the same memory resource.
- Applications that 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. On this tab, you can also enable Dynamic Memory Control (DMC) to allow dynamic reallocation of memory between VMs in the same pool. For more information, see Configuring VM memory.

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

May 25, 2023

XenCenter allows you to assign a dedicated graphics processing unit (GPU) or one or more virtual GPUs to a new VM during VM creation. This feature enables a VM to use the processing power of the GPU, providing better support for high-end 3D professional graphics applications. For example, CAD/CAM, GIS, and Medical Imaging applications.

For detailed information, see Configuring graphics.

Citrix Hypervisor 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. The motherboard must have a chipset which enables GPU functionality, for example, C226 for Xeon E3 v4 CPUs or C236 for Xeon v5 CPUs. For information about supported processors, refer to the Citrix Hypervisor Hardware Compatibility List.

The following table lists whether GPU, shared GPU, and multiple vGPU are supported for guests:

> **Note:**
>
> In Citrix Hypervisor 8.0 and earlier releases, you can only add one vGPU to a VM. From Citrix Hypervisor 8.1, you can add multiple vGPUs to a VM if your NVIDIA GPU supports this feature and the vGPUs are of the same type.

	GPU for Windows VMs	GPU for HVM Linux VMs	Shared GPU For Windows VMs	Virtual GPU for Linux VMs	Multiple vGPU For Windows VMs	Multiple vGPU for Linux VMs
AMD	YES		YES			

	GPU for Windows VMs	GPU for HVM Linux VMs	Shared GPU For Windows VMs	Virtual GPU for Linux VMs	Multiple vGPU For Windows VMs	Multiple vGPU for Linux VMs
Intel	YES		YES			
NVIDIA	YES	YES	YES	YES	YES	YES

You might need a vendor subscription or a license depending on the graphics card used.

When you click **Add**, the **GPU type** 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 add a GPU or a virtual GPU to the VM.

If you are using the virtual GPU feature, 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 necessary. For more information, see Change VM Properties.

> **Note:**
>
> - GPU Pass-through and Graphics Virtualization are only available for Citrix Hypervisor Premium Edition customers, or customers who access Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information, see About Citrix Hypervisor 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** list displays all GPU types available on the host or the pool. However, only NVIDIA GPU pass-through is supported for HVM Linux VMs.

Enabling Intel GPU pass-through

Citrix Hypervisor supports the GPU pass-through feature for Windows 8 (32-/64-bit) VMs using an Intel integrated GPU device. This feature is supported on Haswell (Xeon E3-12xx v3) or newer CPUs that contain an Intel integrated GPU device and have a graphics-capable chipset. For more information on the supported hardware, refer to the Citrix Hypervisor Hardware Compatibility List.

When using an Intel GPU on Intel servers, the Citrix Hypervisor server control domain (dom0) have access to the integrated GPU device. In such cases, the GPU is not available for pass through. To use the Intel GPU pass-through feature on Intel servers, disable the connection between dom0 and the GPU before passing through the GPU to the VM.

To disable the connection:

1. Select the Citrix Hypervisor host on the **Resources** pane.
2. On the **General** tab, click **Properties** and then click **GPU** on the left pane.
3. In the **Integrated GPU pass-through** section, click **This server will not use the integrated GPU**.

   This setting disables the connection between dom0 and the Intel integrated GPU device.
4. Click **OK**.
5. Reboot the Citrix Hypervisor server for the changes to take effect.

   The Intel GPU is now visible on the GPU type list during new VM creation and on the VM's **Properties** tab.

> **Note:**
>
> The Citrix Hypervisor 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

May 25, 2023

Virtual machines created using the **New VM** wizard have at least one virtual disk and the selected VM template might also include extra virtual disks. A VM can have up to seven virtual disks including a virtual CD-ROM.

From the **Storage** page in **New VM** wizard, you can:

- Add more virtual disks
- Remove virtual disks
- Change the size and location of virtual disks

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 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 enables you to quickly create VMs.

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

May 25, 2023

By default, XenCenter includes a predefined set of parameters on the **Cloud-Config Parameters** page. You can modify these parameters based on your requirements. See the 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

May 25, 2023

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 is used for all virtual network interfaces. To enter a different MAC address, click **Properties**. Enter a new address in the **Virtual Interface 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**. For more information, see Change virtual network interface properties.

You can use the **Networking** tab to change the VM's virtual networking configuration later, if necessary. For more information, see Configuring VM networking.

Complete New VM Creation

May 25, 2023

On the last page of the **New VM** wizard, choose **Start VM automatically** to ensure the new VM starts up automatically when it is installed.

The process of creating the VM might take some time. The time it takes depends 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 status bar and on the **Events** view under **Notifications**.

> **Note:**
>
> After creating a VM, install Citrix VM Tools to ensure optimized I/O performance. For more information, see Installing Citrix VM Tools.

Express (unattended) VM Creation

May 25, 2023

You can create multiple identical VMs based on a custom VM template by bypassing the **New VM** wizard and using the **Quick Create** feature in XenCenter:

1. Create a custom VM template that specifies all the configuration parameters you want for your new VMs. For more information, see Creating new templates.
2. Choose your custom template in the **Resources** pane. On the **Templates** menu, point to **Create VM From Selection** and then choose **Quick Create**. Alternatively, right-click in the **Resources** pane and choose **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

May 25, 2023

You can create 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

When you create a template using this method, the VM's disks are copied to the new template and the original VM no longer exists. A VM that is currently in a vApp cannot be converted into a template.

1. Shut down the VM as described in Shutdown a VM.
2. In the **Resources** pane, select the VM, right-click and then select **Convert to Template**.
3. Click **Convert** to confirm. You can view the conversion progress in 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 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 select **Create Template from Snapshot** on the shortcut menu.
2. Enter the name of the new template and then click **OK**.

After the template is 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:

1. On the **File** menu, select **Import**
2. Select the XVA file containing the template on the first page of the wizard
3. Follow the same steps as when importing a VM from XVA

For more information, 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 can take some time, depending on the size of the template and the speed and bandwidth of the network connection between XenCenter and the server. When the newly imported template is available, it appears 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

May 25, 2023

You can create VMs and templates by copying (cloning) an existing VM or a template. XenCenter enables you to copy VMs and templates within and across pools.

Citrix Hypervisor 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. This feature uses 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, 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. For information about cloning VMs and running Sysprep, see Prepare to clone a Windows VM by using Sysprep.
> - If the VM you want to copy is a Windows VM, run the Sysprep utility.
> - 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, select **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. Select **Finish**.

To copy a VM to a different pool

1. Select the VM in the **Resources** pane, and on the **VM** menu, select **Copy VM**.
2. On the **Destination** page, select **Cross-pool**. and select **Next**
3. Select a standalone server or a pool from the **Destination** menu.
4. Select a server from the **Home Server** list to assign a home server for the VM and select **Next**
5. On the **Storage** page, specify the storage repository on which to place the virtual disks of the copied VM and select **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.
   - Select **Place migrated virtual disks onto specified SRs** to specify an SR from the **Storage Repository** menu. 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 VM to networks in the destination pool or server. Specify your options using the **Target Network** menu and select **Next**.
7. Select a storage network on the destination pool to use for the live migration of the VM's virtual disks. Select **Next**.

> **Note:**
>
> Due to performance reasons, it is recommended that you do not use the management network for copying VMs.

8. Review the configuration settings and select **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, select **Copy Template**.
2. On the **Destination** page, select **Within Pool**.
3. **Name and Storage** page, enter the name of the new template and 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 are placed.
6. Select **Finish.**

To copy a template to a different pool

1. Select the template in the **Resources** pane, and on the **Templates** menu, select **Copy**.
2. On the **Destination** page, select **Cross-pool** and select **Next**.
3. Select a standalone server or a pool from the **Destination** menu.
4. Select a server from the **Home Server** list to assign a home server for the VM and select **Next**
5. On the **Storage** page, specify a storage repository on which to place the virtual disks of the copied template and select **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.
   - Select **Place migrated virtual disks onto specified SRs** to specify an SR from the **Storage Repository** menu. 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 server. Specify your options using the **Target Network** menu and select **Next**.
7. Select a storage network on the destination pool or server to use to copy the template's virtual disks. Select **Next**.

> **Note:**
>
> Due to performance reasons, it is recommended that you do not use the management network for copying VMs.

8. Review the configuration settings and select **Finish** to start copying the template.

Configuring VMs

May 25, 2023

- Installing Citrix VM Tools
- VM memory configuration
- VM storage configuration
  - Add virtual disks
  - Attach virtual disks
  - Detach virtual disks
  - Move virtual disks
  - Delete virtual disks
  - Change virtual disk properties
- VM networking configuration
  - Add a virtual network interface
  - Activate/deactivate a virtual network interface
  - Remove a virtual network interface
  - Change virtual network interface properties
- Configuring Virtual GPU
- Change VM properties

Installing Citrix VM Tools

February 22, 2024

Citrix VM Tools provide high performance I/O services without the overhead of traditional device emulation.

Citrix VM Tools for Windows

Citrix VM Tools for Windows consist 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 Citrix Hypervisor product family software. While installing a Windows operating system, Citrix Hypervisor uses traditional device emulation to present a standard IDE controller and a standard network card to the VM. This emulation allows the Windows installation to use built-in drivers, but with reduced performance due to the overhead inherent in emulating the controller drivers.

The Management Agent, also known as the Guest Agent, is responsible for high-level virtual machine management features and provides a full set of functions to XenCenter

Get the Citrix VM Tools for Windows installer from the Citrix Hypervisor downloads page.

The version of the Citrix VM Tools for Windows is updated independently of the version of Citrix Hypervisor. For more information about the latest version of the tools, see Updates to Citrix VM Tools for Windows.

Install Citrix VM Tools for Windows on each Windows VM for that VM to have a fully supported configuration, and to be able to use the xe CLI or XenCenter. A VM functions without the Citrix VM Tools for Windows, but performance is hampered when the I/O drivers are not installed. Install Citrix VM Tools for Windows on Windows VMs 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 live migration or storage live migration)
- Create snapshots with memory (checkpoints) or revert to snapshots

For more information, see Install Citrix VM Tools for Windows.

Citrix VM Tools for Linux

Citrix VM Tools for Linux contain a guest agent that provides extra information about the VM to the host.

Get the Citrix VM Tools for Linux installer from the Citrix Hypervisor downloads page.

Install the Citrix VM Tools for Linux on Linux VMs to be able to perform the following operations:

- View VM performance data in XenCenter
- Adjust the number of vCPUs on a running Linux VM

- Enable dynamic memory control

> **Note:**
>
> You cannot use the Dynamic Memory Control (DMC) feature on Red Hat Enterprise Linux 8, Red Hat Enterprise Linux 9, Rocky Linux 8, Rocky Linux 9, or CentOS Stream 9 VMs as these operating systems do not support memory ballooning with the Xen hypervisor.

For more information, see Install Citrix VM Tools for Linux.

> **Important:**
>
> To have a supported configuration when running a VM, ensure that you install Citrix VM Tools. While a Windows VM can function without them, performance is hampered when the I/O drivers are not installed. Run Windows VMs with these drivers to be supported. Some features, such as live relocation across physical hosts, are available only 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 see whether Citrix VM Tools (I/O drivers and the Management Agent) are installed, and whether the VM can install and receive updates from Windows Update. The following section lists the messages displayed in XenCenter:

I/O optimized (not optimized) - displays whether the I/O drivers are installed on the VM.

Management Agent installed (not installed) - displays whether the latest version of the Management Agent is installed on the VM.

Able to (Not able to) receive updates from Windows Update - specifies whether the VM is able to receive I/O drivers from Windows Update.

Install I/O drivers and Management Agent - indicates that the VM does not have the I/O drivers or the Management Agent installed.

> **Note:**
>
> If you have many VMs on your server or a pool, select the server or pool on the Resources pane and select the **Search** tab. From the **Saved Searches** list, select **VMs without Citrix VM Tools Installed**. This search displays a list of VMs that do not have Citrix VM Tools installed.

Updating Citrix VM Tools

Citrix Hypervisor has a simpler mechanism to automatically update I/O drivers (PV drivers) and the Management Agent for Windows VMs. This mechanism 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 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, disable Windows Update on your VM, or specify a group policy.

> **Important:**
>
> - If you are currently using the 8.2.x.x drivers or earlier and want to use the Management Agent MSI file to update to the latest version of the drivers, you must use Device Manager to uninstall the 8.2.x.x drivers from your VM before installing these drivers. If you do not complete this step, the MSI install process fails.
> - 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 update might result in unexpected behavior.

Updating the I/O drivers

If you are running newly created Windows VMs on Citrix Hypervisor or XenServer 7.0 or later, you can get I/O driver updates automatically from Microsoft Windows Update, provided:

- You are running Citrix Hypervisor with Premium Edition, or have access to Citrix Hypervisor through Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement
- You have created a Windows VM using XenCenter issued with Citrix Hypervisor or 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

Citrix Hypervisor enables you to automatically update the Management Agent on both new and existing Windows VMs. By default, Citrix Hypervisor 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 Citrix VM Tools installation. For more information, see Install Citrix VM Tools for Windows. 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 with Premium Edition, or have access to Citrix Hypervisor through Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement.
- You must have installed Citrix VM Tools issued with Citrix Hypervisor or 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 Citrix Hypervisor Premium Edition customers. This feature is also available to those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement.

Citrix VM Tools in Citrix Hypervisor 8.1 and earlier

On Citrix Hypervisor servers that run version 8.1 and earlier, the Citrix VM Tools are included as part of the server installation. For these servers, XenCenter provides the ability to start the Citrix VM Tools installation from within the XenCenter UI.

Installing Citrix VM Tools on Windows VMs

> **Important:**
>
> Installing Citrix VM Tools causes any media in the VM's CD/DVD-drive to be ejected. Do not attempt to install Citrix VM 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 select **Install Citrix VM Tools** on the shortcut menu. Alternatively, on the VM menu, select **Install Citrix VM Tools**.

   Or

   On the **General** tab of the VM, select **Install I/O drivers and Management Agent**.

   > **Note:**
   >
   > When you install Citrix VM Tools on your VM, you are 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, Citrix VM Tools installer displays the installation options. Click **Install Citrix VM Tools** to continue with the installation. This action mounts the Citrix VM Tools ISO (guest-tools.iso) on the VM's CD/DVD drive.

   When prompted, select one of the following options to choose what happens with the Citrix VM Tools ISO:

   Click **Run Setup.exe** to begin Citrix VM Tools installation. This action opens the **Citrix Hypervisor 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 Citrix VM Tools using this method, the Management Agent is configured to get updates automatically. However, the Management Agent update mechanism does not update the I/O drivers. This behavior is the default.

   Alternatively:

   a) Click **Open folders to view files** and then run **Setup.exe** from the CD drive. This option opens the **Citrix Hypervisor Windows Management Agent Setup** wizard and lets you customize the Citrix VM Tools installation and the Management Agent update settings.

   b) Follow the instructions on the wizard to accept the license agreement and choose a destination folder.

   c) Customize your settings on the **Installation and Updates Settings** page. The **Citrix Hypervisor Windows Management Agent Setup** wizard displays the default 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.
      - Sends anonymous usage information to Citrix

      If you do not want to allow the automatic updating of the Management Agent, select **Disallow automatic management agent updates** from the menu.

      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**.

      > **Note:**
      >
      > If you receive I/O driver updates through the Windows Update mechanism, we rec-

> ommend that you do not allow the Management Agent to update the I/O drivers automatically.

If you do not want to share anonymous usage information with Citrix, clear the **Send anonymous usage information to Citrix** check box. The information transmitted to Citrix contains the UUID of the VM requesting the update. No other information relating to the VM is collected or transmitted to Citrix.

d) Click **Next** and then **Install** to begin the installation process.

e) When prompted, perform any actions that are required to complete the Citrix VM 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 many Windows VMs, install `managementagentx86.msi` or `managementagentx64.msi` using your preferred MSI installation tool. These files can be found on Citrix VM Tools ISO.
> - I/O drivers are automatically installed on a Windows VM that can receive updates from Windows Update. However, we recommend that you install the Citrix VM Tools package to install the Management Agent and to maintain a supported configuration.

Installing Citrix VM Tools on Linux VMs

1. Select the VM in the **Resources** pane, right-click, and then click **Install Citrix VM Tools** on the shortcut menu. Alternatively, on the VM menu, click Install Citrix VM Tools.
2. Click **Install Citrix VM Tools** on the message dialog to go to the VM's console.
3. As the root user, mount the image into the VM:

```
mount -o ro,exec /dev/disk/by-label/Citrix\\x20VM\\x20Tools /mnt
<!--NeedCopy-->
```

> **Note:**
>
> If mounting the image fails, you can locate the image by running the command: `blkid -t LABEL="Citrix VM Tools"`

4. Run the installation script as the root user:

```
/mnt/Linux/install.sh
<!--NeedCopy-->
```

5. Unmount the image from the guest by running the command:


```
umount /mnt
<!--NeedCopy-->
```

6. If the kernel has been upgraded, or the VM was upgraded from a previous version, reboot the VM now.

CD-ROM drives and ISOs attached to Linux VMs appear as `/dev/xvdd` (or `/dev/sdd` in Ubuntu), rather than `/dev/cdrom`. This naming is because they are not true CD-ROM devices, but normal devices. When XenCenter ejects the CD, it hot-unplugs the device from the VM and the device disappears. This behavior is different from Windows VMs, where the CD remains in the VM in an empty state.

Configuring VM Memory

February 22, 2024

When a VM is first created, it is allocated a fixed amount of memory. To improve the utilization of physical memory in your Citrix Hypervisor environment, you can use Dynamic Memory Control (DMC). DMC is 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 are 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 configuration includes whether DMC is enabled and the current dynamic minimum and maximum values. You can edit DMC configuration settings in this tab.

VMs with the same memory configuration are grouped in the **Memory** tab, enabling you to view and configure memory settings for individual VMs and for groups of VMs.

> **Note:**
>
> You cannot use the Dynamic Memory Control (DMC) feature on Red Hat Enterprise Linux 8, Red Hat Enterprise Linux 9, Rocky Linux 8, Rocky Linux 9, or CentOS Stream 9 VMs as these operating systems do not support memory ballooning with the Xen hypervisor.

Dynamic Memory Control (DMC)

Dynamic memory control (sometimes known as *dynamic memory optimization*, *memory overcommit*, or *memory ballooning*) works by automatically adjusting the memory of running VMs.

- DMC keeps the amount of memory allocated to each VM between specified minimum and maximum memory values
- DMC guarantees performance
- DMC permits greater density of VMs per server

Without DMC, if you start further VMs when a server is full, the action fails 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, Citrix Hypervisor attempts 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. This dynamic memory range is the range within which memory can be added or removed from the VM without requiring a reboot. You can adjust the dynamic range while the VM is running, without having to reboot it. Citrix Hypervisor always guarantees to keep the amount of memory allocated to the VM within the dynamic range. For example, if the Dynamic Minimum Memory is 512 MB and the Dynamic Maximum Memory is 1,024 MB, the VM has a Dynamic Memory Range of 512–1,024 MB. The VM operates within this range. With DMC, Citrix Hypervisor guarantees to always assign each VM memory within its specified DMR.

When host memory is plentiful, all running VMs receive their Dynamic Maximum Memory level. When host memory is scarce, all running VMs 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 Citrix Hypervisor supports do not fully support dynamically adding or removing memory. As a result, the Citrix Hypervisor server must declare the maximum amount of memory that a VM can be asked to consume when the VM starts. The guest operating system can use this information to size its page tables and other memory management structures accordingly. This feature introduces the concept of a static memory range within the Citrix Hypervisor product. The static memory range cannot be adjusted while the VM is running. The dynamic range is constrained such as to be always contained within this static range until the VM is next rebooted. The static minimum is there to protect the administrator. Set the static minimum to the lowest amount of memory that the OS can run with on a Citrix Hypervisor server.

> **Important:**
>
> Citrix advises you not to change the static minimum level as this value is set at the supported level per operating system. By setting a static maximum level higher than a dynamic max, you can allocate more memory to a VM in the future 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.
- The minimum dynamic memory must be at least 75% of the static maximum. A lower amount can cause in-guest failures and is not supported.

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 systems.

To enable DMC

1. Choose a VM or server in the **Resources** pane and select the **Memory** tab.
2. Select the **Edit** button for the VM or group of VMs you want to configure.
3. For multiple VMs with the same current memory configuration, choose the VMs you want to configure and click **Next**.
4. Select the **Automatically allocate memory within this range** option.
5. Set the required maximum and minimum dynamic memory range values by using the slider or by typing the values directly.
6. Click **OK** to apply the changes and close the dialog box.

To disable DMC

1. Choose the VM or server in the **Resources** pane and select the **Memory** tab.
2. Select the **Edit** button for the VM or group of VMs you want to configure.
3. For multiple VMs with the same current memory configuration, choose the VMs you want to configure and click **Next**.
4. Select the **Set a fixed memory** option.
5. Specify the amount of memory to allocate.
6. Click **OK** to apply the changes and close the dialog box.

Configuring Virtual Storage

May 25, 2023

Storage on Citrix Hypervisor 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 Citrix Hypervisor Storage Repositories (SRs), and can be attached, detached, and reattached to the same or different VMs when needed. New virtual disks can be created at VM creation time (from within the **New VM** wizard). They can also be added after the VM has been created from the VM's **Storage** tab.

Virtual disks on VMs with Citrix VM Tools installed can be *hot plugged*. That is, you can add, delete, attach, and detach virtual disks without having to shut down the VM first. VMs without Citrix VM Tools installed must be shut down before you carry out any of these operations. To avoid this situation, install Citrix VM Tools on all virtual machines. For more information, see Citrix VM Tools.

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

May 25, 2023

To add a new virtual disk, use the **Add Virtual Disk** dialog box.

> **Important:**
>
> If the VM is running without Citrix VM Tools installed, shut it down before you can add any virtual disks. To avoid this situation, install Citrix VM Tools on all virtual machines. For more information, see Citrix VM Tools.

Procedure:

1. Open the **Add Virtual Disk** dialog box by doing any of the following:

- Select the VM or storage repository in the **Resources** pane, select the **Storage** tab and then select **Add**.
- On the **Storage** menu, select **Virtual Disks** and then **New Virtual Disk**.
- On the **Storage** page of the **New VM** wizard, select **Add**.

2. Enter the name of the new virtual disk and, optionally, a description.
3. Enter the size of the new virtual disk. Ensure that the storage repository (SR) on which the virtual disk is to be stored has sufficient space for the new virtual disk.
4. Select the SR where the new virtual disk is stored.
5. Click **Create** to add the new virtual disk and close the dialog box.

Attach Virtual Disks

May 25, 2023

You can add storage to a VM by attaching an existing virtual disk.

1. Select the VM in the **Resources** pane, select the **Storage** tab, and then select **Attach**. Alternatively, on the **Storage** menu, select **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 setting can help prevent data from being overwritten or changed when multiple VMs access the disk. It also allows you to attach the virtual disk 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 situation happens, activate it again from the VM's **Storage** tab by selecting it and clicking **Activate**.

Detach Virtual Disks

May 25, 2023

When you detach a virtual disk from a VM, the virtual disk and the data on it are preserved. The virtual disk is no longer available to the VM. The detached storage device can later be reattached 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 the following conditions are met:

- The VM is not suspended.
- The VM must have Citrix VM Tools installed.
- The virtual disk is not a system disk.
- The virtual disk must be deactivated to be able to detach it cleanly. The term *deactivate* is equivalent to *unplug*, which is the term used for this operation in the product documentation and in the CLI.

If any of these conditions are not satisfied, shut the VM down before you can detach the virtual disk.

To detach the virtual disk:

1. Select the VM in the **Resources** pane and click the **Storage** tab.
2. Select the virtual disk in the list, click **Deactivate**, and click **Detach**.
3. Click **OK** to confirm the operation.

Move Virtual Disks

May 25, 2023

Virtual disks can be moved or migrated from one storage repository (SR) to a different SR within the same pool. The following types of virtual disks can be moved or migrated:

- 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 live migration)

> **Note:**
>
> 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 live migration

Storage live migration 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 environment to a production environment.
- Moving between tiers of storage when a VM is limited by storage capacity.
- Performing storage array upgrades.

Virtual disks with more than one snapshot cannot be migrated.

To move a virtual disk

1. In the XenCenter **Resources** pane, select the SR where the virtual disk is stored, and then select the **Storage** tab. To locate a virtual disk:
   - In the XenCenter **Resources** pane, select the VM to which the virtual disk that you want to move is attached.
   - Click the **Storage** tab and identify the SR on which the virtual disk is stored.
2. From the **Virtual Disks** list, select one or more virtual disks that you would like to move, and then select **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

May 25, 2023

You can delete a virtual disk without shutting down the VM first if the following conditions are met:

- The VM is not suspended.
- The VM must have Citrix VM Tools installed.
- The virtual disk is not a system disk.
- The virtual disk must be deactivated first. The term *deactivate* is equivalent to *unplug*, which is the term that is used for this operation in the product documentation and in the CLI.

If any of these conditions are not satisfied, shut the VM down before you can delete the virtual disk.

> **Important:**
>
> Deleting a virtual disk permanently deletes the disk, destroying any data stored on it.

To delete the virtual disk:

1. On the VM's **Storage** tab, select the virtual disk in the list and select **Deactivate** and then **Delete**.
2. Click **OK** to confirm the deletion.

Change Virtual Disk Properties

May 25, 2023

To change the properties of a virtual disk, select the VM's **Storage** tab, then select the virtual disk and select **Properties**.

General properties - name, description, folder, tags

Property	Description
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 existing custom fields, and remove custom fields.

For information on adding, setting, modifying, and deleting custom fields, see Using 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.

Option	Description
Mode	The disk read/write permissions of a virtual disk can be changed. For example, change this setting to prevent data from being overwritten on a virtual disk that you use 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 can adjust the disk I/O priority. This setting is only available for virtual disks on storage repositories that are LVM-based: local, shared iSCSI, or hardware HBA. This option is visible only after the disk scheduler is changed to `cfq`. It is not available by default. For more information, see Manage storage repositories.

Configuring VM Networking

May 25, 2023

Each virtual machine (VM) can have one or more virtual network interfaces that act as virtual NICs.

A virtual network interface has the following properties:

Property	Description
Network	The (physical) network location of the virtual network interface.
MAC address	The MAC address of the virtual network interface.

Property	Description
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 slows the memory processing and helps make the system more stable by preventing crashes.

All the virtual network interfaces for a VM are listed on the VM's **Networking** tab. Here, you can add new virtual network interfaces. You can also edit, activate, deactivate, and remove existing virtual network interfaces.

- Networking
- Add a Virtual Network Interface
- Activate/deactivate a Virtual Network Interface
- Remove a Virtual Network Interface
- Change Virtual Network Interface Properties

Add a Virtual Network Interface

July 19, 2023

You can add up to seven Virtual Network Interfaces (VNIs) for a VM. For more information about the maximum number of VNIs supported for a VM, see the Citrix Hypervisor Configuration Limits.

1. Open the **Add Virtual Interface** dialog box by doing one of the following:
   - Select the VM in the **Resources** pane. Go to the **Networking** tab. 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.
   - To use a generated MAC address, select **Auto-generate a MAC address**.
   - To explicitly enter a MAC address, select **Use this MAC address**. 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).
5. Click **Add**.

Activate/deactivate a Virtual Network Interface

May 25, 2023

You can activate or deactivate a virtual network interface on a running VM when the VM meets the following conditions:

- The VM is not suspended.
- The VM must have Citrix VM Tools installed.

To activate ("plug") or deactivate ("unplug") a virtual network interface:

1. On the VM's **Networking** tab, select the interface.
2. Click the button labeled **Activate** or **Deactivate**.

Remove a Virtual Network Interface

May 25, 2023

You can remove a virtual network interface from a VM without shutting down the VM ("hot unplug") if the VM meets the following conditions:

- The VM is not suspended.
- The VM must have Citrix VM Tools installed.

If one or both of these conditions are not satisfied, 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.
2. Select the **Networking** tab.
3. On the **Networking** tab, choose the virtual network interface in the list. Click **Remove**.

Change Virtual Network Interface Properties

May 25, 2023

To change properties of a virtual network interface, open the **Virtual Interface Properties** dialog box by doing one of the following:

- Click the VM's **Networking** tab, then select the virtual network interface and click **Properties**.
- On the **Network** page of the **New VM** wizard, click **Edit**.

You can change the network location and MAC address of a virtual network interface, and you might also be able to set its I/O priority.

Property	Description
Network	The network location of the virtual network interface.
MAC address	The MAC address of the virtual network interface. You can autogenerate this value or enter it manually in the form XY:XX:XX:XX:XX:XX where X is any hexadecimal digit, and Y is 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 slows the memory processing and helps make the system more stable by preventing crashes.

Configuring Virtual GPU

May 25, 2023

GPUs are grouped 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 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 select 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. This information includes: 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 select **OK**. If you would like to pass-through the whole GPU, select **Pass-through whole GPU**.

Change VM Properties

December 28, 2023

Select a virtual machine in the **Resources** pane. On the **General** tab, select 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 the 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 choose a folder. For more information, see 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. For more information, see Using custom fields.

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. Shut down the VM before you modify 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** list. To get the best performance out of your VM, ensure the number of vCPUs does 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** list. To get the best performance out of your VM, ensure the maximum number of vCPUs does 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** 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 is 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 list 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 list.

> **Note:**
>
> 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 the other VMs running on the same host server. To adjust the vCPU priority for the VM, move the vCPU slider.

The Citrix Hypervisor templates provide typical VM configurations and set reasonable defaults for the memory, based on the type of guest operating system. Consider the following factors when deciding how much memory you give to a VM:

- The kinds of applications that run on the VM.
- Other virtual machines that use the same memory resource.
- Applications that run on the server alongside the virtual machine.

Boot options

The available boot options on this tab can vary, depending on the guest operating system. For example, on some VMs, you can change the boot order (or boot sequence), or specify extra boot parameters.

- To change the boot order, select an item in the **Boot Order** list and select **Move Up** or **Move Down**.
- To specify extra 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**, the start delay interval, and the **HA restart priority** for the selected VM.

Start order

Specifies the order in which individual VMs are started up within a vApp or during a high availability recovery operation, allowing certain VMs to be started before others. VMs with a start order value of 0 (zero) are started first, then VMs with a start order value of 1, and so on.

Attempt to start next VM after

This value is a delay interval that specifies how long to wait after the VM starts before starting the next group of VMs in the startup sequence. This setting applies to VMs within a vApp and to individual VMs during a high availability recovery operation.

HA restart priority

In a pool with high availability enabled, this setting specifies which VMs are restarted automatically if the underlying hardware fails or their host server is lost.

- VMs with an HA restart priority of **Restart** are guaranteed to be restarted if sufficient resources are available within the pool. They are 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. However, one attempt to restart them is 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** are not restarted automatically.

For more information, see VM startup 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 provides resources for the VM. The VM is started up on that server if possible. If it is not possible to start eh VM on that server, an alternate server within the same pool is selected automatically. For more information, see Creating a new VM.

In pools with Workload Balancing (WLB) enabled, you cannot set a home server. Instead, XenCenter nominates the best server for the VM by analyzing Citrix Hypervisor resource pool metrics and recommending optimizations. You can decide if you want these recommendations geared towards resource performance or hardware density. 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 one or more virtual GPUs to a VM. This configuration gives the VM direct access to the graphics hardware. The VM can 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.

> **Note:**
>
> In Citrix Hypervisor 8.0 and earlier releases, you can only add one vGPU to a VM. From Citrix Hypervisor 8.1, you can add multiple vGPUs to a VM if your NVIDIA GPU supports this feature and the vGPUs are of the same type.

Click **Add** to add a GPU to the VM. 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 Citrix Hypervisor Premium Edition customers, or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about licensing, see About Citrix Hypervisor Licensing.

USB

On the VM's **USB** properties tab, the right side pane displays the list of USBs attached to the VM. You can attach extra (maximum of 6) USBs to the VM. You can also choose to detach USBs from the VM.

For more information, see Tabs.

> **Note:**
>
> - USB pass-through is available for Citrix Hypervisor Premium Edition customers.
> - USB pass-through is supported for the following USB versions: 1.1, 2.0, and 3.0.

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, extra shadow memory is required to achieve the 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, select **Optimize for Citrix Virtual Apps**.
- To manually adjust the VM's shadow memory allocation, select **Optimize manually** and enter a number in the **Shadow memory multiplier** box.

- To restore the default settings for the VM's shadow memory, select the **Optimize for general use** option.

Cloud-Config parameters

> **Note:**
>
> 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, see Cloud-Config Parameters.

Managing VMs

May 25, 2023

- Start a VM
- Suspend and resume a VM
- Shut down a VM
- Reboot a VM
- Run a remote console session
- Migrate VMs
- Delete VMs
- Changed Block Tracking

Start a VM

May 25, 2023

For VMs in a pool, you can choose where to start your VMs. This choice is 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 start the VM on any server in the pool (subject to available storage on that server). For more information, see Creating a new VM.

When the VM is up and running, its status indicator changes to the **VM running** icon in the **Resources** pane.

To start a VM on a specific server

1. Select the VM in the **Resources** pane.
2. Right-click and select **Start on Server** and then select the server you want on the shortcut menu. Alternatively, on the **VM** menu, select **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 select **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, select **Start**.

In a WLB-enabled pool, this action starts the VM on the optimal server.

In a pool without Workload Balancing configured, this action starts the VM on its home server. If no home server has been set, the VM starts on the first available server.

Suspend and Resume a VM

May 25, 2023

When you suspend a VM, its current state is stored in a file on the default storage repository (SR). This feature allows you to shut down the VM's host server. After rebooting the server, you can resume the VM and return it to its original running state.

> **Note:**
>
> It might 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 might 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 select **Suspend** on the Toolbar.

   

   Alternatively:

   - Right-click and select **Suspend** on the shortcut menu.
   - On the **VM** menu, select **Suspend**.

   When a VM has been suspended, its status indicator changes to the Suspended VM icon 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 resume the VM on any server in the pool (subject to available storage on that server). For more information, see Home server.

When a suspended VM has been successfully resumed, its status indicator changes to the **VM Running** icon in the **Resources** pane.

To resume a suspended VM on a specific server

1. Select the VM in the **Resources** pane.
2. Right-click and select **Resume on Server** and then select the server you want on the shortcut menu. Alternatively, on the **VM** menu, select **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 select **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, select **Resume**.

In a WLB-enabled pool, the VM starts on the optimal server.

In a pool without Workload Balancing configured, the VM starts on its home server. If no home server has been set or if the nominated server is unavailable, the VM starts on the first available server.

Shut Down a VM

May 25, 2023

You might need to shut down a running VM for several different reasons. For example:

- to free up its resources
- to reconfigure its virtual network hardware
- to reconfigure its 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 might not always shut down all running processes and you risk losing data if you shut down a VM in this way. Only use a forced shutdown when a soft shutdown is not possible.

A VM running in HVM mode (that is, VMs without Citrix VM Tools installed) can only be shut down using a forced shutdown. To avoid this situation, install Citrix VM Tools on all HVM virtual machines. For more information, see Citrix VM Tools.

To perform a soft shutdown

Select the VM in the **Resources** pane and then choose **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 the lifecycle icon and then click **Shut Down**.

The VM's console displays shutdown messages as running processes are stopped. When the shut-down is complete, the VM status indicator in the **Resources** pane changes to the stopped VM icon.

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 the lifecycle icon and then click **Force Shut Down**.

When the shutdown is complete, the VM status indicator changes to the stopped VM icon in the **Resources** pane.

Reboot a VM

May 25, 2023

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 action works like pulling the plug on a physical server and then plugging it back in and turning it back on.

Only do a forced reboot as a last resort to forcibly retrieve the system from instances such as a critical error.

An HVM-mode VM without Citrix VM Tools installed can only be rebooted using a forced reboot. To avoid this situation, install Citrix VM Tools on all HVM virtual machines. For more information, see Citrix VM Tools.

To reboot a VM cleanly

In the **Resources** pane, select the VM and then select **Reboot** on the toolbar.

Alternatively:

- Right-click and select **Reboot** on the **Resources** pane shortcut menu.
- On the **VM** menu, select **Reboot**.

The VM is shut down and rebooted. When this process is complete, its status indicator in the **Resources** pane changes back to the **VM start** icon.

To do a forced reboot

In the **Resources** pane, select the VM and then select **Force Reboot** on the toolbar.

Alternatively:

- Right-click and select **Force Reboot** on the **Resources** pane shortcut menu.
- On the **VM** menu, select **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 the **VM start** icon.

Run a Remote Console Session

May 25, 2023

To open a remote console session on a VM, select the VM and then select 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. To use the graphical console ensure that the VNC server and an X display manager are installed and configured on the VM. For information about configuring VNC for Linux virtual machines, see Enable VNC for Linux VMs.

To switch between the two types of remote console, use the **Switch to Graphical Console**/**Switch to Text Console** button on the **Console** tab.

> **Note:**
>
> For HVM Linux guests, screen blanking can take effect after a period of inactivity (typically 10 minutes). When screen blanking happens, the console is black and remains blank until a key is pressed at which point the text reappears.

You can disable this behavior within the guest by adding `consoleblank=0` to the kernel boot parameters.

SSH console

XenCenter allows you to initiate SSH connections to Linux VMs using the **Open SSH Console** button on the VM's Console tab. This action launches 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 Linux VMs.
- 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 are 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 Citrix Hypervisor developed to provide remote access to your VM console. The Remote Desktop console uses RDP (Remote Desktop Protocol) technology. Switch

between a standard graphic console and a Remote Desktop console by using the **Switch to Remote Desktop/Switch to Default Desktop** button on the XenCenter **Console** tab.

To use a Remote Desktop console connection, ensure that the following requirements are met:

- Remote Desktop must be enabled on the virtual machine - see Windows VMs for information on how to enable Remote Desktop on a Windows virtual machine.
- Citrix VM Tools must be installed.
- The virtual machine must have a network interface and be able to connect to XenCenter.
- The Credential Security Support Provider protocol (CredSSP) update must be applied to either both or neither of the client and server in the RDP connection. For more information, see https://support.microsoft.com/en-gb/help/4295591/credssp-encryption-oracle-remediation-error-when-to-rdp-to-azure-vm.

There are various 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

April 18, 2024

This topic contains information about migrating and moving virtual machines within and across pools and standalone servers.

Definitions:

- Migrate a VM: Move a running or a suspended VM to a different server or a pool.
- Move a VM: Move a shut-down VM to a different server or pool.

Live migration

Live migration is available in all versions of Citrix Hypervisor. This feature allows you to move a running or a suspended VM between Citrix Hypervisor servers, when the VM's disks are on storage shared by both servers. This capability allows for pool maintenance features such as Workload Balancing (WLB), high availability, and Rolling Pool Upgrade (RPU) to automatically move VMs. Storage can only be shared between hosts in the same pool. As a result, you can move VMs only within the same pool.

Live migration enables the following to occur without any VM downtime:

- Workload leveling
- Infrastructure resilience
- Upgrade of the server software

Storage live migration

Storage live migration also allows VMs to be moved from one host to another, where the VMs are not on storage shared between the two hosts. As a result, you can migrate VMs stored on local storage without downtime and you can move VMs 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. 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 live migration enables system administrators to:

- Rebalance VMs between Citrix Hypervisor pools (for example from a development environment to a production environment)
- Upgrade and update standalone Citrix Hypervisor servers without any VM downtime
- Upgrade the Citrix Hypervisor server hardware

> **Note:**
>
> You cannot use storage live migration to migrate VMs that have changed block tracking enabled. Disable changed block tracking before attempting storage live migration. For more information, see Changed Block Tracking.
>
> Moving a VM from one host to another preserves the VM state. The state includes information that defines the VM and the historical performance metrics, such as CPU and network usage.

Storage live migration 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 live migration or storage live migration, the new VM and server must meet the following compatibility requirements:

- Citrix VM Tools must be installed on each VM that you want to migrate.
- The destination server must have the same or a more recent version of Citrix Hypervisor installed as the source.
- (Storage live migration only.) 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. Therefore, it is unlikely to be possible to move a VM between, for example, AMD and Intel processors.
- For storage live migration, VMs with more than one snapshot cannot be migrated.
- VM with checkpoint cannot be migrated.
- For storage live migration, 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 fails to complete.
- Storage migration only: A host in the source pool must have sufficient spare memory capacity to run a halted VM that is being migrated. This requirement enables the halted VM to be started at any point during the migration process.
- For storage live migration, 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 fails to complete.
- The source storage must have enough free disk space to create temporary snapshots of the VM's VDIs during the migration. If there is not enough space, the migration fails to complete. The free space required can be up to two times the size of the VM's disk.

Live migration and storage live migration limitations

Live migration and storage live migration are subject to the following limitations:

- Storage live migration cannot be used with VMs created by Machine Creation Services.
- VMs using SR-IOV cannot be migrated. For more information, see Use SR-IOV enabled NICs
- VM performance is reduced during migration.
- If using the high availability feature, ensure the VM being migrated is not marked as protected.
- Time to completion of VM migration depends on the memory footprint of the VM, and its activity. In addition, the size of the VDI and the storage activity of the VDI can affect VMs being migrated with storage live migration.
- Intel GVT-g is not compatible with live migration and storage live migration. For more information, see Graphics overview

- VMs that have the `on-boot` option set to `reset` cannot be migrated. For more information, see Intellicache.

For step-by-step instructions on using live migration or storage live migration to migrate your VMs, see the section *To Migrate or Move a VM*.

Move VMs

XenCenter allows you to move shut-down VMs to a new storage repository in the same pool by using the **Move VM** wizard. For step-by-step instructions, see the following section.

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 live migration or storage live migration: On the **VM** menu, select **Migrate to Server** and then **Migrate VM wizard**. This action opens the **Migrate VM** wizard.

     > **Note:**
     >
     > For pools with 16 or fewer members, the right-click menu displays the list of available servers to migrate the VM to. However, for larger pools, the servers are not listed in the menu. Instead you must open the **Migrate to Server** wizard.

   - To move a stopped VM: On the **VM** menu, select **Move VM**. This action opens the **Move VM wizard**.
2. Select a standalone server or a pool from the **Destination** list.
3. Select a server from the **Home Server** list to assign a home server for the VM and select **Next**.
4. On the **Storage** page, specify a storage repository to place the virtual disks of the migrated VM on. Select **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.
   - Select **Place migrated virtual disks onto specified SRs** to specify an SR from the **Storage Repository** 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 server. Specify your options using the **Target Network** list and select **Next**.

6. Select a storage network on the destination pool that is to be used for the migration of the VM's virtual disks. Select **Next**.

   > **Note:**
   >
   > Due to performance reasons, it is recommended that you do not use the management network for live migration.

7. Review the configuration settings and select **Finish** to start migrating or moving the VM.

If you are upgrading from 7.1 CU2 to 8.2 CU1, you might need to shut down and boot all VMs after migrating your VMs, to ensure that new virtualization features are picked up.

Delete a VM

May 25, 2023

Deleting a virtual machine removes its configuration and its filesystem from the server. When you delete a VM, you can choose to delete or preserve any virtual disks attached to the VM, in addition to 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 select **Delete** on the shortcut menu. Alternatively, on the **VM** menu, select **Delete**.
3. To delete an attached virtual disk, select its check box.

   > **Important:**
   >
   > Any data stored in the VM's virtual disk drives is 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. These snapshots can be exported, deleted, or used to create VMs and templates. To view snapshots in the **Resources** pane, select **Objects** in the Navigation pane and then expand the **Snapshots** group in the Resources pane.

Changed Block Tracking

May 25, 2023

The Citrix Hypervisor changed block tracking feature offers incremental backup capabilities for customers using Citrix Hypervisor. This feature is available only for Citrix Hypervisor Premium Edition.

When you enable changed block tracking for the virtual disk images (VDIs) of a VM, any blocks that are changed in a VDI are recorded in a log file. Every time you take a snapshot of the VDI, this log file can be used to identify the blocks that have changed since the last snapshot of the VDI. With this feature, you can back up only those blocks that have changed.

Changed block tracking can be enabled by using the CLI or API. The third-party product that you use to take incremental backups usually enables this feature. When changed block tracking is enabled for a VDI, additional information is calculated and stored that lists the changed blocks for the VDI. This process uses resources such as memory and space.

For more information about changed block tracking, see the developer documentation.

Disabling changed block tracking on a VM

You can disable changed block tracking for all VDIs associated with a VM by using XenCenter. Before disabling changed block tracking for a VM or VMs, consider the following:

- To use storage live migration to move a VM, you must disable changed block tracking on that VM.
- Disabling changed block tracking prevents your backup solution from taking incremental backups of the VDIs associated with the VM. To take another set of incremental backups, you must enable changed block tracking again.
- Changed block tracking cannot be enabled again by using XenCenter.

To disable changed block tracking, complete the following steps:

1. In the left panel, choose the VM or VMs that you want to disable changed block tracking on.
2. From the main menu, select **VM** then **Disable Changed Block Tracking**.
3. In the confirmation dialog that opens, select **Yes** to continue.

Viewing the changed block tracking status for a VDI

You can see whether changed block tracking is enabled for a VDI on the **Storage** tab for an SR.

If the SR is part of a pool where changed block tracking is available, XenCenter displays the **Changed Block Tracking** column. This column shows whether changed block tracking is **Enabled** or **Disabled** for a VDI.

Importing and Exporting VMs

February 8, 2024

You can import VMs from OVF/OVA packages, from disk images, and from Citrix Hypervisor XVA files. VMs can be exported as OVF/OVA packages and as Citrix Hypervisor XVA files. Import and export VMs in XenCenter using the **Import** and **Export** wizards.

When importing VMs created on hypervisors other than Citrix Hypervisor, use the Operating System Fixup tool to ensure that imported VMs can boot on a Citrix Hypervisor server.

You can import or export a UEFI-enabled VM created on a Citrix Hypervisor server as an OVA, OVF, or an XVA file. Importing a UEFI-enabled VM from other hypervisors is not supported.

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 VMs. For more information about XenCenter support for OVF and OVA file formats, see Open Virtualization Format.
Disk image formats	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)
Citrix Hypervisor 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 name extension is `.xva`.

Which format to Use?

Use OVF/OVA to:

- Share Citrix Hypervisor 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:

- 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 interoperability for VMs that are imported to a Citrix Hypervisor server. Use Operating System Fixup when importing VMs created on other hypervisors from OVF/OVA packages and disk images.

Operating System Fixup configures a guest OS to boot by enabling boot devices critical for booting in a Citrix Hypervisor server and disabling any services or tools for other hypervisors. Guest OSes include all versions of Windows that Citrix Hypervisor supports and some Linux distributions.

> **Note:**
>
> 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. It 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. Specify a location to copy the Fixup ISO to so that Citrix Hypervisor 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.

Importing VMs: overview

When you import a VM, you are effectively creating a VM. The import process involves many of the same steps as creating a VM, such as nominating a home server and configuring storage and networking for the VM. For detailed information about each of these steps, see Creating a New VM.

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 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 to copy the file to. The **Import** wizard continues 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 must 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. For more information, see VM CPU and Memory Allocation. VM names are not checked for uniqueness within XenCenter, so it makes it easier for you to manage different VMs if you give them meaningful, memorable names. For more information, see VM Name and Description.

3. **(OVF/OVA only) Review/accept EULAs.**

   If the package you are importing includes any EULAs, accept them and then select **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 VMs. To nominate a Home Server for the incoming VMs, select a server in the list.

5. **Configure storage.**

   Next, choose the SRs where the virtual disks in the imported VMs are to be placed:

   For VMs in XVA format, you select an SR where the imported VM's virtual disks are to be placed.

   For VMs in OVF/OVA packages or in disk image files, you can place all imported virtual disks onto the same SR. Alternatively, you can place individual virtual disks onto specific SRs.

6. **Configure networking.**

   Next, map the virtual network interfaces in the imported VMs 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 must specify the necessary information.

8. **(OVF/OVA and disk image only) Configure OS Fixup**

   If the VMs you import are built on a hypervisor other than Citrix Hypervisor, you must configure the Operating System Fixup feature to enable the imported VM to boot correctly on a Citrix Hypervisor server.

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 when 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. 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 VMs 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. When exporting to XVA, only one VM can be selected.

3. **(OVF/OVA only) Configure EULA and Advanced Options**

   When exporting VMs as an OVF of OVA package, various extra settings can be configured. For more information, see Export VMs as OVF/OVA.

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 VMs and close the wizard.

Open Virtualization Format (OVF and OVA)

May 25, 2023

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** contains metadata and file elements that describe VMs, plus additional information important to the deployment and operation of the applications in the OVF package. Its file name extension is `.ovf`.

An **Open Virtual Appliance (OVA)** is an OVF Package in a single file archive with the `.ova` extension.

In Citrix Hypervisor environments 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. For more information about RBAC roles, see RBAC overview.

What's in an OVF package?

An **OVF Package** always includes a descriptor file (`*.ovf`) and might also include several other files.

File type	Description
Descriptor	The descriptor specifies the virtual hardware requirements of the service. This descriptor can also include information such as virtual disks descriptions, the service itself, guest OSes, a EULA, instructions to start and stop appliance VMs, and instructions to install the service. The descriptor file name 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 name 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. It allows the package author to be verified. The signature file name extension is `.cert`.

File type	Description
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. Citrix Hypervisor 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 do I use?

OVF packages contain a series of uncompressed files that make it easier to access individual disk images in the file. 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 a single file, making the package easier to handle, such as 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 problem happens for various reasons, including: different interpretations 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 a Citrix Hypervisor server. Run Operating System Fixup on imported VMs that were created on other hypervisors to ensure they boot correctly on a Citrix Hypervisor server.

For more information about the Operating System Fixup feature, see About VM Import and Export.

More information about OVF

For more information about OVF, see the following documents on the DMTF website:

Open Virtualization Format Specification

Disk Image Formats (VHD and VMDK)

May 25, 2023

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 scenario 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. This information includes 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.

In Citrix Hypervisor environments where Role-Based Access Control (RBAC) is implemented, only users with the RBAC role of Pool Admin can import disk images. For more information, see RBAC overview.

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 name 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 name 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. This feature aims to ensure a basic level of interoperability for VMs created on hypervisors other than a Citrix Hypervisor server. Run Operating System Fixup when importing VMs from disk images that were created on other hypervisors to ensure that they boot correctly on a Citrix Hypervisor server.

To find out more, see About VM Import and Export.

Import VMs From OVF/OVA

February 8, 2024

You can import virtual machines (VMs) that have been saved as OVF/OVA files using the **Import** wizard. The wizard takes you through many of the usual steps needed to create a 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.

- Specifying **Operating system fixup** settings if importing VMs that were built on a hypervisor other than a Citrix Hypervisor server.

For an overview of the steps involved in creating a VM, see Creating a New VM.

Imported OVF packages appear as vApps when imported using XenCenter. When the import is complete, the new VMs appear in the XenCenter **Resources** pane and the new vApp appears in the Managing vApps dialog box.

> **Note:**
>
> It might 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 CPU, then exported, might not run when imported to a server with an AMD-V CPU.

Prerequisites

- You need an RBAC role of Pool Admin to import OVF/OVA packages. The **Import** wizard performs checks to ensure that you have a Pool Admin role in the destination pool before allowing you to continue. For more information about RBAC roles, see RBAC overview.
- When importing an OVF Package that was compressed or contains compressed files, extra free disk space is necessary on your Citrix Hypervisor 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 name extension), then select **Next** to continue.
   - If you select a compressed OVA file (`*.ova.gz`), on clicking **Next**, the file is decompressed to an OVA file and the old `*.ova.gz` file is deleted.
   - If you enter a URL location (HTTP, HTTPS, file, FTP) in the **Filename** box, on clicking **Next**, a **Download Package** dialog opens. Use this dialog to specify a folder on your XenCenter host where the package is to be copied.
3. **Review/accept EULAs.** Accept the EULAs and then select **Next** to continue.

   If no EULAs are included in the package, the wizard skips this step and moves 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 VMs** to list. Optionally, assign the VMs a home server:

- To nominate a home server for a VM, select the server from the list in the **Home Server** column. Citrix Hypervisor always attempts to start up a VM on its home server if it can. For more information, see Home Server.
- 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 are to be placed, then select **Next** to continue.

   - To place all the imported disk images on the same SR, select **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, select **Place imported VMs on the specified SR targets**. 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 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. Select **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; select 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 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.

   When the 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 causes the import fail.

8. **Enable Operating System Fixup.** If the VMs in the import package are built on a hypervisor other than the Citrix Hypervisor server, select **Use Operating System Fixup**. Choose an ISO SR where the Fixup ISO can be copied so that Citrix Hypervisor can use it.

   If the ISO library you want is not listed, select **New ISO Library** to create an ISO SR. For more information, see ISO Storage.

   Click **Next** to continue.

9. On the **Finish** page, review all the import settings, and then select **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 can take some time. The import time depends 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 appear in the **Resources** pane and the new vApp appears in the Managing vApps dialog.

> **Note:**
>
> After using XenCenter to import an OVF package that contains Windows operating systems, you must set the **platform** parameter:
>
> ```
> xe vm-param-set uuid=<VM UUID> platform:device_id=0002
> xe vm-param-set uuid=<VM UUID> platform:viridian=true
> ```

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. For more information, see About VM Import and Export.

Import Disk Images

April 16, 2024

Using the XenCenter **Import** wizard, you can import a disk image into a resource pool or a specific host, creating a VM. The wizard takes you through many of the usual steps needed to create a 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 VM.
- Specifying settings if importing disk images that were built on a hypervisor other than Citrix Hypervisor.
- Specifying a boot mode for the new VM.

For more information, see VM Import and Export.

Requirements for importing disk images

You need an RBAC role of Pool Admin to import disk images. The **Import** wizard performs checks to ensure that you have a Pool Admin role in the destination pool before allowing you to continue. For more information, see RBAC overview.

Procedure

1. Open the **Import** wizard: on the **File** menu, select **Import**.
2. 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 opens. In this dialog, specify a folder on your XenCenter host where the disk image is to be copied.
3. **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. For more information, see VM CPU and Memory Allocation.

   Click **Next** to continue.
4. **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.

- Select a pool or standalone server from the Import VMs to** list.
- To nominate the home server for the VM, select a server from the list in the **Home Server** column. Citrix Hypervisor always attempts to start up a VM on its home server if it can. For more information on assigning a home server to new VMs, see Home Server.
- 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.

5. **Configure storage for the new VM.**

   On the **Storage** page, select a storage repository (SR) where the imported virtual disk is placed, then click **Next** to continue.

6. **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.

7. **Specify the boot option.**

   On the **Boot options** page, specify the boot mode for the new VM. Select **BIOS Boot** to boot the VM in legacy BIOS mode.

   Citrix Hypervisor supports UEFI boot only on newly created Windows 10 (64-bit), Windows Server 2016 (64-bit), Windows Server 2019 (64-bit), and Windows Server 2022 (64-bit) VMs. Guest UEFI boot is an experimental feature. You can create UEFI-enabled VMs on hosts that are in a production environment. However, UEFI-enabled VMs must not be used for production purposes. You might have to re-create the VMs when you upgrade the host to a newer version of Citrix Hypervisor.

   For detailed information about Guest UEFI boot, see What's new.

   Select **UEFI Boot** to boot the VM in UEFI mode.

8. **Enable Operating System Fixup.**

   If the disk image that you are importing was built on a hypervisor other than Citrix Hypervisor, select **Use Operating System Fixup**. Choose an ISO SR where the Fixup ISO can be copied so that Citrix Hypervisor can use it. For more information, see About VM Import and Export.

- On the **Finish** page, review all the import settings. 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 might take some time. The time it takes depends 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 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 varies according to the version of Windows contained in the disk image:
>
> - For Windows Server, set the platform parameter to device_id=0002. For example:
>
> ```
> xe vm-param-set uuid=<VM UUID> platform:device_id=0002
> <!--NeedCopy-->
> ```
>
> - For all other versions of Windows, set the platform parameter to viridian=true. For example:
>
> ```
> xe vm-param-set uuid=<VM UUID> platform:viridian=true
> <!--NeedCopy-->
> ```

Import VMs From XVA

May 25, 2023

You can import VMs, templates, and snapshots that have been exported and stored locally in XVA format (`.xva`) by using the XenCenter **Import** wizard.

Importing a VM from an XVA file involves the same steps as creating and provisioning a new VM using the **New VM** wizard. These steps can include nominating a home server and configuring storage and networking for the new VM. For more information, see Creating a New VM.

> **Note:**
>
> It might 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 CPU, then exported, might 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 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 opens. In this dialog, specify a folder on your XenCenter host where the files are to 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. Click **Next** to continue.
   - To place the imported VM in a pool and assign it to a specific server (or to place it on a standalone server), select a server. Click **Next** to continue.

4. On the **Storage** page, select a storage repository (SR) where the imported virtual disks are to 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 when 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 can take some time. The time it takes depends 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

February 8, 2024

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, select **Export**. The VMs must be shut down or suspended before they can be exported.

You need an RBAC role of Pool Admin to export to OVF/OVA. The **Export** wizard performs checks when it starts up to ensure that you have a Pool Admin role before allowing you to continue. For more information, see RBAC overview.

> **Note:**
>
> It might 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 CPU, then exported, might not run when imported to a server with an AMD-V CPU.

Procedure

1. Open the **Export** wizard: select the pool or server containing the VMs you want to export, then on the **VM** menu, select **Export**.
2. On the first page of the wizard, enter the name of the export file. Specify the folder where you want the files to be saved. Select **OVF/OVA Package (*.ovf, *.ova)** from the **Format** list. Select **Next**.
3. Select the VMs you want to export. Select **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 select **View**. If you do not want to include a EULA in the package, select **Next** to continue.
5. On the Advanced options page, specify any manifest, signature, and output file options, or select **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. The manifest 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. 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 an X.509 certificate that 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. For more information on the different file formats, see Open Virtualization Format (OVF and OVA).

    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 GiB of space has a hard disk image that consumes 26 GiB of space, regardless of whether the VM actually requires that much space. Compressing the VHD files makes the export process take longer to complete, and importing a package containing compressed VHD files takes longer as the **Import** wizard must extract 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 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. Select **Finish** to begin exporting the selected VMs and close the wizard.

   The export process can take some time. The export speed depends 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, select **Notifications** > **Events**, find the export in the list of events, and select **Cancel**.

Export VMs as XVA

May 25, 2023

You can export a single VM as an XVA file using the **Export** wizard. Shut down or suspended before you attempt to export them.

> **Note:**
>
> It might 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 CPU, then exported, might 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, select **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. Click **Next**.

3. On the **Virtual Machines** page, the VM to be exported is selected in the list. When exporting as XVA, you can only select one VM from 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, check **Verify export on completion**. Click **Finish** to begin exporting the selected VM and close the wizard.

   The export process can take some time. The speed of the export depends 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, select **Notifications** > **Events**. Find the export in the list of events. Select **Cancel**.

About Snapshots

May 25, 2023

A virtual machine (VM) snapshot is a record of a running VM 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 on the VM temporarily halts while you take the snapshot to ensure that the snapshot captures a self-consistent disk image.

You can create snapshots without first shutting down the VM. This behavior is different to VM exports. 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 to export for backup purposes and then restore, or to use to quickly create VMs.

Snapshots are supported on all storage types.

Types of snapshots

XenCenter supports the following types of VM snapshot:

- disk-only
- disk and memory

> **Note:**
>
> In Citrix Hypervisor 8.0 and earlier versions, quiesced snapshots are also supported.

For more information, see Take a VM Snapshot.

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.

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 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 in the following cases:

- If you are upgrading or patching software
- If you 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

Quiesced snapshots

> **Important:**
>
> In Citrix Hypervisor 8.1 and later, quiesced snapshots are not supported.

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 performance impact on a system while they are being taken. They might also fail under load, so more than one attempt to take the snapshot might be required.

Accessing orphaned snapshots

If you take snapshots of a VM and later 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

May 25, 2023

Perform the following steps to take a snapshot of a Virtual Machine:

1. Select the VM in the **Resources** pane and then click the **Snapshots** tab.
2. Click **Take Snapshot**. Alternatively,
   - Right-click in the **Properties** pane and select **Take Snapshot**.
   - On the **VM** menu, select **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 disk and memory snapshot, select **Snapshot the virtual machine's disks and memory**.
   - **Note:** In Citrix Hypervisor 8.0 and earlier versions, quiesced snapshots of Windows VMs are also supported. You can choose to **Quiesce the VM before taking the snapshot**.
5. Click **OK** to begin creating the snapshot. Progress is displayed on the status bar and on the **Events** view under **Notifications**.

   When the new snapshot has been created, it appears 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

May 25, 2023

Reverting to a snapshot restores the VM to the state it was in at the point in time when you created the snapshot. All changes made to the VM since the snapshot was taken are discarded. The current state of the VM is lost.

The **Revert to Snapshot** dialog box includes an option to take a snapshot of the current VM state before you revert to the earlier snapshot. This option allows you to easily restore the VM to its current state again if you need to.

1. On the **Snapshots** tab, select the snapshot and select **Revert To**.

   If the snapshot you want to revert to is a scheduled snapshot, make scheduled snapshots visible on the **Snapshots** tab before you can select it. To make these snapshots visible, select **View** > **Scheduled Snapshots**.

2. To take a new snapshot of the current state of the VM before reverting it back to the earlier snapshot, select the check box.
3. Select **Yes**.

Create a New VM From a Snapshot

May 25, 2023

You create a VM from a snapshot in the same way as you create a VM from a regular VM template, by using the **New VM** wizard.

To create a VM from a snapshot

1. On the **Snapshots** tab, select the snapshot you want to use, then right-click and select **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 VM. For more information, see Creating a new VM.

To create a 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. Select to expand the **Types** group and then expand the **Snapshots** group.
3. Select the snapshot, then right-click and select **New VM from Snapshot** on the shortcut menu.

Create a New Template From a Snapshot

May 25, 2023

While you cannot copy a VM snapshot directly, you can create a VM template from a snapshot and then use that to make copies of the snapshot. Templates are a "gold image"- ordinary VMs that you use as master copies from which to create VMs. After you have set up a VM the way you want it and taken a snapshot of it, save the snapshot as a new template. Use this new template to create copies of your specially configured VM in the same resource pool. The snapshot's memory state is not saved when you do this action.

To save a snapshot as a new template

1. On the **Snapshots** tab, select the snapshot, right-click, and select **Create Template from Snapshot** on the shortcut menu.
2. Enter the name of the new template and then click **OK**.

   After the new template has been successfully created, it appears as a custom template in the **Resources** pane. It also appears 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. Select to expand the **Types** group and then expand the **Snapshots** group.
3. Select the snapshot, right-click, and select **Create Template from Snapshot** on the shortcut menu.

Export a Snapshot to a File

May 25, 2023

When you export a VM snapshot, it is saved as a VM template in an XVA file on your XenCenter system. This template contains a complete copy of the snapshot (including disk images). You can then import the template and use it to create a 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, select **Actions**, and then select **Export Snapshot as Template**.

2. Browse to locate the folder where you want to create the XVA file, enter the file name, then select **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. Select 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 file name, then select **Save** to begin the export.

Delete a Snapshot

May 25, 2023

To delete a snapshot

1. On the **Snapshots** tab, select the snapshot. Select **Delete**.
2. Select **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. Select to expand the **Types** group and then expand the **Snapshots** group.
3. Select the snapshot, right-click and then select **Delete Snapshot** on the shortcut menu.

Scheduled Snapshots

May 25, 2023

The scheduled snapshots 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.

Scheduled snapshots work by having pool-wide snapshot schedules for selected VMs in the pool. When a snapshot schedule is enabled, snapshots of the specified VM are taken at the scheduled time each hour, day, or week.

Several scheduled snapshots can be enabled in a pool, covering different VMs and with different schedules. A VM can be assigned to only one snapshot schedule at a time.

XenCenter provides a range of tools to help you use this feature:

- To define a Scheduled Snapshot, use the **New snapshot schedule** wizard.
- To enable, disable, edit, and delete scheduled snapshots for a pool, use the **VM Snapshot Schedules** dialog box.
- To edit a snapshot schedule, open its Properties dialog box from the **VM Snapshot Schedules** dialog box.
- To revert a VM to a scheduled snapshot, select the snapshot on the **Snapshots** tab and revert the VM to it.

Create Scheduled Snapshots

May 25, 2023

Use the **New snapshot schedule** wizard to create a **Snapshot Schedule** that lets you specify the following information:

- The VMs in the pool to snapshot
- The type of snapshot to be created (disk-only or disk and memory)
- The snapshot schedule.

To open the **New snapshot schedule** wizard: on the **Pool** menu, select **VM Snapshot Schedules**, and then select **New** to start the wizard.

- **Schedule name**: Enter a name for the snapshot schedule. Optionally, provide a description.
- **VMs in the snapshot schedule**: Select the VMs you would like to add to the snapshot schedule.
- **Snapshot Type**: Choose the type of snapshot you would like to take.

  Scheduled 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). This type of snapshot can be large.
  - **Note:** In Citrix Hypervisor 8.0 and earlier versions, quiesced snapshots of Windows VMs are also supported. You can choose to **Quiesce the VM before taking the snapshot**.

- **Snapshot schedule**: Choose how often you would like to schedule a snapshot.

  Snapshot schedule options:

  - **Hourly snapshots**: A snapshot of the specified VM or VMs is taken each hour at the specified time.
  - **Daily snapshots**: A snapshot of the specified VM or VMs is taken each day at the specified time.
  - **Weekly snapshots**: A snapshot of the specified VM or VMs is taken at the specified time on the specified days of the week. You can select one or more days.

  In the **Number of snapshots to keep** section, specify the number of snapshot schedules to retain. When the number of scheduled snapshots taken exceeds this value, the oldest snapshot is deleted automatically.

  > **Note:**
  >
  > You can take up to 10 scheduled snapshots per VM.

Quiesced snapshots

In Citrix Hypervisor 8.1 and later, scheduled quiesced snapshots are no longer supported. If you have an existing snapshot schedule for quiesced snapshots that you created with an earlier version, this scheduled snapshot fails in Citrix Hypervisor 8.1 and later. Delete this snapshot schedule and create a snapshot schedule that creates a supported type of snapshot.

Manage scheduled snapshots

May 25, 2023

To enable, disable, edit, and delete **scheduled snapshots** for a pool, use the **VM Snapshot Schedules** dialog box on the **Pool** menu, select **VM Snapshot Schedules**.

Enabling a snapshot schedule

When you enable a **Snapshot Schedule**, you turn it "on". Automated snapshots of the specified VMs are generated at the scheduled time. Scheduled snapshots are taken until the schedule is disabled.

To enable a 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 Snapshot Schedules**.
2. Select a snapshot schedule from the list of snapshot schedules defined in the pool and select **Enable**.

Disabling a snapshot schedule

If you want to stop automated snapshots from being taken, you can disable the **Snapshot Schedules** using the **VM Snapshot Schedules** dialog box: a disabled snapshot schedule can be enabled again at any time.

To disable a 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 Snapshot Schedules**.
2. Select the snapshot schedule from the list of snapshot schedules defined in the pool and select **Disable**.

Editing a snapshot schedule

You can change the properties of a snapshot schedule, 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. From the **Pool** menu, select **VM Snapshot Schedules**.
2. Select the snapshot schedule from the list of snapshot schedules defined in the pool.
3. Click **Properties** and go to the tab you need:
    - **General** Change the name and description of the snapshot schedule.
    - **Custom Fields** Define extra text and date/time fields for your snapshot schedules.
    - **VMs in the snapshot schedule** Select a VM from the list to add it to the snapshot schedule.
    - **Snapshot Type** Change the type of snapshot to create.
    - **Snapshot Schedule** Change the schedule used to take VM snapshots and change the snapshot retention value.
4. Click **OK** to save your changes and close the **Properties** window.

Deleting a snapshot schedule

To delete a Snapshot Schedule:

1. Select the pool, or any server or VM in the pool in the **Resources** pane. On the **Pool** menu, select **VM Snapshot Schedules**.
2. Select schedule from the list of snapshot schedules and select **Delete**.

Revert VMs to Snapshots

May 25, 2023

To revert a VM to a scheduled snapshot:

1. Select the VM and select the **Snapshots** tab.
2. To view scheduled snapshots, select **View** and then select **Scheduled Snapshots**. The Snapshots tab does not display scheduled snapshots by default.
3. Select the scheduled snapshot you want to revert the VM to and then select **Revert To**.
4. To take a new snapshot of the current state of the VM before reverting it back to the snapshot schedule, select the check box.
5. Click **Yes** to revert the VM back to the selected snapshot.

Citrix Hypervisor vApps

May 25, 2023

Citrix Hypervisor vApp: 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 Citrix Hypervisor server and are distributed within a pool using the normal rules.

When a vApp is started, the VMs contained within it start in a user-defined order, allowing VMs which depend upon one another to be automatically sequenced. This capability means that you do not need to manually sequence the startup of dependent VMs when the whole service requires a restart.

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 are not 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 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.
![Shutdown icon - a red circle with a power icon overlaid in white.]](/en-us/xencenter/media/001-shut-down-h32bit-16.png) **Shut Down**	Shut down all 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

May 25, 2023

To create a vApp, use the Manage vApps dialog box.

1. Select the pool and, from the **Pool** menu, select **Manage vApps**.

   Alternatively, right-click in the **Resources** pane and select **Manage vApps** on the shortcut menu.

2. Select **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 select **Next**. You can choose any name you like, but a descriptive name is best. Although it is advisable to avoid having multiple vApps with the same name, it is not a requirement. XenCenter does not enforce any unique-

ness 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 select **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 select **Next**.

Value	Description
Start order	Specifies the order in which individual VMs are started up within the vApp, allowing certain VMs to be restarted before others. VMs with a start order value of 0 (zero) start first, then VMs with a start order value of 1, and so on.
Attempt to start next VM after	This value is a delay interval that specifies how long to wait after the VM starts before the next group of VMs in the startup sequence start.

> **Note:**
>
> 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. Select **Previous** to go back and modify any settings, or **Finish** to create the vApp and close the wizard.

Modify vApps

May 25, 2023

You can use the **Manage vApps** dialog to take the following actions:

- Change the name or description of a vApp
- Add or remove VMs from the vApp
- Change the startup sequence of the VMs in the vApp

Using the Manage vApps dialog

1. Select the pool and, on the **Pool** menu, select **Manage vApps**.

   Alternatively, right-click in the **Resources** pane and select **Manage vApps** on the shortcut menu.

2. Select the vApp and select **Properties** to open its **Properties** dialog box.

3. Select the **General** tab to change the vApp name or description.
4. Select the **Virtual Machines** tab to add or remove VMs from the vApp.
5. Select 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 are started up within the vApp, allowing certain VMs to be restarted before others. VMs with a start order value of 0 (zero) start first, then VMs with a start order value of 1, and so on.
Attempt to start next VM after	This value is a delay interval that specifies how long to wait after the VM starts before the next group of VMs in the sequence start.

> **Note:**
>
> The shutdown order of VMs in a vApp is always the reverse of the configured start order.

6. Select **OK** to save your changes and close the **Properties** dialog box.

Delete a vApp

May 25, 2023

To delete a vApp from a pool, use the **Manage vApps** dialog box.

1. Select the pool and, on the **Pool** menu, select **Manage vApps**.

   Alternatively, right-click in the **Resources** pane and select **Manage vApps** on the shortcut menu.
2. Select the vApp you want to delete from the list, then select **Delete**.

   The VMs in the vApp are not deleted.

Start and Shut Down vApps

May 25, 2023

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. You can set these values when you create the vApp and change them at any time from the **vApp Properties** dialog or from the individual **VM Properties** dialog.

For more information, see:

- Create vApps
- vApp Properties
- VM Properties

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, select **Manage vApps**.

   Alternatively, right-click in the **Resources** pane and select **Manage vApps** on the shortcut menu.

2. Select the vApp and select **Start** to start all 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, select **Manage vApps**.

   Alternatively, right-click in the Resources pane and select **Manage vApps** on the shortcut menu.

2. Select the vApp and select **Shut Down** to shut down all VMs in the vApp.

   

   A soft shutdown is attempted on all VMs. If a soft shutdown is not possible, then a forced shutdown is performed. For more information about soft and forced VM shutdowns, see Shut down a VM.

Export and Import vApps

May 25, 2023

To export a vApp

vApps can be exported as OVF/OVA packages.

1. Open the **Manage vApps** dialog box: on the **Pool** menu, select **Manage vApps**.
2. Select the vApp you want to export in the list and select **Export**.

   

3. Follow the procedure described in Export VMs as OVF/OVA.

   Exporting a vApp can 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, select **Manage vApps**.
2. Select **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

May 25, 2023

Citrix Hypervisor offers a range of features to enable you to protect your VMs and vApps.

High Availability

High availability protects against downtime of critical VMs caused by the failure of individual servers in a pool. This feature guarantees that VMs are automatically restarted on an alternate server in the same pool, with minimal service interruption. Citrix Hypervisor constantly replicates the pool database across all nodes. The pool database is also backed up to shared storage on the heartbeat SR for extra safety.

For more information, see the following articles:

- About high availability
- High availability requirements
- VM startup settings
- Configure high availability
- Disable high availability
- Change high availability settings

Disaster recovery

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 through mirrored storage. If a disaster occurs at your primary site, DR can recover your VMs and vApps from the mirrored storage to a pool on a secondary (DR) site.

For more information, see the following articles:

- About Citrix Hypervisor DR
- Configuring DR
- Failover
- Failback
- Test Failover

High availability

May 25, 2023

Citrix Hypervisor high availability enables VMs to restart automatically in the event of an underlying hardware failure or loss of any server. High availability is about making sure important VMs are always running in a resource pool. With high availability enabled, if one of your servers fails, its VMs restart on other servers in the same pool. This capability allows essential services to be restored with minimal service interruption in the event of system or component failure.

If the pool master server fails, Citrix Hypervisor high availability selects a new server to take over as master. Any server in a pool can be a master server. Citrix Hypervisor replicates the pool database constantly across all nodes. It also backs up the database to shared storage on the heartbeat SR for extra safety.

There are two key aspects to Citrix Hypervisor high availability:

- Reliably detecting server failure
- Computing a failure plan to enable swift recovery

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 high availability incorrectly decides a master server has broken down and elects a new master, there might be unpredictable results if the original server returns. Similarly, if a network issue causes the pool to split into two equal halves, we must ensure only one half accesses the shared storage and not both simultaneously. Citrix Hypervisor solves all these problems by having two mechanisms: a storage heartbeat and a network heartbeat.

When you enable high availability in a pool, you nominate an iSCSI, Fibre Channel or NFS storage repository to be the heartbeat SR. Citrix Hypervisor 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 high availability starts up, all servers exchange data over both network and storage channels. This action indicates which servers they can see over both channels and demonstrates which I/O paths are working and which are not. This information is exchanged until a fixed point is reached and all servers in the pool agree about what they can see. When this agreement happens, high availability is enabled and the pool is protected. This high availability arming process can take a few minutes to settle for larger pools, but is only required when you first enable high availability.

After high availability is active, each server regularly writes storage updates to the heartbeat virtual disk, and network packets over the management interface. Ensure that network adapters are bonded for resilience, and that storage interfaces are using dynamic multipathing where it is supported. This configuration ensures that any single adapter or wiring failures do not result in any availability issues.

For more information, see:

- Network bonding
- Storage multipathing

Server fencing

The worst-case scenario for high availability is one where a server is thought to be off-line but is still writing to the shared storage. This scenario can result in corruption of persistent data. Citrix Hypervisor uses server fencing to prevent this situation. The server is automatically powered off and isolated from accessing any shared resources in the pool. Fencing prevents the failing server from writing to shared disks. This behavior prevents damage to the stored data during an automated failover, when protected virtual machines are being moved to other servers in the pool.

Servers 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 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 must be restarted on a server with working networking. If the network partitions are the same size, then only one of them self-fences according to a stable selection function.
- If the storage heartbeat goes away but the network heartbeat remains, the servers check to see if they can see all other servers over the network. If this condition holds true, the servers remain running on the assumption that the storage heartbeat server has gone away. This action doesn't compromise VM safety, but any network glitches 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 high availability: capacity planning for failure.

A resource pool consists of several servers (say, 32), each with potentially different amounts of memory and a different number of running VMs. Citrix Hypervisor high availability dynamically computes a failure plan that calculates the actions to be taken on any server failure. This failure plan ensures that no single server failure makes it impossible to restart its VMs on another server (for example, due to insufficient memory on other servers). In addition to dealing with failure of a single server, Citrix Hypervisor high availability can deal with the loss of multiple servers in a pool. For example, high availability can handle when failure of a network partition takes out an entire group of servers.

In addition to calculating what actions are 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 high availability plan for a pool:

- Maximum failure capacity. This value is the maximum number of servers that can fail before

there are insufficient resources to run all the protected VMs in the pool. To calculate the maximum failure capacity, Citrix Hypervisor considers:

  - The restart priorities of the VMs in the pool
  - The number of servers in the pool
  - The server CPU and memory capacity

- Server failure limit. You can define this value as part of the high availability configuration which specifies the number of server failures to allow in the pool, within the plan. For example, when the server failure limit for a pool is 3, Citrix Hypervisor calculates a failover plan that allows any 3 servers to fail and all protected VMs can still run in the pool. You can configure the server failure limit to a value that is lower than the maximum failure capacity, making it less likely that the pool becomes overcommitted. This configuration can be useful in an environment with RBAC enabled. For example, this setting allows RBAC users with lower permissions than Pool Operator to bring more VMs online without breaking the high availability plan. For more information, see the *High availability and Role-Based Access Control (RBAC)* section.

A system alert is generated when the maximum failure capacity value falls below the value specified for the server failure limit.

Overcommit protection

When high availability is first enabled on a pool, a failure plan is calculated based on the resources available then. Citrix Hypervisor high availability 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, the pool becomes overcommitted. Examples of insufficient resources might be not enough free memory or changes to virtual disks and networks that affect which VMs might be restarted on which servers.

High availability restart priority is used to determine which VMs to start 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, the maximum failure capacity for the pool is recalculated dynamically. This information enables you to try various combinations of VM restart priorities depending on your business needs. You can 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 is displayed in XenCenter. The message can also be sent to an email address, if configured. You are given the option to cancel the operation, or proceed anyway, causing the pool to become overcommitted.

Working with an HA-enabled pool

The best practice for high availability is not to make configuration changes to the pool while high availability is enabled. Instead, it is intended to be the "2am safeguard"which restarts 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, disable high availability during these changes.

- If you try to shut down a protected VM from XenCenter, XenCenter offers the option of removing the VM from the failure plan and then shutting it down. This option ensures that accidental VM shutdowns do not result in downtime, but that you can still stop a protected VM if you really want to.
- If you must reboot a server when high availability is enabled, XenCenter automatically uses the VM restart priorities to determine if this reboot invalidates the pool failure plan. If it doesn't affect the plan, then the server is shut down normally. If the plan is violated, but the maximum failure capacity is greater than 1, XenCenter provides the option of lowering the pool's server failure limit by 1. This action reduces the overall resilience of the pool, but always ensures that at least one server failure is 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 high availability on the pool by selecting **Turn HA off**. You can re-enable high availability after the update has been installed. If you do not disable high availability, the update does not proceed. Monitor the pool manually while updates are being installed to ensure that server failures do not disrupt the operation of the pool.
- When high availability is enabled, some operations that can compromise the plan for restarting VMs might be disabled, such as removing a server from a pool. To perform these operations, temporarily disable high availability or you can shut down the protected VMs before continuing.

High availability and role-based access control (RBAC)

In Citrix Hypervisor environments where role-based access control (RBAC) is implemented, not all users are permitted to change a pool's high availability configuration settings. For example, VM Operators do not have sufficient permissions to adjust the failover capacity for an HA-enabled pool. If starting a VM reduces the maximum number of allowed server failures to a value lower than the current value, a VM Operator cannot start the VM. Only Pool Administrator or Pool Operator-level users can configure the number of server failures allowed.

In this case, the Pool Administrator or Pool Operator can set the server failure limit to a number that is lower than the maximum number of failures allowed. This setting creates slack capacity and so ensures that less privileged users can start up new VMs. It reduces the pool's failover capacity without threatening the failure plan.

High availability Requirements

May 25, 2023

Before you can configure high availability 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. The shared storage must include at least one iSCSI, Fibre Channel or NFS LUN of 356 MiB or greater that is used for the heartbeat SR. If you are using a NetApp or EqualLogic storage array, manually provision an iSCSI LUN on the array to use for the heartbeat SR.
- We strongly recommend that you use a bonded management interface on the servers in the pool.
- We strongly recommend that you use multipath storage for the heartbeat SR.
- Adequate licenses must be installed on all servers.
- All the virtual machines you want to protect with high availability must be agile. This means:
  - Virtual disks must be on shared storage. You can use any type of shared storage to store the virtual disks. The iSCSI, Fibre Channel or NFS LUN is only required for the storage heartbeat. These SRs can be used for virtual disk storage, if you prefer, but it is not necessary.
  - Virtual network interfaces must be on pool-wide networks.
  - Do not configure a connection to any local DVD drive.

  VMs that are not agile can be assigned only **Restart if possible** restart priority. These VMs are tied to one server. For example, a VM with a physical CD drive mapped in from a server can only run on the server with the CD drive.

VM Restart Settings

May 25, 2023

If more servers fail than have been planned for, then a high availability recovery operation begins. The **HA restart priority** is used to determine which VMs are restarted. The start order and delay interval values determine the order in which individual VMs are started. These settings ensure that the most important VMs are restarted first.

High availability restart priority

The **HA restart priority** specifies which VMs restart under the high availability failure plan for a pool:

- **Restart** - VMs with this priority are guaranteed to be restarted if sufficient resources are available within the pool. They restart 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 reliably restart, the pool is overcommitted.

- **Restart if possible** - VMs with this restart priority are not considered when calculating a failure plan. However, an attempt to restart these VMs is made if a server that is running them fails. This restart is attempted after all higher-priority VMs are restarted. If the attempt to start a restart-if-possible VM fails because there is not capacity to start the VM, it is not retried.

  This setting is useful for test/development VMs which aren't critical to keep running, but would be nice to do so.

- **Do not restart** - No attempts are made to restart VMs with this priority.

Start order

The **Start order** property specifies the order in which individual VMs start up during a recovery operation. This setting allows certain VMs to be started before others. VMs with a start order value of 0 (zero) start first, then VMs with a start order value of 1, and so on.

Delay interval (Attempt to start next VM after)

The **Attempt to start next VM after** property specifies how long the recovery process waits after the VMs start before starting the next VMs in the startup sequence. The next group of VMs are those VMs with a later start order.

Configure high availability

May 25, 2023

You enable high availability for a resource pool using the **Configure HA** wizard. The wizard takes you through the high availability configuration process, step-by-step. During this process, the wizard calculates the server failure limit for the pool given the available resources and the high availability restart priorities you specify.

To open the **Configure HA** wizard: in XenCenter, select the pool, select the **HA** tab, and then select **Configure HA**.

Alternatively:

- On the **Pool** menu, select **High Availability**.
- Right-click in the **Resources** pane and then select **High Availability** on the shortcut menu.

To configure high availability on a pool:

1. Ensure that the high availability requirements are satisfied. For more information, see High availability requirements.
2. Open the **Configure HA** wizard.
3. Click **Next** on the first page of wizard. The wizard scans the pool for a shared iSCSI, Fibre Channel or NFS LUN to use as the heartbeat SR. If no suitable SR is found, configure some appropriate new storage before you continue.
4. On the **Heartbeat SR** page, choose an SR from the list and then click **Next**.
5. 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.

   Set the following options:

   - **HA restart priority**: Choose a restart priority for each VM:
     - Choose **Restart** to ensure the selected VMs 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 start during the recovery operation, allowing certain VMs to be started before others. VMs with a start order value of 0 (zero) start first, then VMs with a start order value of 1, and so on.
   - **Attempt to start next VM after**: Set a delay interval to wait after starting the VM before attempting to start the next group of VMs in the startup sequence. The next group of VMs are those VMs with a lower start order.
6. Also on the **HA Plan** page, under **Server failure limit**, you can set the number of server failures to allow within this high availability plan. Ensure this value is 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 cannot continue until you resolve the situation. To stop the pool being overcommitted, either adjust the high availability restart priorities or make more resources available within the pool. For more information, see To increase the maximum failure capacity for a pool. Click **Next** when you have finished high availability plan configuration.

7. On the last page of the wizard, review your high availability configuration settings. Click **Back** to go back and change any of the settings or click **Finish** to enable high availability and close the wizard.

Disable high availability

May 25, 2023

When high availability is enabled, some operations that can compromise the plan for restarting VMs might be disabled, such as removing a server from a pool. To perform these operations, you can temporarily disable high availability.

To disable high availability:

1. Select the pool in the **Resources** pane, select the **HA** tab, and then select **Disable HA**.
2. Click **OK** to confirm. The **VM startup** settings specified for each VM in the pool are stored and remembered if you turn high availability back on again later.

Change high availability Settings

May 25, 2023

When a pool has high availability enabled, use the **Configure HA** dialog box to change VM startup settings and the server failure limit for the pool.

To change high availability restart priority and VM startup sequence settings

1. Select the pool in the **Resources** pane, choose the **HA** tab, and then select **Configure HA**. Alternatively:

   - On the **Pool** menu, select **High Availability**.
   - Right-click in the **Resources** pane and then select **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.

   Set the following options:

   - **HA restart priority**: Choose a restart priority for each VM:

  - Choose **Restart** to ensure the selected VMs 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 start up during the recovery operation, allowing certain VMs to be started before others. VMs with a start order value of 0 (zero) start first, then VMs with a start order value of 1, and so on.
- **Attempt to start next VM after**: Set the delay interval to wait after starting the VM before starting the next group of VMs in the startup sequence. The next group of VMs are those 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, select the **HA** tab, and then click **Configure HA**. Alternatively:
   - On the **Pool** menu, select **High Availability**.
   - Right-click in the **Resources** pane and then select **High Availability** on the shortcut menu.
2. Under **Server failure limit**, enter the number of server failures to allow. Ensure this value is 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 cannot save the change. To be able to save the change, either adjust the high availability restart priorities or make more resources available within the pool. For more information, see the following section.
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 high availability 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)

November 16, 2023

The disaster recovery (DR) feature allows you to recover 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 High Availability. High availability restarts VMs on an alternate server in the same pool.

Understanding DR

Disaster recovery stores all the information needed to recover your business-critical VMs and vApps on storage repositories (SRs). These storage repositories are then replicated from your primary (production) environment to a backup environment. When a protected pool at your primary site fails, the VMs and vApps in that pool can be recovered from the replicated storage and recreated on a secondary (DR) site. The result is minimal application or user downtime.

After the recovered VMs are up and running in the DR pool, the DR pool metadata must also be saved on storage that is replicated. This action allows recovered VMs and vApps to be restored back to the primary site when it is back online.

> **Note:**
>
> Disaster Recovery can only be used with LVM over HBA or LVM over iSCSI storage types.

Citrix Hypervisor VMs consist 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. The metadata contains all the information required to recreate the VM if the original VM is unavailable or corrupted. Most metadata is written when the VM is created and is updated only when you change 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, networking, and storage configuration. It also includes the VM startup options used when restarting the VM in a high availability or DR environment: start order, delay interval, and restart priority. For example, when recovering VMs, the VMs within a vApp restart in the DR pool in the order and with the delay intervals specified in the metadata.

> **Note:**
>
> To use Disaster Recovery, you must be logged in as root or have a role of Pool Operator or higher.

Disaster recovery terminology

vApp: A logical group of related VMs which are managed as a single entity.

Site: A physical group of Citrix Hypervisor 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 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. For VMs, the configuration information includes UUID, memory, virtual CPU, networking and storage configuration, and startup options. 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.

Disaster recovery infrastructure

To use Disaster Recovery, set up the appropriate DR infrastructure 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, varies from device to device. We recommend you use your storage solution to handle storage replication.
- After recovered VMs and vApps are up and running on a pool on your DR site, replicate the SRs containing the DR pool metadata and virtual disks. This action allows 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. However, the Citrix Hypervisor environment must be at the same release and patch level. Also, sufficient

resources must 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. Ensure 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, disable the mirror functionality before VMs are restarted on the recovery site.

Failover, failback, and test failover with the Disaster Recovery wizard

The **Disaster Recovery** wizard makes failover and failback simple. The steps involved in these processes are outlined here:

Failover

1. Choose a target pool on your secondary DR site to which you want to recover your VMs and vApps.
2. Provide details of the storage targets containing the replicated SRs from your primary site. The wizard scans the targets and lists all SRs found there.
3. Select the SRs containing the metadata and virtual disks for the VMs and vApps you want to recover. The wizard scans the SRs and lists all the VMs and vApps found.
4. Select which VMs and vApps you want to recover to the DR site. Specify whether you want the wizard to start them up automatically when they have been recovered, or whether you prefer to wait and start them up manually yourself.

   The wizard performs prechecks to ensure that the selected VMs and vApps can be recovered to the target DR pool. For example, the wizard checks that all the storage required by the selected VMs and vApps is available.

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. 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. Provide details of the storage targets containing the replicated SRs from your DR site. The wizard scans the targets and lists all SRs found.
3. Select the SRs containing the metadata and virtual disks for the VMs and vApps you want to restore. The wizard scans the SRs and lists all the VMs and vApps found.
4. Select which VMs and vApps you want to restore back to the primary site. Specify whether you want the wizard to start them up automatically when they have been recovered, or whether you prefer to wait and start them up manually yourself.

   The wizard then performs prechecks to ensure that the selected VMs and vApps can be recovered to the target pool on the primary site. For example, the wizard checks that all the storage required by the selected VMs and vApps is available.

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 in a two or more places, it uses only the most recent information per VM. For example, the information might be stored on the primary site storage, the DR site storage, and in the pool the data is imported into.

> **Tip:**
>
> To make recovering VMs and vApps easier, name your SRs to indicate 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, the steps are the same as for failover, but recovered VMs and vApps are started up in a paused state on the DR site. Cleanup is performed when the test is finished to remove all VMs, vApps, and storage recreated on the DR site. For more information, see Test Failover.

Configuring disaster recovery

November 16, 2023

Use the XenCenter **Configure DR** dialog box to choose the storage repositories (SRs) to store the metadata for a pool. This metadata includes configuration information about all the VMs and vApps in the pool. Pool metadata is updated whenever you change the VM or vApp configuration.

> **Note:**
>
> Citrix Hypervisor DR supports only LVM over HBA or LVM over iSCSI storage types.

To configure disaster recovery on the primary pool:

1. On your primary site, select the pool that you want to protect.
2. On the **Pool** menu, choose **Disaster Recovery** and click **Configure**.
3. Select up to eight SRs where the pool metadata is stored. A small amount of space is required on this storage for a new LUN which contains the pool recovery information.
4. Click **OK**.

> **Important:**
>
> To configure your DR environment fully, replicate the SRs that contain the pool metadata and VM virtual disks from your production environment to a backup environment. Storage replication cannot be configured from within XenCenter. Use your storage solution to handle storage replication, for example using mirroring. The replication varies from device to device.

Failover

November 16, 2023

Failover recovers VMs and vApps to a secondary site in the event of disaster at your primary site. To fail over your critical VMs and vApps, use the **Disaster Recovery** wizard.

> **Important:**
>
> The **Disaster Recovery** wizard does not control any storage array functionality. Disable duplication (mirroring) of the metadata storage and the storage used by the VMs to be restarted 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, break mirroring before you attempt to recover data. This action gives the secondary site 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 want to recover. Use the **Power state after recovery** option to specify whether you want the wizard to start the recovered VMs and vApps immediately. Alternatively, you can wait and start the VMs and vApps manually after failover is complete.

5. Click **Next** to progress to the next wizard page and begin failover prechecks.

   The wizard performs pre-checks before starting failover. For example, the wizard ensures 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.

6. Resolve any issues on the prechecks page, and then click **Failover** to begin the recovery process.

   A progress page is displayed showing whether recovery was successful for each VM and vApp.

   Failover can take some time depending on the number of VMs and vApps you are recovering. During this process, the following actions occur:

   - Metadata for the VMs and vApps is exported from the replicated storage
   - VMs and vApps are re-created in the primary pool
   - SRs containing the virtual disks are attached to the re-created VMs
   - VMs are started, if specified

7. When the failover is complete, click **Next** to see the summary report.

8. Click **Finish** on the summary report page to close the wizard.

After your primary site is available again, you can return the VMs and vApps to that site. To complete this process, follow the **Disaster Recovery** wizard again, but instead select the **Failback** option. For more information, see Failback.

Failback

November 16, 2023

Failback restores VMs and vApps from replicated storage back to a pool on your primary site. Failback occurs 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. Disable duplication (mirroring) of the metadata storage and the storage used by the VMs which are to be restored 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, break mirroring before you attempt to recover data. This action gives the primary site 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. Choose the VMs and vApps that you want to restore. Use the **Power state after recovery** option to specify whether to start the restored VMs and vApps automatically. Alternatively, you can wait and start the VMs and vApps manually after failback is complete.
5. Click **Next** to progress to the next wizard page and begin failback prechecks.

   The wizard performs pre-checks before starting failback. For example, the wizard ensures 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.

6. Resolve any issues on the prechecks page, and then click **Failback** to begin the recovery process.

   A progress page is displayed showing whether restoration was successful for each VM and vApp. Failback can 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.
8. Click **Finish** on the summary report page to close the wizard.

Test Failover

November 16, 2023

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. However, instead of starting after they are recovered to the DR site, the VMs and vApps are placed in a paused state. At the end of test failover, the wizard automatically removes all VMs, vApps, and storage recreated on the DR site.

After initial DR configuration, verify that failover works correctly by performing a test failover. We recommend you also do a test failover after you make significant configuration changes in a DR-enabled pool.

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, break mirroring before you attempt to recover data. This action gives the secondary site 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 want to recover.
5. Click **Next** to progress to the next wizard page and begin failover prechecks.

   Before beginning the test failover process, the wizard performs pre-checks. For example, the checks ensure 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, disable high availability on the secondary pool. This action ensures that high availability does not automatically start the recovered VMs and vApps after recovery. To disable high availability on the secondary pool, click **Disable HA** on this page. (If the wizard disables high availability at this point, it is enabled again automatically at the end of the test failover process.)

6. Resolve any issues on the pre-checks page, and then click **Failover** to begin the test failover.

   A progress page is displayed showing whether recovery was successful for each VM and vApp. Failover can take some time depending on the number of VMs and vApps you are recovering. The following actions occur during this process:

   - Metadata for the VMs and vApps is recovered from the replicated storage
   - VMs and vApps are re-created in the DR pool
   - SRs containing the virtual disks are attached to the re-created VMs
   - **The recovered VMs are placed in a paused state. The VMs are not started 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 are removed.
   - Storage that was recovered during the test failover is detached.
   - If the wizard disabled high availability on the DR pool at the prechecks stage to allow the test failover to take place, it is enabled again.

   The progress of the cleanup process is displayed in the wizard.

8. Click **Finish** to close the wizard.

Access Control (AD and RBAC)

May 25, 2023

Managing Users

May 25, 2023

When you first install Citrix Hypervisor, a user account is added to Citrix Hypervisor automatically. This account is the local super user (LSU), or `root`, which the Citrix Hypervisor system authenticates locally. You can create extra users by adding Active Directory accounts from the **Users** tab in XenCenter.

> **Note:**
>
> The term "user"refers to anybody with a Citrix Hypervisor account, that is, anyone administering Citrix Hypervisor hosts, regardless of the 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 feature allows Citrix Hypervisor 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 use Active Directory and some don't.

When you create a user in Citrix Hypervisor you must first assign a role to the newly created user before they can use the account. Citrix Hypervisor **does not** automatically assign a role to the newly created user. As a result, these accounts do not have any access to the Citrix Hypervisor 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 Citrix Hypervisor environment

Even though the Citrix Hypervisor servers are Linux-based, Citrix Hypervisor lets you use Active Directory accounts for Citrix Hypervisor 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 Citrix Hypervisor, Active Directory users and groups become Citrix Hypervisor subjects, called users in XenCenter. When a subject is registered with Citrix Hypervisor, users and groups are authenticated with Active Directory on login. These users and groups do not need to qualify their user name with a domain name.

To qualify a user name, you must enter the user name in Down-Level Logon Name format, for example, `mydomain\myuser`.

> **Note:**
>
> By default, if you don't qualify the user name, XenCenter attempts to log users in to Active Directory authentication servers using the domain to which it is joined. The exception to this rule is the LSU account, which XenCenter always authenticates locally (that is, on Citrix Hypervisor) 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 Citrix Hypervisor, 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.

Role Based Access Control overview

November 16, 2023

The Role Based Access Control (RBAC) feature lets you assign predefined roles or sets of permissions to Active Directory users and groups. These permissions control the level of access Citrix Hypervisor administrators have to servers and pools. RBAC is configured and deployed at the pool level. Because users acquire permissions through their assigned role, assign a role to a user or their group to give them the required permissions.

Using Active Directory accounts for Citrix Hypervisor user accounts

RBAC lets you restrict which operations different groups of users can perform. This control 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. For more information, see Audit Changes.

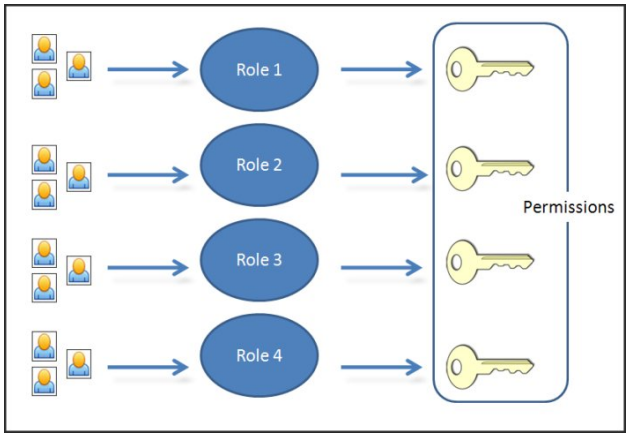

RBAC depends on Active Directory for authentication services. Specifically, Citrix Hypervisor keeps a list of authorized users based on Active Directory user and group accounts. As a result, you must join the pool to the domain and add Active Directory accounts before you can assign roles.

RBAC process

The standard process for implementing RBAC and assigning a user or group a role consists of the following steps:

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 Citrix Hypervisor, the local super user is the default account at installation. The LSU is authenticated by Citrix Hypervisor and not an external authentication service. If the external authentication service fails, the LSU can still log in and manage the system. The LSU can always access Citrix Hypervisor physical server through SSH.

RBAC roles

Citrix Hypervisor 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 Citrix Hypervisor features and settings. They can perform all operations, including role and user management. They can grant access to Citrix Hypervisor console. As a best practice, Citrix recommends assigning this role to a 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.
  >
  > If you remove the Pool Admin role from a user, consider also changing the server root password and rotating the pool secret. For more information, see Pool Security.

- **Pool Operator (Pool Operator)**. This role is designed to let the assignee manage pool-wide resources. Management actions include 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, 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 Citrix Hypervisor 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 user, you must first assign a role to the newly created user before they can use the account. Citrix Hypervisor **does not** automatically assign a role to the newly created user.

Definitions of RBAC roles and permissions

March 6, 2024

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					

Permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Server backup/re-store	X					
Install a TLS certificate on a server	X					
Rolling Pool Upgrade	X					
Import/export OVF/OVA packages; import disk images	X					
Set cores per socket	X	X	X	X		
Convert VMs using Citrix Hypervisor Conversion Manager	X					
Switch-port locking	X	X				
Multipathing	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				
Live migration	X	X	X			

Permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Storage live migration	X	X	X			
VM advanced operations	X	X	X			
VM create/destroy operations	X	X	X	X		
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 Workload Balancing (WLB)	X	X				
Apply WLB optimization recommendations	X	X				

Permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Accept WLB placement recommendations	X	X	X			
Display WLB configuration	X	X	X	X	X	X
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					
Gather diagnostic information	X	X				
vCPU Hotplug	X	X	X	X		
Configure Changed Block Tracking	X	X	X	X		
List changed blocks	X	X	X	X	X	

Permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Configure PVS-Accelerator	X	X				
View PVS-Accelerator configuration	X	X	X	X	X	X
Scheduled Snapshots (Add/Remove VMs to existing Snapshots Schedules)	X	X	X			
Scheduled Snapshots (Add/Modify/Delete Snapshot Schedules)	X	X				

Definitions of permissions

This section provides more 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 can arbitrarily reconfigure the entire system, including RBAC.

Server backup/restore

- Back up and restore servers
- Back up and restore pool metadata

The ability to restore a backup lets the assignee revert RBAC configuration changes.

Install a TLS certificate on a server

This permission enables an administrator to install a TLS certificate on a server that runs Citrix Hypervisor 8.2 or later.

Rolling Pool Upgrade

- Upgrade all hosts in a pool using the Rolling Pool Upgrade wizard.

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 Citrix Hypervisor Conversion Manager

- Convert VMware ESXi/vCenter VMs to Citrix Hypervisor VMs

This permission lets the user convert workloads from VMware to Citrix Hypervisor. Convert these workloads by copying batches of VMware ESXi/vCenter VMs to the Citrix Hypervisor environment.

Switch-port locking

- Control traffic on a network

This permission lets the user block all traffic on a network by default, or define specific IP addresses from which a VM can send traffic.

Multipathing

- Enable multipathing
- Disable multipathing

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 Citrix Hypervisor cancel an in-progress task initiated by any user.

Pool management

- Set pool properties (naming, default SRs)
- Create a clustered pool
- 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 recovery of 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 high availability
- Configure server management interfaces and secondary interfaces
- Disable server management
- Delete crashdumps
- Add, edit, and remove networks
- Add, edit, and remove PBDs/PIFs/VLANs/Bonds/SRs

Live migration

- Migrate VMs from one host to another host when the VMs are on storage shared by both hosts

Storage live migration

- Migrate from one host to another host when the VMs are not on storage shared between the two hosts
- 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

> **Note:**
>
> The VM Admin role can import XVA files only into a pool with a shared SR. The VM Admin role does not have permission to import an XVA file into a server or a pool without shared storage.

VM change CD media

- Eject current CD
- Insert new CD

Import/export OVF/OVA packages; import disk images

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

View VM consoles

Cancel own tasks

- Enables users to cancel their own tasks

Read audit log

- Download Citrix Hypervisor audit log

Apply WLB Optimization Recommendations

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

Modify WLB Report Subscriptions

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

Gather diagnostic information from Citrix Hypervisor

- Initiate GC collection and heap compaction
- Gather garbage collection statistics
- Gather database statistics
- Gather network statistics

Configure changed block tracking

- Enable changed block tracking
- Disable changed block tracking
- Destroy the data associated with a snapshot and retain the metadata
- Get the NBD connection information for a VDI
- Export a VDI over an NBD connection

Changed block tracking can be enabled only for licensed instances of Citrix Hypervisor Premium Edition.

List changed blocks

- Compare two VDI snapshots and list the blocks that have changed between them.

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

Scheduled snapshots

- Add VMs to existing snapshot schedules
- Remove VMs from existing snapshot schedules
- Add snapshot schedules
- Modify snapshot schedules
- Delete snapshot schedules

Join a domain and add users

November 16, 2023

Before you can assign a user or group account an RBAC role, you must add the account to Citrix Hypervisor through RBAC. This process consists of the following tasks:

1. Join the pool or server to the domain. The domain can be one of the following:
   - The domain that the user or group belongs to
   - A domain that is in the same Active Directory forest
   - A domain that has a trust relationship with the user's domain
2. Add the user's Active Directory account or group to Citrix Hypervisor.

After you add the user's Active Directory account or group to Citrix Hypervisor, the user is assigned a fixed role of Pool Admin. In Citrix Hypervisor Premium Edition, you must assign a role to the user or group manually. For more information, see Assign roles to users and groups.

To change domains, leave the current domain and then join the new domain.

To join the Citrix Hypervisor or pool to a domain

1. In the **Resources Pane**, select the pool or server for which you want to grant somebody permissions.
2. Select the **Users** tab.
3. Select **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, specify `other_domain\jsmith`. Alternatively, you can enter a fully qualified domain name (FQDN). For example, specify `jsmith@other_domain.com`.
3. Select **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, 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. Select **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**. If you didn't fill the username/password or know an account with sufficient privileges, select this option 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

November 16, 2023

All Citrix Hypervisor users must have an RBAC role. In Citrix Hypervisor, you must first assign a role to the newly created user before they can use the account. Citrix Hypervisor does not automatically assign a role to the newly created user. As a result, these accounts do not have any access to the Citrix Hypervisor 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 the Citrix Hypervisor pool. Add the AD account after joining the associated domain. For more information, see 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 action requires the **Assign/modify role** permission, which is only available to a Pool Administrator.
2. Modify the user's group membership in your Active Directory to make the user part of a group that is assigned a different role.

If an administrator indirectly applies multiple roles to a user, Citrix Hypervisor grants the user the permissions from the highest role that the user is assigned to.

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. Select 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. Select **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 logged on to the pool and you want them to receive their new permissions immediately, click **Logout User**. This action disconnects the user's sessions on the pool so the user receives a new session with the modified role.

   > **Note:**
   >
   > When changing a role, the user must log out and log back in again for the new role to take effect. Force this log out by clicking the **Logout User** button.
   > To force a logout, the user requires the **Logout active user connections** permission. This permission is available to a Pool Administrator or Pool Operator.

> **Note:**
>
> If you remove the Pool Admin role from a user, consider also changing the server root password and rotating the pool secret. For more information, see Pool Security.

Calculating RBAC roles

May 25, 2023

When I log in, how does Citrix Hypervisor 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. Citrix Hypervisor then verifies the following information:
   - The roles assigned to the subject
   - The roles assigned to any Active Directory groups that the subject is a member of.
3. Citrix Hypervisor applies the highest level of permissions to the subject. Because subjects can be members of multiple Active Directory groups, they inherit all permissions of the associated roles.

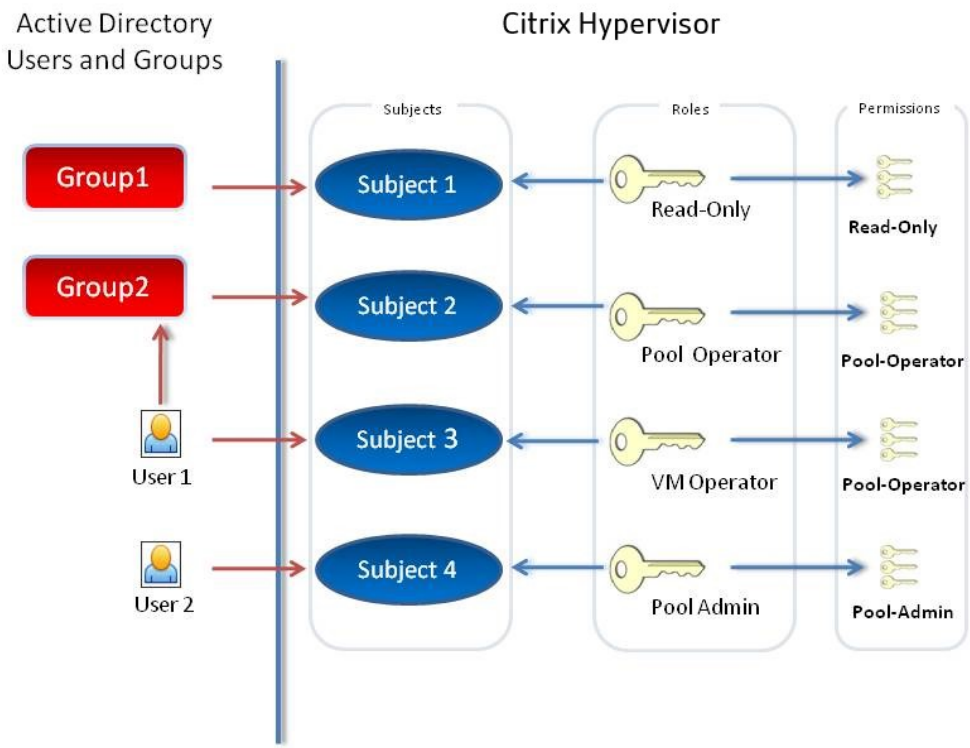

This illustration shows the following information:

- Subject 2 (Group 2) is the Pool Operator.
- User 1 is a member of Group 2.
- When Subject 3 (User 1) tries to log in, they inherit both Subject 3 (VM Operator) and Group 2 (Pool Operator) roles.
- The Pool Operator role is higher, so the resulting role for Subject 3 (User 1) is Pool Operator and not VM Operator.

Audit changes

May 25, 2023

The Citrix Hypervisor audit log, which is enabled by default, records any operation with side-effects performed by a known user. The operation is recorded whether it is successful or unsuccessful. This audit log includes:

- The user's name who performed the action. If the user's name is not available, Citrix Hypervisor logs the user ID instead.
- The server name that the action targeted.
- 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 Citrix Hypervisor `syslog` command to duplicate the audit log to a safe box. The syslog command is available from the xe CLI and documented in Command line interface.

If you are concerned with auditing, we recommend you implement Role Based Access Control. However, the audit log does not require that users are assigned RBAC roles nor does it require Active Directory integration.

Citrix Hypervisor 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

May 25, 2023

Workload Balancing is an 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 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 Citrix Hypervisor Premium Edition customers, or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about licensing, see About Citrix Hypervisor 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 can send notifications in XenCenter regarding the actions it takes. For more information on how to configure the alert level for Workload Balancing alerts by using the XenAPI, see Configuring Workload Balancing alerts in XenCenter.

Workload Balancing evaluates the utilization of VM workloads across a pool. When a host reaches one of its thresholds, WLB relocates the VM to a different host in the pool.

To ensure the rebalancing and placement recommendations align with your environment's needs, you can configure WLB to optimize workloads for resource performance or to maximize the density. These optimization modes can be configured to change automatically at predefined times or stay the same always. For more 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 VM performance, and VM motion history.

For more information about Workload Balancing, you can see the Citrix Hypervisor product documentation.

Getting Started with Workload Balancing

May 25, 2023

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, select **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.

For more information, see the Citrix Hypervisor product documentation - Get started with Workload Balancing.

> **Note:**
>
> The **WLB** tab appears on the **Properties** pane after you license your hosts with Citrix Hypervisor Premium Edition or a Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about licensing, see About Citrix Hypervisor Licensing.

Workload Balancing Basic Concepts

May 25, 2023

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 VMs on as few hosts as possible, while ensuring they maintain adequate computing power. This behavior enables you to 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 do not conflict with High Availability settings. High Availability settings always take precedence.

Connecting to Workload Balancing

May 25, 2023

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 CLI.

> **Note:**
>
> The **WLB** tab appears on the **Properties** pane after you license your hosts with Citrix Hypervisor Premium Edition or a Citrix Virtual Apps and Desktops license. For more information about licensing, see About Citrix Hypervisor 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 account is often known as the Workload Balancing user account. Citrix Hypervisor 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, 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:**
>
> If you don't receive optimal placement recommendations after WLB runs for a time, Citrix recommends you evaluate your performance thresholds. For more information, see 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, select the **WLB** tab. The WLB tab appears in the Properties pane after licensing your Citrix Hypervisor hosts.
3. In the **WLB** tab, select **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. Citrix Hypervisor uses this port to communicate with Workload Balancing.

      By default, Citrix Hypervisor 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:**
      >
      > Change the default port number only if 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 that Citrix Hypervisor uses to connect to the Workload Balancing virtual appliance. This account must be the account you created during Workload Balancing Configuration. By default, the user name for this account is `wlbuser`.
6. In the **Citrix Hypervisor Credentials** section, enter the user name and password for the pool you are configuring. Workload Balancing uses these credentials to connect to each of the hosts in that pool. To use the credentials with which you are currently logged into Citrix Hypervisor, 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

May 25, 2023

Workload Balancing is a powerful Citrix Hypervisor component that includes many features designed to optimize the workloads in your environment. These features include:

- Host power management
- The ability to schedule optimization-mode changes

- 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 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

Determining the best host on which to start a VM

Workload Balancing can provide recommendations about the host. Determining the host to start a VM on is useful when you want to restart a powered off VM and when you want to migrate a VM to another host. It might also be useful in Citrix Virtual Desktops environments.

For more information, see Choosing an Optimal Server for VM Initial Placement, Migrate, and Resume.

Accepting Workload Balancing 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 issues a recommendation to consolidate VMs on a host. Assuming you are not running in automated mode, you can choose to either apply this recommendation or simply ignore it.

For more information, see Accepting Optimization Recommendations.

Both of these tasks, and how 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 some time, if you don't receive optimal placement recommendations, evaluate your performance thresholds as described in the Workload Balancing documentation. 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

May 25, 2023

When Workload Balancing is enabled and you start a VM, XenCenter provides recommendations to help you determine the optimal host in the pool to start a VM on. The recommendations are also known as star ratings since stars are used to indicate the best host.

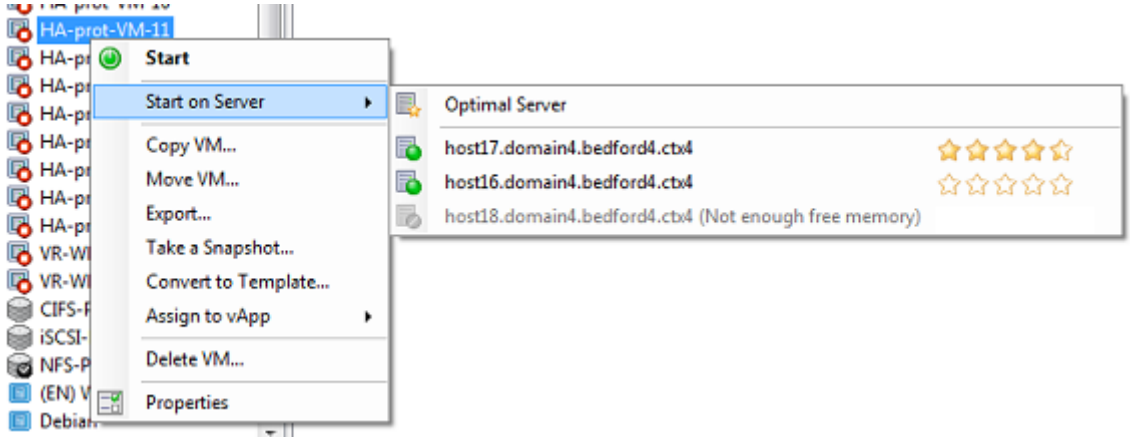

More stars appear beside host17 since this server is the optimal host on which to start the VM. host16 does not have any stars beside it, which indicates that the host is not recommended. However, since this host is enabled the user can select that host. host18 is grayed out due to insufficient memory, 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 VM. These ratings also apply when you want to start a VM that is powered off or suspended and when you want to migrate the VM to another server.

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 indicate the lowest-rated (least optimal) server. When it is not possible to start or move a VM to a host, the host name is grayed out. The reason that the host cannot accept the VM appears beside it.

In general, Workload Balancing functions more effectively and makes better, less frequent optimization recommendations if you start VMs 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 runs in Maximum Performance mode, Workload Balancing tries to balance the VMs across the hosts in the pool so that all VMs have good performance. When a pool is running in Maximum Density mode, Workload Balancing tries to place VMs onto hosts as densely as possible while ensuring the VMs 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 VMs on hosts with the most resources available of the type the VM requires. In Maximum Density mode, Workload Balancing tries to place VMs on hosts that already have VMs running so that VMs are running on as few hosts as possible.
- The amount and type of resources the VM requires. After WLB monitors a VM for a while, it uses the VM metrics it gathered to make placement recommendations according to the type of resources the VM requires. For example, WLB might select a host with less available CPU but more available memory if it is what the VM requires (based on its past performance history). However, 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 indicate 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 indicate the most-recommended (optimal) server and five empty stars indicates the least-recommended server.

Accepting Optimization Recommendations

May 25, 2023

Workload Balancing provides recommendations about ways you can migrate virtual machines to optimize your environment. Optimization recommendations appear in the **WLB** tab in XenCenter.

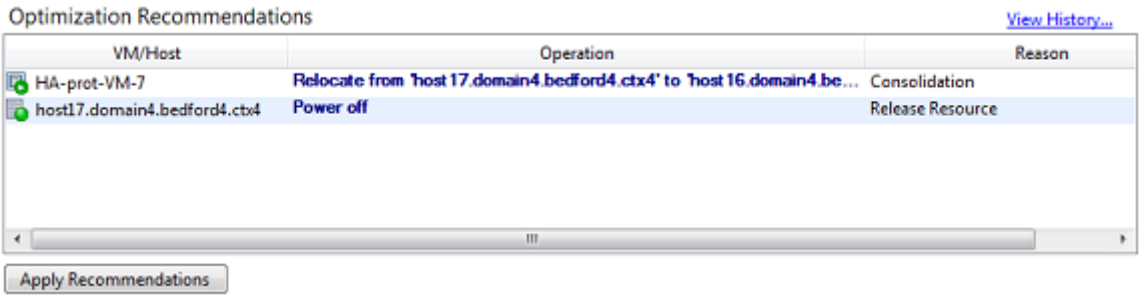

This illustration shows a screen capture of the Optimization Recommendations list, which appears on the **WLB** tab. The **Reason** column displays the purpose of the recommendation. The **Operation** column displays the behavior change suggested for that optimization recommendation. This screen capture 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 following factors:

- 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 VMs.)

The optimization recommendations display the following information:

- The name of the VM that Workload Balancing recommends relocating
- The host it currently resides on
- The host Workload Balancing recommends as the machine's new location
- The reason Workload Balancing recommends moving the VM

  For example, "CPU"to improve the CPU utilization.

After you accept an optimization recommendation, Citrix Hypervisor 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 select 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, select **Apply Recommendations**. Citrix Hypervisor begins moving all virtual machines listed in the **Optimization Recommendations** section to their recommended servers.

   After you select **Apply Recommendations**, you can select **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 Citrix Hypervisor 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 might still make recommendations to power off more hosts than the number of host failures to tolerate. (For example, Workload Balancing still makes 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 might display an error message stating that High Availability is no longer 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 see a message that says a power-management recommendation was not applied because High Availability is enabled.

Working with Workload Balancing Reports

May 25, 2023

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

May 25, 2023

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. Workload 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, the Workload Balancing recommendations might not be appropriate for your environment.

Troubleshooting administrative changes

You can use the Pool Audit Trail report to determine the source (that is, user account) of problematic changes and the event or task that user performed.

Generating and Managing Workload Balancing Reports

May 25, 2023

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 have 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 the toolbar icon.

Report Toolbar Buttons:

	**Document Map**. Displays a document map that helps you navigate through
	**Page Forward/Back**. Moves one page ahead or back in the report.
	**Back to Parent Report**. Returns to the parent report when working with dri
	**Stop Rendering**. Stops the report generation.
	**Print**. Prints a report and specify general printing options, such as the printe
	**Print Layout**. Displays a preview of the report before you print it. To exit Pri
	**Page Setup**. Specifies printing options such as the paper size, page orientat
	**Export**. Exports the report as an Acrobat (.pdf) file or as an Excel file with a
	**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, select **Print Layout**.

   

2. (Optional.) To change the paper size/source, page orientation, or margins, select **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, select **Export** and select one of the following:

- Excel
- Acrobat (`.pdf`) file

> **Note:**
>
> The amount of data included when you export a report might differ depending on the export format. Reports exported to Excel include all the data available for the report, including "drill-down"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

May 25, 2023

Workload Balancing provides the following 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 department within your organization used. Specifically, the report shows information about all the virtual machines in your pool, including their availability 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. You can 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 want to bill internal or external customers for virtual machine usage, consider incorporating department or customer names in your VM naming conventions. This convention makes reading chargeback reports easier.

The resource reporting in the chargeback report is, sometimes, based on the allocation of physical resources to individual VMs.

Likewise, because Citrix Hypervisor 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 VM, whether it is through a fixed memory allocation or an automatically adjusting memory allocation (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 CPUs 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 the virtual CPU utilization in the reporting period. This value is expressed as a percentage of the VM's virtual CPU capacity. The capacity is based on the number of virtual CPUs allocated to the VM. For example, if you allocated one virtual CPU to the VM, Minimum CPU Usage represents the lowest percentage of virtual CPU usage Citrix Hypervisor recorded. If you allocated two virtual CPUs to the VM, the value in this column represents the lowest usage of the combined capacity of both virtual CPUs. The value is expressed as a percentage.

  Ultimately, the percentage of CPU usage represents the lowest recorded workload that the virtual CPU handled. For example, if you allocate one virtual CPU to a VM and the physical CPU on the host is 2.4 GHz, you are allocating one-eighth of 2.4 GHz to the VM. This behavior means that if the VM's allocated capacity is 0.3 GHz, or 300 MHz, and the Minimum CPU Usage for the virtual machine was 20%, the VM's lowest usage of the physical host's CPU during the reporting period was 60 MHz.
- **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 you allocated one virtual CPU to the VM, 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, 1,024 MB), 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 Citrix Hypervisor to adjust a VM's memory automatically based on a range, this column shows the minimum amount of memory specified in the range. 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 1,024 MB as the minimum memory and 2,048 MB as the maximum memory, then 1,024 MB appears in the **Current Minimum Dynamic Memory (MB)** column.
- **Current Maximum Dynamic Memory (MB)**.
  - **Dynamic memory allocation**. If Citrix Hypervisor is set to adjust a VM's memory automatically based on a range, this column shows the maximum amount of memory specified in the range. For example, if the memory range you provided was 1,024 MB minimum and 2,048 MB maximum, then 2,048 MB appears in the Current Maximum Dynamic Memory (MB) column.
  - **Fixed memory allocation**. If you assign a VM a fixed amount of memory (for example, 1,024 MB), 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 Citrix Hypervisor is allocating to the virtual machine when the report is run.
  - **Fixed memory allocation**. If you assign a virtual machine a fixed amount of memory (for example, 1,024 MB), 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 Citrix Hypervisor 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, 1,024 MB), 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 might 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 a 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 the 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 15 days at 12:00 hours combined as an average. That is, if CPU utilization was 82% at 12 PM on May 1, 88% at 12 PM on May 2, 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 declines or remains steady after an optimization event. If you do not see improved resource usage after optimization, consider readjusting threshold values. Also, consider whether the resource pool has too many virtual machines and whether new virtual machines were added or removed during the time frame you specified.

Pool audit trail

This report displays the contents of Audit Log, a feature designed to log attempts to perform unauthorized actions and select authorized actions. These actions include import/export, host and pool

backups, and guest and host console access. The report gives more meaningful information when Citrix Hypervisor 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 documentation.

> **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, select the **WLB** tab, and then select **Settings**.
2. In the left pane, select **Advanced**.
3. On the Advanced page, select the **Pool Audit Trail Report Granularity** list and select an option from the list.

   > **Important:**
   >
   > Select the granularity based on 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 contains detailed information about the audit log report. 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 Citrix Hypervisor recorded the user's action.
- **User Name**. The name of the person who created the session in which the action was performed. Sometimes, this value might 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 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 value provides the status of the action (that is, whether 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 the 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 so, adjust your performance threshold values until most 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 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 most 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 the 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.

Except for the Average Free Memory graph, the data points never average above the Critical threshold line (red). For the Average Free Memory graph, the data points 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 so, adjust the pool's Critical threshold values until most 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 intersection 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 most 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 the 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**: Whether 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. 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 VMs 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 VM is 67%, your VM was using, on average, 67% of its vCPU for the specified period.

The initial view of the report displays an average value for the 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

May 25, 2023

The Audit Log report logs Citrix Hypervisor events, event objects, and actions. These actions include 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 Citrix Hypervisor Audit Log and Pool Audit Trail report. The table 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 extra event definitions, see the Management API. 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 Citrix Hypervisor 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 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.

Pool Audit Trail Granularity	Event Action	User Action
Medium	VIF.plug	Hot plugged 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 is canceled.
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. This record is initialized when the VM boots and updated to reflect runtime configuration changes, for example, CPU hotplug.

Editing Workload Balancing Settings

May 25, 2023

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
- Configuring Workload Balancing alerts in XenCenter

> **Note:**
>
> After connecting or reconnecting to Workload Balancing, wait at least 60 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, 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 so, try adjusting the thresholds until Workload Balancing begins providing suitable optimization recommendations.
- Not getting recommendations when you think your network has insufficient bandwidth. If so, 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 CLI to modify the configuration settings.

To update the credentials that Citrix Hypervisor 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 documentation.

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

May 25, 2023

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.

You can select one of two optimization modes:

- **Maximize Performance**. (Default.) Workload Balancing attempts to spread the 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 an 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 an optimization when a virtual machine reaches the Critical threshold.

Workload Balancing also lets you apply these optimization modes always, *Fixed*, or switch between modes for specified time periods, *Scheduled*.

- **Fixed**. Fixed optimization modes always set Workload Balancing to a specific optimization behavior - either to try to create the best performance or the highest density.

- **Scheduled**. Scheduled optimization modes let you schedule for Workload Balancing to apply different optimization modes depending on the time of day. For example, you can configure Workload Balancing to optimize for performance during the day when you have users connected. Then you can save energy by specifying 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 that you specified.

To set an optimization mode for all time periods

1. Select the pool in the **Resources** pane, select the **WLB** tab, and then select **Settings**.
2. In the left pane, select **Optimization Mode**.
3. Select **Fixed**, and select one of these optimization modes:
   - **Maximize Performance**. (Default.) Attempts to spread the 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. (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 changes automatically

1. Select the pool in the **Resources** pane, select the **WLB** tab, and then select **Settings**.
2. In the left pane, select **Optimization Mode**.
3. Select **Scheduled**.
4. Select **Add New** to open the **Optimization Mode Scheduler** dialog box.
5. Select an optimization mode in the **Change to** list box:
   - **Maximize Performance**. Attempts to spread the 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.
6. Select the day of the week and the time when you want Workload Balancing to begin operating in this mode.
7. Create more scheduled mode changes (that is, "tasks") until you have the number you need. If you only schedule one task, Workload Balancing switches to that mode and never switch back.
8. Select **OK**.

To delete or pause a scheduled optimization mode task

1. Select the pool in the **Resources** pane, select the **WLB** tab, and then select **Settings**.
2. Select **Optimization Mode**.
3. Select a scheduled task and select one of the following:

   - **Delete the task permanently**. Select the **Delete** button.
   - **Stop the task from running temporarily**. Right-click the task and select **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, select the **WLB** tab, and then select **Settings**.
2. Select a scheduled task.
3. Select **Edit**.
4. In the **Change to** box, select a different mode or make other changes as desired.

Optimizing and Managing Power Automatically

May 25, 2023

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. 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 does not automatically apply recommendations to hosts or virtual machines when the recommendations conflict with High Availability settings. If a pool becomes overcommitted

by applying Workload Balancing optimization recommendations, XenCenter prompts you whether you want to continue applying the recommendation. When Automation is enabled, Workload Balancing does 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

Power management is the ability to 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 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 WLB detects the pool has insufficient host capacity to shut down servers, it recommends leaving the servers on until the pool workload decreases enough. When you configure Workload Balancing to power off extra servers automatically, it applies these recommendations automatically and, therefore, 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.

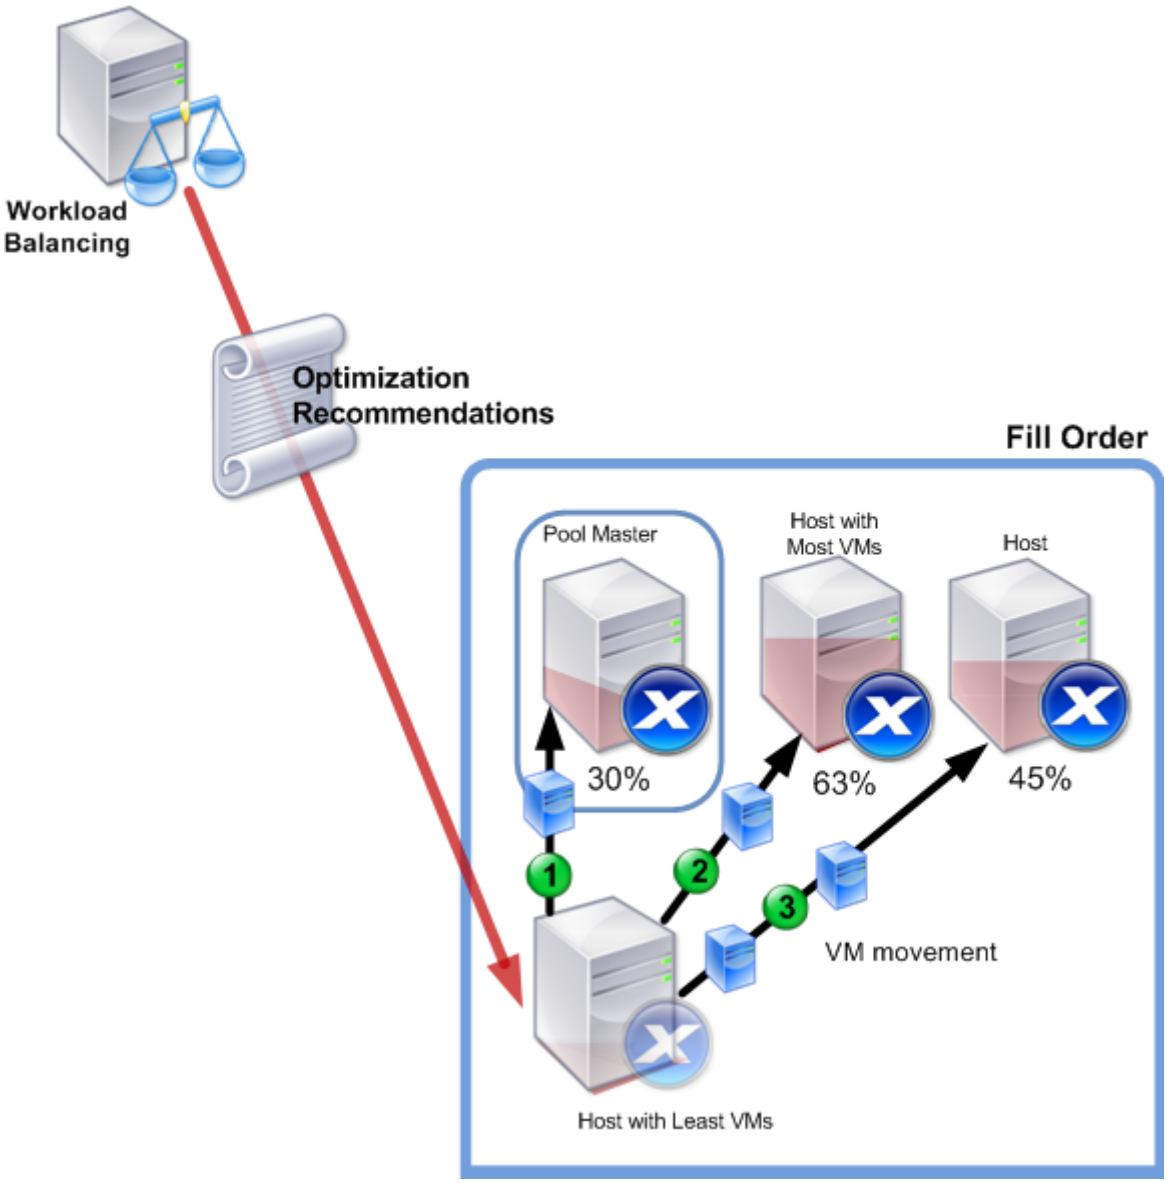

This illustration shows how, when consolidating VMs on hosts in Maximum Density mode, Citrix Hypervisor fills the pool master first, the most loaded server second, and the least loaded server third.

If Workload Balancing detects a performance issue while the pool is in Maximum Density mode, it ad-

dresses 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 hosts to power on by applying the same criteria it would if the optimization mode was set to Maximum Performance.)

When WLB runs in Maximum Performance mode, it powers on hosts until resource utilization on all hosts in the pool falls below the High threshold.

While migrating one or more VMs, if WLB determines that increasing capacity would benefit the pool's overall performance, it powers on hosts automatically or recommends doing so.

> **Important:**
>
> Workload Balancing only recommends powering on a host that Workload Balancing powered off.

Designing environments for Power Management and VM consolidation

When you are planning Citrix Hypervisor implementations and you intend to configure automatic VM consolidation and power management, consider your workload design. For example, you might want to:

- **Place Different Types of Workloads in Separate Pools.** If you have distinct types of workloads or types of applications that perform better with certain types of hardware, consider whether to locate the VMs hosting these workloads in different pools.

  Because power management and VM consolidation are managed at the pool level, design pools so they contain workloads that you want consolidated at the same rate. Factor in considerations such as those discussed in the Advanced Settings topic.

- **Exclude Hosts from Workload Balancing.** Some hosts might have to be always on. For more information, see Excluding Hosts from Recommendations.

To apply optimization recommendations automatically

1. In the **Resources** pane of XenCenter, select **XenCenter** > your resource pool.
2. In the **Properties** pane, select the **WLB** tab.
3. In the **WLB** tab, select **Configure WLB**.
4. In the left pane, select **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 the 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 the 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 off hosts automatically during low usage periods.

6. Do one of the following:
   - If you want to configure power management, select **Automation/Power Management** and proceed to the following section.
   - If you do not want to configure power management and you are finished configuring automation, select **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 **Automatically apply Power Management recommendations** 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

May 25, 2023

This topic provides guidance about how to modify the default Critical thresholds and how to set values for the 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 documentation.

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.

Workload Balancing determines whether to produce recommendations based on if the averaged historical utilization for a resource violates its threshold. **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.)

For more information, see Workload Balancing documentation.

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	51 MB	63.75 MB	510 GB	1,020 GB
Network Read	25 MB/s	21.25 MB/s	12.5 MB/s	6.25 MB/s
Network Write	25 MB/s	21.25 MB/s	12.5 MB/s	6.25 MB/s

Setting	Default	High	Medium	Low
Disk Read	25 MB/s	21.25 MB/s	12.5 MB/s	6.25 MB/s
Disk Write	26 MB/s	21.25 MB/s	12.5 MB/s	6.25 MB/s

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 behavior means that if you increase, for example, a Critical threshold to 95%, WLB automatically resets the other thresholds to the following:

- High threshold to 80.75%
- Medium threshold to 47.5%
- Low threshold to 23.75%

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:

```
1 High, Medium, or Low Threshold = Critical Threshold \* Threshold Factor
```

For example, if you change the Critical threshold for Network Reads to 40 MB/s and you want to know its other thresholds:

- To obtain the Low threshold, multiply 40 by 0.25
- To obtain the Medium threshold, multiply 40 by 0.50
- To obtain the High threshold, multiply 40 by 0.85

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 a lower resource utilization.
- **Maximum Density**. Workload Balancing doesn't recommend placing a VM on host if doing so causes the utilization of any host resource 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. This signal 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 resources reaches its High threshold.

  However, if after a VM is relocated, utilization of a resource on the new host exceeds its Critical threshold, WLB temporarily uses an algorithm similar to the Maximum Performance algorithm to find a new host for the VMs. 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, select the **WLB** tab, and then select **Settings**.
2. In the left pane, select **Critical Thresholds**.
3. In the **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

May 25, 2023

Workload Balancing uses metric weightings, a method of assigning importance to resources, to determine what hosts to optimize first.

> **Note:**
>
> Before tuning metric weightings, Citrix recommends reading about the optimization and consolidation process in the Workload Balancing documentation. The information in this article is a subset of that information and is only meant to be used as a reference when changing the user interface.

When **Workload Balancing** is processing optimization recommendations, it creates an optimization order. Workload Balancing determines the optimization order by ranking 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.

How 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. In general, metric weightings are 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:

- Which hosts' performance to address first
- Which VMs to recommend migrating first

For example, if you rank Network Writes as the most important resource, Workload Balancing first addresses performance issues and makes optimization recommendations for the host with the most Network Writes per second.

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 resources 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. 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. The metric weighting for memory utilization is set to the far left of the slider **(Less Important)**.

Workload Balancing recommends 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. These recommendations begin 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 more 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 behavior means that if all metrics are set to the same level, even if that level is **Less Important**, they are all 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, select the **WLB** tab, and then select **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

May 25, 2023

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 might 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 VM workloads must always 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, excluded hosts remain excluded even when the optimization mode changes. To prevent Workload Balancing from powering down a host automatically, consider not enabling Power Management for that host. For more information, see Optimizing and Managing Power Automatically.

To exclude hosts from placement and optimization recommendations

1. Select the pool in the **Resources** pane, select the **WLB** tab, and select **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

May 25, 2023

The settings in the **Advanced** dialog primarily fine-tune how Workload Balancing applies recommendations when it is running in automated mode.

> **Important:**
>
> After Workload Balancing is running for a period, if you do not receive optimal placement recommendations, evaluate your performance thresholds as described in the Workload Balancing documentation. It is critical to set Workload Balancing to the correct thresholds for your environment or its recommendations might not be appropriate.

When WLB runs in automated mode, the frequency of optimization and consolidation recommendations and how soon they are automatically applied is a product of multiple factors. These factors include:

- **VM migration interval**: How long you specify Workload Balancing waits before applying another optimization recommendation.
- **Recommendation count**: The number of recommendations Workload Balancing must make before applying a recommendation automatically
- **Recommendation severity**: The severity level a recommendation must achieve before the optimization is applied automatically
- **Optimization aggressiveness**: 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 how many minutes WLB waits after the last time a particular VM was moved, before it generates another optimization recommendation that includes that particular VM.

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, therefore, 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 perform the following checks:

- Vet if the recommendation is truly needed
- Determine if the destination host has stable enough performance over a prolonged period to accept a relocated VM (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 can reset the Recommendation Count to zero, depending on the factors described in the Workload Balancing documentation
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 to set varies according to your environment. Consider these scenarios:

- If server loads and activity increase quickly in your environment, you can increase the 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 can 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. This setting encourages Workload Balancing to apply more conservative, stable recommendations and can decrease the potential for spurious virtual machine moves. However, the Recommendation Count is set to a conservative value by default.

Because of the potential impact adjusting this setting might have on your environment, change it with extreme caution. We advise that you test and iteratively change the value, or change it 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 the configuration options you set, such as:

- Thresholds and metric tunings
- Resources available for the workload
- Resource-usage history.

When you configure Workload Balancing to apply optimization recommendations automatically, you can set the minimum severity level to associate with a recommendation before Workload Balancing automatically applies it.

Optimization aggressiveness

To provide more assurance when running in automated mode, Workload Balancing has consistency criteria for accepting optimizations automatically to prevent moving VMs 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 VM 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 VMs on that host continue to trigger recommendations.

Workload Balancing determines if a 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 an **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 the 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 to 3, Workload Balancing waits until it sees all the criteria for Low are met and repeated in three consecutive recommendations. This behavior helps ensure that the VM actually needs to be moved and that the destination host Workload Balancing recommends has consistently stable resource utilization over a longer period. 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.

For information about the criteria for the Low aggressiveness level, see Workload Balancing documentation.

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 documentation.

If you find that Workload Balancing is not automatically applying optimization recommendations frequently enough, you might want to increase the aggressiveness setting. However, Citrix strongly recommends reviewing the information in the Workload Balancing documentation before doing so.

To configure virtual machine recommendation intervals

1. Select the pool in the **Infrastructure** view, select the **WLB** tab, and then select **Settings**.
2. In the left pane, select **Advanced**.
3. In the **VM Migration Interval** section, do one or more of the following:

- In **Minutes to wait**, enter 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 can take one of the following actions:
>
> - Change the granularity setting to Medium or Minimum.
> - Consider expanding your WLB server's memory.
> - Consider expanding the size of the hard disk.

For more information, see Workload Balancing Report Glossary and Audit Log Events.

Administering Workload Balancing

May 25, 2023

Some administrative tasks you might 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 Citrix Hypervisor 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.

> **Note:**
>
> For information about configuring Workload Balancing to use a different certificate or configuring Citrix Hypervisor to verify the identity of a certificate, see the Workload Balancing documentation.

Disconnecting from Workload Balancing

May 25, 2023

If you want to stop Workload Balancing (WLB) 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 WLB virtual appliance, Workload Balancing permanently deletes pool information from the WLB database and stops collecting data for that pool. If you want to use the same WLB virtual appliance to manage the pool again, reenter the appliance's information in the **Connect to WLB Server** dialog box.

> **Important:**
>
> If you only want to stop Workload Balancing temporarily, select the **WLB** tab and select the **Pause** button.

To disconnect from Workload Balancing

1. In the **Resource** pane of XenCenter, choose 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 server** 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

May 25, 2023

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, 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 select **Disconnect** when prompted. For instructions, see Disconnecting from Workload Balancing.
2. In the **WLB** tab, select **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:**
   >
   > 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

May 25, 2023

After initial configuration, to update the credentials the Citrix Hypervisor server and the Workload Balancing appliance use to communicate, complete the following process:

1. Disconnect from Workload Balancing, as described in a following section.
2. Change the WLB credentials by editing the WlbConfig file (run the `WlbConfig` command in the console on the Workload Balancing virtual appliance). For more information, see the Workload Balancing documentation.
3. Re-enable Workload Balancing and specify the new credentials, as described in a following section.

Situations when you might want to use these steps include:

- If you have to change the user account Citrix Hypervisor 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 server** dialog box appears.
3. Select **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, select **Connect**. The Connect to WLB Server dialog box appears.
2. Select **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 Citrix Hypervisor uses to connect to Workload Balancing.

      By default, Citrix Hypervisor connects to Workload Balancing on port 8012.

> **Note:**
>
> Only edit this port number if 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 systems running Citrix Hypervisor use to connect to the Workload Balancing server.
5. In the **Citrix Hypervisor Credentials** section, enter the user name and password for the pool you are configuring (typically the password for the pool master). Workload Balancing uses these credentials to connect to the computers running Citrix Hypervisor in that pool. To use the credentials with which you are currently logged into Citrix Hypervisor, select the **Use the current XenCenter credentials** check box.

Entering maintenance mode with Workload Balancing Enabled

May 25, 2023

When WLB is enabled, if you take a host offline for maintenance, Citrix Hypervisor automatically migrates the VMs running on that host to their optimal servers when available. Citrix Hypervisor 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.

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 select **Enter Maintenance Mode** on the shortcut menu.
   - On the **Server** menu, select **Enter Maintenance Mode**.

2. Select **Enter Maintenance Mode**. The VMs 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 select **Exit Maintenance Mode** on the shortcut menu.
   - On the **Server** menu, select **Exit Maintenance Mode**.
2. Select **Exit Maintenance Mode**.

   When you remove a server from maintenance mode, Citrix Hypervisor automatically restores that server's original virtual machines to that server.

Troubleshooting Workload Balancing

May 25, 2023

While Workload Balancing usually runs smoothly, this help system provides a series of topics with guidance in case you encounter issues. Extra troubleshooting topics are provided in the Workload Balancing documentation.

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 the Workload Balancing virtual appliance in this location (by default):

```
/var/log/wlb
<!--NeedCopy-->
```

Also, you can also see the event logs in the XenCenter **Navigation** pane, select **Notifications** and then **Events** for more information.

Error messages

Workload Balancing displays error messages in the **Alerts** view in XenCenter and, sometimes, on the screen as dialog boxes.

Issues Entering Workload Balancing Credentials

May 25, 2023

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:

- Ensure the Workload Balancing virtual appliance is imported and was configured correctly and all of its services are running by using the following command: `service workloadbalancing start`
- 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 fully qualified domain name.

You can enter the host name of the Workload Balancing appliance in the **Address** box, but it must be a fully qualified domain name. For example, `yourcomputername.yourdomain.net`.

Issues Starting Workload Balancing

May 25, 2023

After importing and configuring the Workload Balancing appliance, you might receive an error message that Citrix Hypervisor and Workload Balancing cannot connect to each other. If so, 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 the Citrix Hypervisor server
- 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

May 25, 2023

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.

For more information, see CTX231579 - Troubleshooting Workload Balancing (WLB) issues when connecting via XenCenter.

Issues changing Workload Balancing servers

May 25, 2023

If you connect a pool to a different Workload Balancing virtual appliance without first disconnecting the pool from the original Workload Balancing appliance, both appliances monitor the pool.

To solve this problem, you can do one of the following actions:

- Shut down and delete the old Workload Balancing appliance
- 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 the Workload Balancing server configuration.

XenServer Conversion Manager

March 6, 2024

Use the XenServer Conversion Manager (formerly Citrix Hypervisor Conversion Manager) virtual appliance to migrate your entire VMware environment to XenServer quickly and efficiently. You can convert up to 10 VMware ESXi/vCenter VMs in parallel at the same time. After converting your VMs, the Conversion Manager automatically shuts down by itself, saving resources on the host.

As part of the migration, XenCenter helps you prepare the VMs for networking and storage connectivity. After converting your VMs to a XenServer environment, they're almost ready to run.

> **Note:**
>
> In Citrix Hypervisor 8.0 and earlier, a separate Conversion Manager console is provided. From

> Citrix Hypervisor 8.1, this capability is integrated into XenCenter.

Overview

Citrix Hypervisor allows you to:

- Convert up to 10 VMware ESXi/vCenter VMs in parallel using one simple wizard
- Map network settings between VMware and Citrix Hypervisor so your converted VMs can be up and running with the proper network settings
- Select a storage location where you would like your new Citrix Hypervisor VMs to run

> **Notes:**
>
> - XenCenter does not remove or change your existing VMware environment. VMs are duplicated onto your Citrix Hypervisor environment and not removed from VMware.
> - XenServer Conversion Manager virtual appliance supports converting VMware ESXi/vCenter VMs with different storage such as thin provisioning, thick provisioning, IDE, and SCSI.
> - XenServer Conversion Manager virtual appliance does not require the source VMs to have VMware Tools installed. You can perform conversion on VMware ESXi/vCenter VMs regardless of whether they have VMware Tools installed.
> - XenServer Conversion Manager virtual appliance cannot convert VMware ESXi/vCenter VMs with four or more disks into Citrix Hypervisor VMs. Your VMware ESXi/vCenter VMs must have three or fewer disks.
> - XenServer Conversion Manager virtual appliance is available for Citrix Hypervisor Premium Edition customers or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about Citrix Hypervisor licensing, see Licensing. To upgrade, or to buy a Citrix Hypervisor 8.2 license, visit the Citrix website.

Understand Citrix Hypervisor

Before you can convert your environment, it is suggested that you become familiar with Citrix Hypervisor concepts. For more information, see Technical overview.

To successfully convert VMware virtual machines to Citrix Hypervisor, perform the following tasks:

- Set up a basic Citrix Hypervisor environment, including installing Citrix Hypervisor. For more information, see Quick start and Install.

- Create a network in Citrix Hypervisor, assigning an IP address to a NIC. For more information, see Quick start.
- Connect to storage. For more information, see Quick start.

Compare VMware and Citrix Hypervisor terminology The following table lists the approximate Citrix Hypervisor equivalent for common VMware features, concepts, and components:

VMware Term	Citrix Hypervisor Equivalent
VMware vSphere Client	XenCenter (the management console for Citrix Hypervisor)
Cluster / Resource Pool	Resource Pool
Data Store	Storage Repository
vMotion	Live migration
Distributed Resource Scheduling (DRS)	Workload Balancing
High Availability (HA)	High Availability (HA)
vCenter Converter	XenServer Conversion Manager virtual appliance
Role Based Access Control (RBAC)	Role Based Access Control (RBAC)

Conversion overview

XenCenter and XenServer Conversion Manager virtual appliance create a copy of each targeted VM. After converting the targeted VM to a Citrix Hypervisor VM with comparable networking and storage connectivity, XenCenter imports the VM into your Citrix Hypervisor pool or host. You can convert as few as one or two VMs or perform batch conversions of an entire environment of up to 10 VMware ESXi/vCenter VMs in parallel at the same time.

> **Note:**
>
> Before converting the VMs from vSphere, you must shut down the VMs (intended for conversion) on vSphere. XenServer Conversion Manager virtual appliance does not support converting a running VM using memory copied from vSphere to Citrix Hypervisor.
>
> Also, before converting, ensure that a network and a storage controller exist in your VMware VM.

The conversion process requires four items:

- **XenCenter** - the Citrix Hypervisor management interface includes a conversion wizard where you set conversion options and control conversion. You can install XenCenter on your Windows

desktop. XenCenter must be able to connect to Citrix Hypervisor and the XenServer Conversion Manager virtual appliance.

- **XenServer Conversion Manager virtual appliance** - a pre-packaged VM you import into the Citrix Hypervisor host or pool where you want to run the converted VMs. The virtual appliance converts the copies of the VMware ESXi/vCenter VMs into Citrix Hypervisor virtual machine format. After conversion, it imports these copies into the Citrix Hypervisor pool or host.
- **Citrix Hypervisor standalone host or pool** - the Citrix Hypervisor environment where you want to run the converted VMs.
- **VMware server** - XenServer Conversion Manager requires a connection to a VMware server that manages the VMs you want to convert. This connection can be to a vCenter Server, ESXi Server, or ESX Server. The VMs are not removed from the VMware server. Instead, the XenServer Conversion Manager Virtual Appliance makes a copy of these VMs and converts them to Citrix Hypervisor virtual-machine format.

The following illustration shows the relationships between these components:

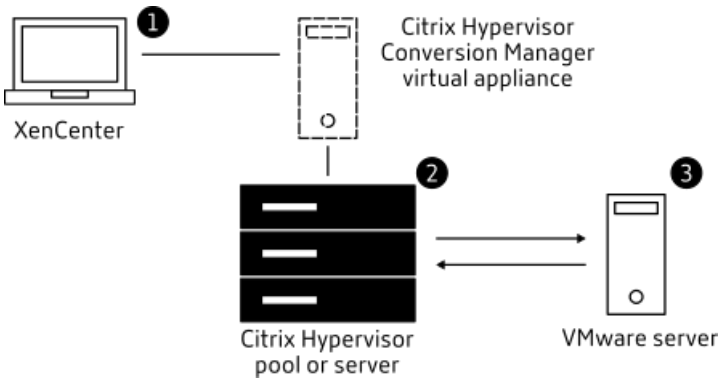

This illustration shows:

1. How XenCenter communicates with XenServer Conversion Manager virtual appliance.
2. How the XenServer Conversion Manager virtual appliance authenticates with the VMware server.
3. How the VMware server responds to the XenServer Conversion Manager virtual appliance during conversion.

The VMware server communicates with the XenServer Conversion Manager virtual appliance only when the appliance queries the VMware server for environment information and disk data throughout the conversion.

Summary of how to convert VMs You can configure the XenServer Conversion Manager virtual appliance and start to convert VMs in just a few easy steps:

1. Download the XenServer Conversion Manager virtual appliance from the Citrix Hypervisor 8.2 Premium Edition page.
2. Import the XenServer Conversion Manager virtual appliance into Citrix Hypervisor using XenCenter.
3. Configure the XenServer Conversion Manager virtual appliance by using XenCenter.
4. From XenCenter, launch the conversion wizard and start to convert VMs.
5. Complete the post-conversion tasks which include installing XenServer VM Tools (formerly Citrix VM Tools) for Windows on your Windows VMs. For Linux VMs, the XenServer Conversion Manager automatically installs XenServer VM Tools for Linux during the conversion process.

After converting your VMs, the Conversion Manager automatically shuts down by itself, saving resources on the host. For more information on how to convert VMware ESXi/vCenter VMs, see Get started with XenServer Conversion Manager.

What's new in XenServer Conversion Manager

March 6, 2024

The latest version of the XenServer Conversion Manager (formerly Citrix Hypervisor Conversion Manager) virtual appliance is version 8.3.1. You can download this version of the XenServer Conversion Manager virtual appliance from the Citrix Hypervisor Downloads page.

What's new in 8.3.1

Released Feb 01, 2024

This update includes the following improvement:

- You can now convert up to 10 VMware ESXi/vCenter VMs in parallel at the same time.

Get started with XenServer Conversion Manager

April 18, 2024

Using the XenServer Conversion Manager (formerly Citrix Hypervisor Conversion Manager), you can easily convert your VMware ESXi/vCenter virtual machines (VMs) to Citrix Hypervisor in just a few steps:

1. Prepare your Citrix Hypervisor environment and review the prerequisite information.
2. Import and configure the XenServer Conversion Manager virtual appliance by using XenCenter.
3. From XenCenter, launch the conversion wizard and begin converting your VMware ESXi/vCenter VMs to Citrix Hypervisor.
4. Complete the post-conversion tasks.
5. Review other conversion tasks.

Prepare your environment

Before converting your VMware environment, you must create and prepare the target Citrix Hypervisor standalone host or pool to run the converted VMware ESXi/vCenter VMs. Preparing your environment includes the following activities:

1. Defining a strategy of how you convert your VMware environment. Do you want to convert 1 or 2 VMs? Do you want to convert your entire environment? Do you want to create a pilot first to ensure that your configuration is correct? Do you run both environments in parallel? Do you want to maintain your existing cluster design when you convert to Citrix Hypervisor?
2. Planning your networking configuration. Do you want to connect to the same physical networks? Do you want to simplify or change your networking configuration?
3. Installing Citrix Hypervisor on the hosts you want in the pool. Ideally, plug the NICs on the hosts into their physical networks before you begin installation.
4. Creating a pool and performing any basic networking configuration. For example, do the following:
    - Configure a network to connect to the VMware cluster on the Citrix Hypervisor host (if the cluster is not on the same network as the Citrix Hypervisor host).
    - Configure a network to connect to the storage array. That is, if you use IP-based storage, create a Citrix Hypervisor network that connects to the physical network of the storage array.
    - Create a pool and add hosts to this pool.
5. (For shared storage and Citrix Hypervisor pools.) Preparing the shared storage where you store the virtual disks and creating a connection to the storage, known as a Storage Repository (SR) on the pool.
6. (Optional.) Although not a requirement for conversion, you might want to configure the administrator accounts on the Citrix Hypervisor pool to match those accounts on the VMware server. For information about configuring Role-based Access Control for Active Directory accounts, see Role-based access control.

Install Citrix Hypervisor and create a pool

Before you can convert VMware ESXi/vCenter VMs, ensure that you create a Citrix Hypervisor pool or host where you want to run the converted VMs. This pool must have networking configured so it can connect to the VMware server. You might also want to configure the same physical networks on the Citrix Hypervisor pool that you have in the VMware cluster, or simplify your networking configuration. If you want to run the converted VMs in a pool, create a storage repository before conversion and add the shared storage to the pool.

If you are new to Citrix Hypervisor, you can learn about Citrix Hypervisor basics, including basic installation and configuration, by reading Quick start.

Citrix Hypervisor environment considerations

Before installing Citrix Hypervisor and importing the virtual appliance, consider the following factors that might change your conversion strategy:

Selecting the host where you want to run the XenServer Conversion Manager virtual appliance. Import the virtual appliance into the stand-alone host or into a host in the pool where you run the converted VMs.

For pools, you can run the virtual appliance on any host in the pool, provided its storage meets the storage requirements.

> **Note:**
>
> We recommend that you run only one XenServer Conversion Manager in a pool at a time.

The storage configured for the pool or host where you want to run the converted VMs must meet specific requirements. If you want to run your newly converted VMs in a pool, their virtual disks must be stored on shared storage. However, if the converted VMs run on a single standalone host (not a pool), their virtual disks can use local storage.

If you want to run the converted VMs in a pool, ensure that you add the shared storage to the pool by creating a storage repository.

Guest operating systems supported for conversion:

You can convert VMware ESXi/vCenter VMs running the following Windows guest operating systems:

- Windows 10 (32-bit) Enterprise edition [Latest tested version is 22H2]
  - Only boot from BIOS mode is supported
- Windows 10 (64-bit) Enterprise edition [Latest tested version is 22H2]
- Windows Server 2016 (64-bit) Standard (Desktop) edition

- Windows Server 2019 (64-bit) Standard (Desktop) edition
- Windows Server 2022 (64-bit) Standard (Desktop) edition

> **Note:**
>
> Only the listed Windows SKUs are supported for conversion.

The following Linux operating systems are also supported:

- Red Hat Enterprise Linux 7.9 (64-bit) with the following configuration:
  - File system: EXT3 or EXT4
  - Boot partition type: btrfs, lvm, or plain
- Red Hat Enterprise Linux 8.x (64-bit) with the following configuration:
  - File system: EXT3 or EXT4
  - Boot partition type: lvm or plain
- Ubuntu 20.04 with the following configuration:
  - File system: EXT3 or EXT4
  - Boot partition type: lvm or regular

For more information about the guest operating systems supported by Citrix Hypervisor, see Guest operating system support.

Meet networking requirements To convert VMware ESXi/vCenter VMs, the XenServer Conversion Manager virtual appliance needs connectivity to a physical network or VLAN that can contact the VMware server. (In the following sections, this network is referred to as the "VMware network".)

If the VMware server is on a different physical network than the hosts in the Citrix Hypervisor pool, add the network to Citrix Hypervisor before conversion.

> **Note:**
>
> - The time it takes for your VMs to be converted depends on the physical distance between your VMware and Citrix Hypervisor networks and also the size of your VM's virtual disk. You can estimate how long the conversion will last by testing the network throughput between your VMware server and XenServer.
> - By default, the XenServer Conversion Manager uses HTTPS to download the VM's virtual disk during VM conversion. To speed up the migration process, you can switch the download path to HTTP. For more information, see VMware's article Improving transfer speed of task with library items.

Map your existing network configuration XenServer Conversion Manager virtual appliance includes features that can reduce the amount of manual networking configuration needed after you convert from your existing VMware ESXi/vCenter VMs to Citrix Hypervisor. For example, XenServer Conversion Manager virtual appliance will:

- Preserve virtual MAC addresses on the VMware ESXi/vCenter VMs and reuse them in the resulting Citrix Hypervisor VMs. Preserving the MAC addresses associated with virtual network adapters (virtual MAC addresses) may:
  - Help preserve IP addresses in environments using DHCP
  - Be useful for software programs whose licensing references the virtual MAC addresses
- Map (virtual) network adapters. XenServer Conversion Manager virtual appliance can map VMware networks onto Citrix Hypervisor networks so that after the VMs are converted, their virtual network interfaces are connected accordingly.

  For example, if you map VMware 'Virtual Network 4'to Citrix Hypervisor 'Network 0', any VMware VM that had a virtual adapter connected to 'Virtual Network 4'is connected to 'Network 0'after conversion. XenServer Conversion Manager virtual appliance does not convert or migrate any hypervisor network settings. The wizard only alters a converted VM's virtual network interface connections based on the mappings provided.

  > **Note:**
  >
  > You do not need to map all of your VMware networks on to the corresponding Citrix Hypervisor networks. However, if you prefer, you can change the networks the VMs use, reduce, or consolidate the number of networks in your new Citrix Hypervisor configuration.

  To gain the maximum benefit from these features, Citrix recommends the following:
  - Before installing Citrix Hypervisor, plug the hosts into the networks on the switch (that is, the ports) that you would like to configure on the host.
  - Ensure that the Citrix Hypervisor pool can see the networks that you would like to be detected. Specifically, plug the Citrix Hypervisor hosts into switch ports that can access the same networks as the VMware cluster.

  Though it is easier to plug the Citrix Hypervisor NICs into the same networks as the NICs on the VMware hosts, it is not required. If you would like to change the NIC/network association, you can plug a Citrix Hypervisor NIC into a different physical network.

Prepare for the XenServer Conversion Manager virtual appliance networking requirements When you perform a conversion, you must create a network connection to the network where the

VMware server resides. XenServer Conversion Manager virtual appliance uses this connection for conversion traffic between the Citrix Hypervisor host and the VMware server.

To create this network connection, you must perform two tasks:

- When you import the XenServer Conversion Manager virtual appliance, specify the network you added for conversion traffic as a virtual network interface. You can do so by configuring **interface 1** so it connects to that network.
- Before you run the conversion wizard, add the network connecting VMware and Citrix Hypervisor to the Citrix Hypervisor host where you want to run the converted VMs.

By default, when you import the XenServer Conversion Manager virtual appliance, XenCenter creates one virtual network interface associated with Network 0 and NIC0 (eth0). When adding a network for conversion, select a network other than XenServer's management network to improve performance in busy pools. For more information about the management interface, see Networking.

Inside the XenServer conversion manager, you might see multiple network interfaces (eth0 and eth1). eth0 attaches to the host's internal network which is used to communicate with the local dom0. eth1 attaches to the routable network which is used to communicate with XenCenter.

To add a network to Citrix Hypervisor:

1. In the **Resource** pane in XenCenter, select the pool where you would like to run XenServer Conversion Manager virtual appliance.
2. Click the **Networking** tab.
3. Click **Add Network**.
4. On the **Select Type** page, select **External Network**, and click **Next**.
5. On the **Name** page, enter a meaningful name for the network (for example, "VMware network") and a description.
6. On the **Interface** page, specify the following:
    - **NIC**. The NIC that you want Citrix Hypervisor to use to create the network. Select the NIC that is plugged in to the physical or logical network of the VMware server.
    - **VLAN**. If the VMware network is a VLAN, enter the VLAN ID (or "tag").
    - **MTU**. If the VMware network uses jumbo frames, enter a value for the Maximum Transmission Unit (MTU) between 1500 and 9216. Otherwise, leave the MTU box t its default value of 1500.

      > **Note:**

> Do not select the **Automatically add this network to new virtual machines** check box.

7. Click **Finish**.

Meet storage requirements Before you convert batches of VMware ESXi/vCenter VMs, consider your storage requirements. Converted VM disks are stored on a Citrix Hypervisor storage repository.

This storage repository must be large enough to contain the virtual disks for all the converted VMs you want to run in that pool. For converted machines that only run on a standalone host, you can specify either local or shared storage as the location for the converted virtual disks. For converted machines running in pools, you can only specify shared storage.

To create a storage repository:

1. In the **Resource** pane in XenCenter, select the pool where you intend to run the XenServer Conversion Manager virtual appliance.
2. Click the **Storage** tab.
3. Click **New SR** and follow the instructions in the wizard. For more instructions, press **F1** to display the online help.

Citrix Hypervisor requirements You can run VMs converted with this release of XenServer Conversion Manager on the following versions of Citrix Hypervisor:

- Citrix Hypervisor 8.2 Cumulative Update 1

VMware requirements XenServer Conversion Manager virtual appliance can convert VMware ESXi/vCenter VMs from the following versions of VMware:

- vCenter Server 6.7.x, 7.x, and 8.x
- vSphere 6.7.x, 7.x, and 8.x
- ESXi 6.7.x, 7.x, and 8.x

> **Note:**
>
> XenServer Conversion Manager virtual appliance cannot convert VMware ESXi/vCenter VMs with four or more disks into Citrix Hypervisor VMs. Your VMware ESXi/vCenter VMs must have three or fewer disks.
>
> Your VMware ESXi/vCenter VMs must also have a network and a storage controller configured.

Prepare to import the virtual appliance Before importing the virtual appliance, note the following information and make the appropriate changes to your environment, as applicable.

Download the virtual appliance The XenServer Conversion Manager virtual appliance is packaged in XVA format. You can download the virtual appliance from the Citrix Hypervisor 8.2 Premium Edition page. When downloading the file, save it to a folder on your local hard drive (typically, but not necessarily, on the computer where XenCenter is installed). After the `.xva` file is on your hard drive, you can import it into XenCenter.

> **Note:**
>
> XenServer Conversion Manager virtual appliance is available for Citrix Hypervisor Premium Edition customers or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about Citrix Hypervisor licensing, See Licensing. To upgrade, or to buy a Citrix Hypervisor 8.2 license, visit the Citrix website.

Virtual appliance prerequisites The XenServer Conversion Manager virtual appliance requires a minimum of:

- Citrix Hypervisor 8.2 Cumulative Update 1
- Disk space: 30 GB of disk space
- Memory: 6 GB
- Virtual CPU allocation: 2 vCPU

Import and configure the virtual appliance

The XenServer Conversion Manager virtual appliance is a single pre-installed VM designed to run on a Citrix Hypervisor host. Before importing it, review the prerequisite information and considerations in the section called *Preparing to import the virtual appliance*.

Import the virtual appliance into Citrix Hypervisor

To import the XenServer Conversion Manager virtual appliance into the pool or host where you want to run the converted VMs, use the XenCenter **Import** wizard:

1. Open XenCenter. Right-click on the pool (or host) into which you want to import the virtual appliance package, and select **Import**.
2. Browse to locate the virtual appliance package.
3. Select the pool or a *home server* where you want to run the XenServer Conversion Manager virtual appliance.

> **Note:**
>
> A home server is the host that provides the resources for a VM in a pool. While it can, a Citrix Hypervisor attempts to start the VM on that host, before trying other hosts. If you select a host, the XenServer Conversion Manager virtual appliance uses this host as its home server. If you select the pool, the virtual appliance automatically starts on the most suitable host in that pool.

4. Choose a storage repository on which to store the virtual disk for the XenServer Conversion Manager virtual appliance and then click **Import**. To add a storage repository to the pool, see the section called "Meet Storage Requirements."You can choose either local or shared storage.
5. Ensure the network to be used for conversion (which connects the VMware server to the Citrix Hypervisor host) is selected as the network associated with **interface 1** ("virtual NIC 1").

   - If the correct network does not appear beside interface 1, use the list in the **Network** column to select a different network.
   - If you have not added the VMware network that is on a different physical network than the pool, do the following:

     a) Exit the wizard.
     b) Add the network to the pool.
     c) Rerun the wizard.

     For more information, see **To add a network to Citrix Hypervisor**.

     > **Warning:**
     >
     > Do NOT configure NIC0 to your customer network. Assign NIC0 only to "Host internal management network."

6. Leave the **Start VM after import** check box enabled, and click **Finish** to import the virtual appliance.
7. After importing the `.xva` file, the XenServer Conversion Manager virtual appliance appears in the **Resources** pane in XenCenter.

Configure the XenServer Conversion Manager virtual appliance

Before you can use the XenServer Conversion Manager virtual appliance to convert VMware ESXi/vCenter VMs, configure it using the XenCenter **Console** tab:

1. After importing the XenServer Conversion Manager virtual appliance, click the **Console** tab.

2. Read the license agreement. To view the contents of the license agreement, open the URL in a web browser. Press any key to continue.
3. Enter and confirm a new root password for the XenServer Conversion Manager virtual appliance. Citrix recommends selecting a strong password.
4. Enter a host name for the XenServer Conversion Manager virtual appliance.
5. Enter the domain suffix for the virtual appliance. For example, if the fully qualified domain name (FQDN) for the virtual appliance is `citrix-migrate-vm.domain4.example.com`, enter `domain4.example.com`.
6. Enter **y** to use DHCP to obtain the IP address automatically for the XenServer Conversion Manager virtual appliance. Otherwise, enter **n** and then enter a static IP address, subnet mask, and gateway for the VM.
7. Review the host name and network setting and enter **y** when prompted. This step completes the XenServer Conversion Manager virtual appliance configuration process.
8. When you have successfully configured the appliance, a login prompt appears. Enter the login credentials and press Enter to log in to the XenServer Conversion Manager virtual appliance.

Convert VMware ESXi/vCenter VMs

When you convert VMware ESXi/vCenter VMs, they are imported into the Citrix Hypervisor pool or standalone host where you are running the XenServer Conversion Manager virtual appliance. Converted VMs retain their original VMware settings for the virtual processor and virtual memory.

Before you start the conversion procedure, ensure that the following is true:

- You have the credentials for the Citrix Hypervisor pool (or standalone host). Either the root account credentials or a Role-Based Access Control (RBAC) account with the Pool Admin role configured is acceptable.
- You have the credentials for the VMware server containing the VMs you want to convert. The conversion procedure requires you connect the XenServer Conversion Manager Console to the VMware server.
- The VMware virtual machines to convert are powered off.
- The VMware virtual machines to convert have a network and a storage controller configured.
- The Citrix Hypervisor pool (or host) that run the converted VMs is connected to a storage repository. The storage repository must contain enough space for the converted virtual disks.
- If you want to run your newly converted VMs in a pool, the storage repository must be shared storage. However, if the converted VMs run on a single standalone host (not a pool), you can use local storage.
- The virtual disks of the VM to convert are less than 2 TiB.

- Citrix Hypervisor pool (or host) has networks that the converted VMs use.

To convert your VMware ESXi/vCenter VMs into VMs that can run in a Citrix Hypervisor environment:

1. Ensure that the virtual appliance is installed and running on the Citrix Hypervisor server or pool where you want to import the VMs.
2. In XenCenter, go to **Pool** > **Conversion Manager**.

   The **Conversion Manager** window opens. Wait while the wizard connects to your virtual appliance.
3. Click **New Conversion**.
4. In the **New Conversion** wizard, enter the credentials for the VMware server:

   - **Server**. Enter the IP address or FQDN for the VMware server that contains the VMs you want to convert to Citrix Hypervisor.
   - **Username**. Enter a valid user name for this VMware server. This account must either be a VMware admin account or have a Root role.
   - **Password**. Enter the password for the user account you specified in the **Username** box.

   Click **Next**. XenCenter connects to the VMware server.
5. In the **Virtual Machines** page, select from the list of VMs hosted in the VMware server the VMs that you want to convert. Click **Next**.
6. In the **Storage** page, select the storage repository you want to use during conversion. This storage repository is where the VMs and the virtual disks that you are creating are stored permanently.

   This tab indicates the proportion of available storage that the virtual disks of the converted VMs consume.
7. On the **Networking** page, for each VMware network listed, select the Citrix Hypervisor network to map it to. You can also select whether to preserve virtual MAC addresses. Click **Next**.
8. Review the options you configured for the conversion process. You can click **Previous** to change these options. To proceed with the configuration shown, click **Finish**.

   The conversion process begins. Conversion from ESXi or vSphere can take several minutes depending on the size of the virtual disks.

   After converting your VMs, the Conversion Manager automatically shuts down by itself, saving resources on the host. Start a VM by selecting the VM's host and then clicking **Pool** > **Conversion Manager**.

The **Conversion Manager** window displays conversions in progress and completed conversions.

Steps after conversion

For Windows VMs, you must install XenServer VM Tools (formerly Citrix VM Tools) for Windows. For Linux VMs, you do not need to install XenServer VM Tools for Linux as the Conversion Manager automatically installs it during the conversion process.

In XenCenter, perform the following steps on your newly converted VMs:

On Windows machines:

1. On Windows VMs, depending on your Microsoft licensing model, you might have to reactivate the VM's Windows license. This reactivation happens because the Windows operating system perceives the conversion as a hardware change.
2. On Windows VMs, install XenServer VM Tools (formerly Citrix VM Tools) for Windows to obtain high-speed I/O for enhanced disk and network performance. XenServer VM Tools for Windows also enable certain functions and features, including cleanly shutting down, rebooting, suspending, and live migrating VMs. You can download the XenServer VM Tools for Windows from the Citrix Hypervisor downloads page.

If you are working with a VM that does not have XenServer VM Tools installed, a XenServer VM Tools not installed message appears on the **General** tab in the **General** pane.

> **Note:**
>
> XenServer VM Tools for Windows must be installed on each Windows VM for the VM to have a fully supported configuration. Although Windows VMs function without XenServer VM Tools for Windows, their performance can be impacted.

Enable VNC On Linux machines

On Linux VMs, configure the VNC server. For more information, see Enable VNC for Linux VMs.

> **Note:**
>
> The VNC password must have at least six characters.

Other conversion tasks

The **Manage Conversions** window enables you to perform other tasks related to converting VMs. These tasks include clearing jobs, saving a summary of jobs, retrying jobs, canceling jobs, and displaying the log file.

To clear all jobs:

1. Select **Clear All**.

2. When prompted to confirm this action, click **Yes** to continue.

To save a summary of jobs:

1. Click **Export All**.
2. Specify where to save the CSV file.
3. Click **Save**.

To retry a job:

1. Select the job from the list.
2. Click **Retry**.

> **Note:**
>
> The **Retry** option is only enabled for failed or canceled jobs.

To cancel a job:

1. Select the job from the list.
2. Click **Cancel**.

> **Note:**
>
> Cancel jobs is only enabled for queued or running jobs.

To save the conversion log file for a single job:

1. Select the job from the list.
2. From the logs menu, Click **Fetch Selected Log**.
3. Specify where to save the log file.

To save the conversion log file for all jobs:

1. From the logs menu, Click **Fetch All Logs**.
2. Specify where to save the log file.

To display conversion details:

1. Select the job from the list.

   The information is displayed in the **Details** panel.

Troubleshoot XenServer Conversion Manager

March 6, 2024

This section provides information about troubleshooting the conversion process and converted VMs.

Problems starting a converted VM

In general, conversion runs smoothly and XenServer Conversion Manager (formerly Citrix Hypervisor Conversion Manager) virtual appliance converts VMs without any issues. However, in some rare cases, you might receive errors when attempting to open converted VMs. The following sections provide some guidance on resolving errors and other issues.

Blue screen with Windows STOP code 0x0000007B

This stop code indicates that XenServer Conversion Manager virtual appliance was unable to configure a Windows device that is critical to boot in Citrix Hypervisor for the first time. Save the logs and send them to Citrix Support for further guidance.

Windows product activation

Depending on your licensing model, an error message on system activation might appear when you attempt to start a Windows VM.

Lost network settings in a Windows VM

If you import a Windows VM from an ESXi server to Citrix Hypervisor, the IPv4/IPv6 network settings can be lost. To retain the network settings, reconfigure the IPv4/IPv6 settings after completing the conversion.

Unable to start VMware SCSI disk

If a VMware VM boots from a SCSI disk but also has IDE hard disks configured, the VM might not boot when you convert it to Citrix Hypervisor. This issue occurs because the migration process assigns the IDE hard disks lower device numbers than SCSI disks. However, Citrix Hypervisor boots from the hard disk assigned to device 0. To resolve this issue, rearrange the virtual disk position in XenCenter so that the VM reboots from the virtual disk that contains the operating system.

To change the position of the virtual disk containing the operating system:

1. In the XenCenter **Resources** pane, select the powered off guest VM.
2. Select the **Storage** tab.

3. From the **Virtual Disks** list, select the virtual disk containing the operating system and then click **Properties**.
4. In the virtual disk's **Properties** dialog, click the ***vm_name*** tab to display device options.
5. From the **Device Position** list, select **0** and Click **OK**.

Problems during conversion

If you experience problems or errors when converting VMs, try exporting the VMware VM as an OVF package. If you cannot export the VMware VM as an OVF package, Conversion Manager cannot convert this VM. Use the error messages you receive when attempting to export the VM as an OVF package to troubleshoot and fix the issues with your VMware VM. For example, you might have to configure a network or a storage controller before the VM can be exported as an OVF package or converted. For more information about troubleshooting your VMware ESXi/vCenter VMs, see the VMware documentation.

If you see any errors when converting Linux VMs, remove the converted VM, restart the XenServer Conversion Manager virtual appliance and retry.

Logs of failed conversions are stored in the XenServer Conversion Manager virtual appliance and can be retrieved by clicking **Fetch All Logs** on the **Conversion Manager** window. When you contact Citrix support to raise any issues, we recommend that you provide the conversion log file and, also, a full server status report for troubleshooting. For more information, see Creating a Server Status Report.

Monitoring System Performance

May 25, 2023

The **Performance** tab in XenCenter provides real time monitoring of performance statistics across resource pools. This tab also provides 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 the following resources go over a specified threshold on a managed server, virtual machine, or storage repository:

  - CPU
  - Memory usage
  - Network
  - Storage throughput
  - VM disk activity

For more information, see Configuring performance alerts.

> **Note:**
>
> Full performance data is only available for virtual machines with Citrix VM Tools installed.

Viewing Performance Data

May 25, 2023

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.
- You can add graphs showing extra resource usage data, if necessary. For example, you can include the Control Domain Load. This load is the average (Linux `loadavg`) of the number of processes queued inside the control domain over the last 5 minutes.
- Lifecycle events for all the VMs hosted on the server are shown in the **VM Lifecycle Events** pane.

For VMs, 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 is happening on the machine. This graph also allows you to adjust the time frame that is shown in the other graphs. The time frame can be changed either to show data from a longer or shorter period, or to show data from an earlier period.

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**. Select **1 Hour**, **1 Day**, **1 Week**, **1 Month**, or **1 Year**.
- To resize the time period that is 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:

  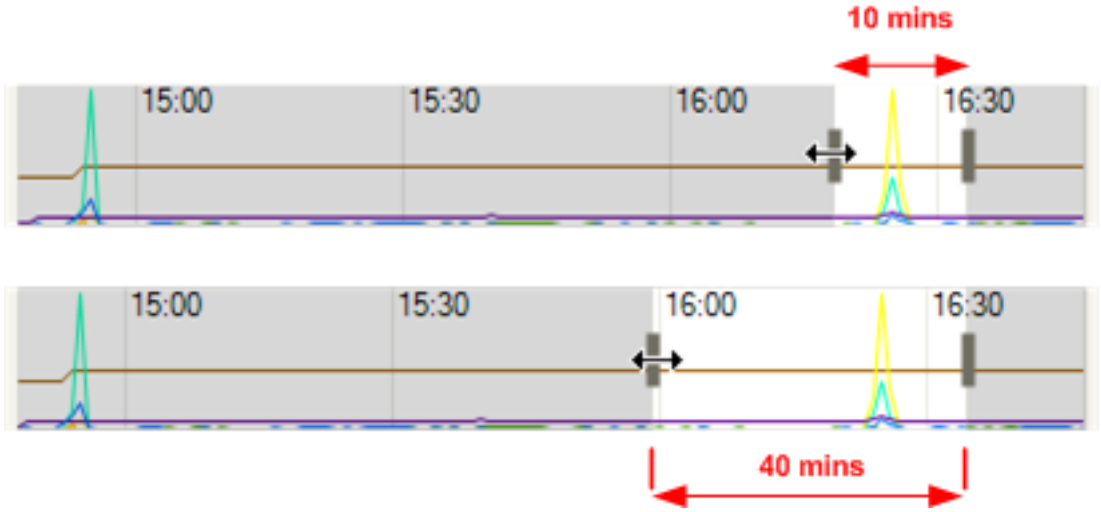

To view data from a different time period

To move the time frame for data displayed in the graphs, point to any graph. When the pointer changes to a move cursor, drag the graph or the sample area in the summary graph to the left or right.

For example:

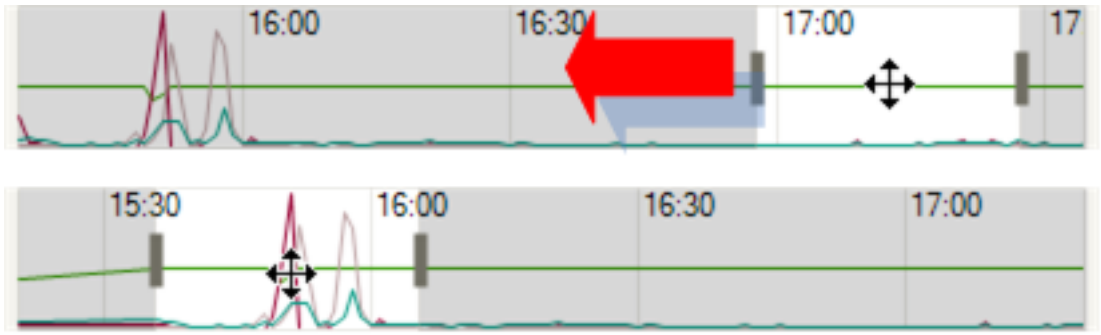

To view VM lifecycle event data on a server

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 the list.
- Double clicking or pressing **Enter** zooms 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

May 25, 2023

To add a graph

1. On the **Performance** tab, select **Actions** and then **New Graph**. The **New Graph** dialog box is displayed.
2. Enter a name for the graph in the **Name** field.
3. From the list of datasources, check 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. Select **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 to remove from the list of graphs displayed on the **Performance** tab.
2. Select **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. Select 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 want to change the color of the datasource. The graph details dialog box is displayed.
3. Check the colored check box located against the required datasource and choose 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:

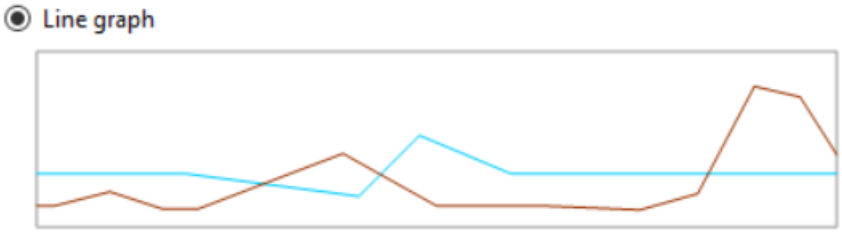

Area Graph:

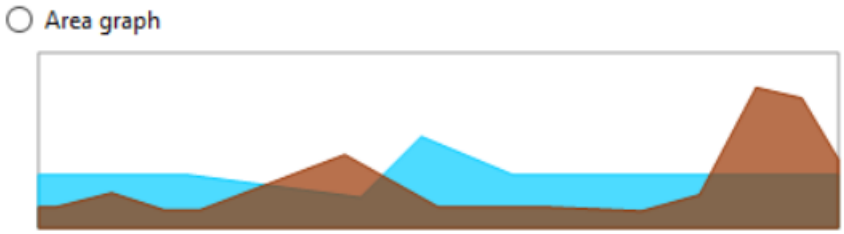

To change the graph type:

1. On the **Tools** menu, select **Options** and then select the **Graphs** tab.
2. To view performance data as a line graph, choose the **Line graph** radio button.
3. To view performance data as an area graph, choose the **Area graph** radio button.
4. Click **OK** to save your changes.

Configuring Performance Alerts

February 5, 2024

Performance alerts can be generated when CPU, memory usage, network, storage throughput, or VM disk activity exceeds a specified threshold on a server, VM, or storage repository (SR). By default, the alert repeat interval is set to 60 minutes, and it can be modified if necessary.

Performance alerts appear in the **Alerts** view (accessed by clicking the **Notifications** button on the left hand pane). You can have performance alerts emailed to you. For more information, see XenCenter Alerts.

You can configure performance alerts for servers, VMs, or SRs. To configure performance alerts:

1. Select the server, VM, or SR in the **Resources** pane. Select the **General** tab and then click **Properties**.
2. Select the **Alerts** tab. The following table summarizes which alerts are available for servers, VMs, or SRs:

Alert name	Server	VM	SR	Description
Generate CPU usage alerts	X	X		Set the CPU usage and time threshold that trigger the alert.
Generate memory usage alerts	X			Set the memory usage and time threshold that trigger the alert.
Generate control domain memory usage alerts	X			Set the control domain memory usage and time threshold that trigger the alert.
Generate disk usage alerts		X		Set the disk usage and time threshold trigger the alert.

Alert name	Server	VM	SR	Description
Generate storage throughput alerts			X	Set the storage throughput and time threshold that trigger the alert. Note: Physical Block Devices (PBD) represent the interface between a specific XenServer server and an attached SR. When the total read/write SR throughput activity on a PBD exceeds the threshold you have specified, alerts are generated on the server connected to the PBD. Unlike other XenServer server alerts, this alert must be configured on the SR.
Generate network usage alerts	X	X		Set the network usage and time threshold that trigger the alert.

To change the alert repeat interval, enter the number of minutes in the **Alert repeat interval** box. After an alert threshold is reached and an alert generated, another alert is not generated

until after the alert repeat interval has elapsed.

3. Click **OK** to save your changes.

Updates and Upgrades

May 25, 2023

XenCenter issues notifications about available Citrix Hypervisor and XenCenter updates and upgrades on the **Updates** tab in the **Notifications** view.

XenCenter is configured by default to automatically check for new Citrix Hypervisor and XenCenter updates and upgrades at regular intervals. You are notified when a new update or product version is available. It is recommended you install all published updates. You can check for available updates manually at any time:

1. Select **Notifications** > **Updates**
2. Select **Refresh**.

Follow this process to verify you are running the latest version of both Citrix Hypervisor and XenCenter.

Applying updates to Citrix Hypervisor hosts

Updates to a version of Citrix Hypervisor 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 Citrix Hypervisor are also delivered as updates. You can apply a Current Release as an update to some previous Current Releases of Citrix Hypervisor. This update moves you to a newer version of Citrix Hypervisor. For more information about the supported update paths for Current Releases, see Install.

Updating your version of XenCenter

The most up-to-date version of XenCenter is supplied on the Citrix Hypervisor product downloads page. Use this file to update your XenCenter installation. For more information, see Updating XenCenter.

Upgrading Citrix Hypervisor hosts

To upgrade Citrix Hypervisor 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

March 4, 2024

You can use the **Rolling Pool Upgrade** wizard to upgrade Citrix Hypervisor - standalone servers or a pool of servers to a newer version.

> **Note:**
>
> Rolling Pool Upgrade is available for licensed Citrix Hypervisor customers, or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement. For more information, see About Citrix Hypervisor 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 prechecks to ensure that certain pool-wide features, such as high availability and WLB, are temporarily disabled. THe wizard also checks that each host in the pool is prepared for upgrade. 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 Citrix Hypervisor 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 Citrix Hypervisor installation media or specify a PXE boot server for each server that you upgrade.
- In automatic mode, the wizard uses network installation files on an HTTP, NFS, or FTP server to upgrade each server in turn. This mode does not require you to insert install media, manually reboot, or step through the installer on each server. If you perform a rolling pool upgrade in this manner, unpack the Citrix Hypervisor 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 servers, that is, servers which do not belong to any resource pool.

Important: Before you upgrade

Upgrading a pool of servers requires careful planning. As you plan your upgrade, it is important to be aware of the following:

- Download and install the latest version of XenCenter. For example, when you are upgrading your hosts to Citrix Hypervisor 8.2, you must use XenCenter issued with Citrix Hypervisor 8.2. Using earlier versions of XenCenter to upgrade to a newer version of Citrix Hypervisor is not supported.
- VMs can only be migrated from a server that runs an older version of Citrix Hypervisor to one that runs the same version or higher. You cannot migrate VMs from an upgraded server to one running an older version of Citrix Hypervisor. Ensure to allow for space on your servers accordingly.
- We strongly advise against running a mixed-mode pool (one with multiple versions of Citrix Hypervisor coexisting) for longer than necessary. In this case, the pool operates in a degraded state during upgrade.
- Do not attempt to perform any key control operations during the upgrade process. Though VMs continue to function as normal, VM actions other than migrate might 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 upgrades the pool master first. Do not place the pool master into maintenance mode using XenCenter before performing the upgrade as this action causes a new master to be designated.
- Take a backup of the state of your existing pool using the `pool-dump-database` xe CLI command. For more information, see Command line interface. Backups allow 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, you might have to shut down VMs 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 might choose to perform them before you begin the upgrade:
  - Empty the CD/DVD drives of the VMs in the pool. For details and instructions, see Upgrade.
  - Disable high availability.
  - Disable WLB

To upgrade Citrix Hypervisor 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 select **Next** to continue.
3. Select the pool or standalone servers that you would like to upgrade and then select **Next**.
4. Choose the **Upgrade Mode**.

    - You can select **Automatic Mode** for an automated upgrade from network installation files on an HTTP, NFS, or FTP server. If you choose the **Automatic Mode**, specify the location of the network installation file, user name, and password before continuing to the next step.
    - You can select **Manual Mode** for a manual upgrade from either a CD/DVD or a server using PXE boot. If you choose the **Manual Mode**, run the Citrix Hypervisor 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 Citrix Hypervisor installation media or specify a PXE boot server for each server that you upgrade.

5. On the **Upgrade Options** page, choose whether you want XenCenter to automatically download and install the minimal set of updates after upgrading the servers to a newer version. The apply updates option is selected by default. However, you must have an internet connection to download and install the updates.

    In addition, to allow XenCenter to install a previously downloaded update or a supplemental pack after upgrading the servers, select **Install an update or supplemental pack from disk**. Select **Browse** to choose the file. Selecting a file incompatible with the upgraded version of Citrix Hypervisor can cause the installation to fail.

6. After choosing the Upgrade Options, select **Run Prechecks**.
7. Follow the on-screen recommendations to resolve any upgrade prechecks that have failed. To allow XenCenter to automatically resolve all failed prechecks, select **Resolve All**. When the prechecks have been resolved, select **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 process fails for any reason, the Rolling Pool Upgrade wizard halts the process. This halt 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. Select **Finish** to close the wizard.

Updating Managed Servers

May 25, 2023

XenCenter issues notifications about available Citrix Hypervisor updates in the **Updates** tab on the **Notifications** view. Updates to Citrix Hypervisor 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 Citrix Hypervisor releases in the Long Term Service Release (LTSR) and Current Release (CR) streams.

  Hotfixes on the latest CR are available to all Citrix Hypervisor 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 Licensing.

- **Cumulative Updates**, which contain previously released hotfixes and might contain support for new guests and hardware. Cumulative Updates are provided for Citrix Hypervisor or XenServer releases in the LTSR stream and are available to customers with an active CSS account.

- **Current Releases**, which are full versions of Citrix Hypervisor from the Current Release (CR) stream.

To provide these update notifications, XenCenter requires internet access. If your XenCenter is behind a firewall, ensure that it has access to the `updates.ops.xenserver.com` domain and subdomains on the `citrix.com` domain through the firewall.

This topic contains information about applying Citrix Hypervisor updates to your managed servers. Pay careful attention to the release note that is published with each update. Each update might have unique installation instructions, particularly regarding preparatory and post-update operations. Some updates might only be available to licensed users or to Customer Success Services customers.

In addition to Citrix Hypervisor updates, the **Updates** tab also notifies users about the availability of new versions of Citrix Hypervisor and new versions of XenCenter. Some new Citrix Hypervisor current releases can be applied as updates to earlier versions of Citrix Hypervisor. However, you can only reach some new releases by the upgrade process. For information about upgrading Citrix Hypervisor, 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 select **Dismiss** and then **Dismiss Selected**. To dismiss all the updates, select **Dismiss All**. If you would like to see updates which were previously dismissed, select **Restore Dismissed Updates**.

Authenticating your XenCenter to receive updates

To provide a more secure service for hotfix downloads, XenCenter now requires that you authenticate it with Citrix to automatically download and apply hotfixes.

Citrix Hypervisor updates are hosted on the Citrix Support site. The support site restricts the download of these updates to customers with a Citrix account. Some downloads are restricted to customers who have an active Citrix Success Services (CSS) agreement.

Prerequisites

If your organization is an existing Citrix customer, ensure that your Citrix account meets the following requirements before using it to generate a client ID file:

- You are registered as a contact for your organization.
- Citrix Customer Service created your Citrix account as a web login associated with the registered contact.

To complete these steps, you can contact Citrix Customer Service.

To be able to download the hotfixes that are restricted to CSS customers, your organization must have an active Citrix Success Services agreement.

> **Note:**
>
> If you are unsure whether your Citrix account can be used to download CSS-only hotfixes through XenCenter, log in to https://support.citrix.com and check that you can download a CSS-only hotfix file through the browser.

Importing a client ID JSON file

To receive updates to Citrix Hypervisor through XenCenter, you must request a client ID JSON file from Citrix and import this file into your XenCenter instance. This setup task is required only once for each user of an instance of XenCenter.

1. In the XenCenter menu, go to **Tools > Options**. The **Options** window opens.
2. In the **Updates** tab, go to the **Client ID** section.
3. Click the provided link to go to the page **Generate and Download a Client ID** in your web browser.
4. You are prompted to log in to your Citrix account. If your organization has an active Citrix Success Services agreement, ensure that you use a Citrix account that is associated with this organization.

> **Note:**
>
> Creating a user account using the **Need an Account** on this login screen does not associate your new user account with any existing organization and its associated Citrix Success Services entitlements.

5. After you log in, click the **Download Client ID** button. The client ID is provided as a JSON file (`xencenter_client_id.json`).
6. Return to XenCenter.
7. In the **Location** field, browse to the location of the JSON file you downloaded (`xencenter_client_id.json`) and select the file.
8. Click **OK**.

If you do not complete these steps in advance, XenCenter prompts you to obtain and install a client ID file when you first use it to install an update.

About the client ID

- The client ID is unique to your Citrix account.
- The client ID does not expire.
- The client ID is not affected by a password change or password expiry in your Citrix account.
- The client ID is not revoked by changes to Citrix account privileges.

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 Disaster recovery and backup.
3. We recommend that you reboot all your servers before installing an update and then verify their configuration. For example, check that the VMs start and that storage is accessible. This recommendation is because some configuration changes take effect only after a server is rebooted. The reboot might uncover configuration problems that might 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. 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. Schedule the updates 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 Citrix Hypervisor and XenCenter updates. Select **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. Select **Check for Updates Now** to manually check for updates.

You can select from the **View** menu 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. You can order these updates 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, select **Export All**. The `.csv` file lists the following information:
  - The update name
  - A description of the update
  - The servers that this update can be applied to
  - The timestamp of the update
  - A reference to the webpage that the update is downloaded from
- To apply an update to a server, from the **Actions** menu for that update select **Download and Install**. This action 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, select the **Actions** menu 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, select **Export All**. The `.csv` file lists the following information:
  - The **Pool** that the server belongs to
  - The **Server** name
  - The **Status** of the installed Citrix Hypervisor
  - The update **Status** of the server
  - The **Required Updates** for this server
  - The **Installed Updates** for this server
- To apply the updates, select **Install Updates**. This action 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. When a cumulative update is available for the currently applied release, XenCenter applies the new cumulative update baseline and all available hotfixes for that cumulative update. If you do not want to update to the cumulative update baseline, instead manually download the hotfixes available for the currently applied release and apply them to your servers.

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. This mechanism also enables you to apply an update to multiple servers and pools using the **Install Update** wizard. During the process, the **Install Update** wizard automatically performs these steps:

1. It migrates VMs off each server
2. It places the server in Maintenance mode
3. It applies the update
4. It reboots the server if necessary
5. It migrates 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 high availability, are reverted.

When you install a Current Release, the **Install Update** mechanism offers to apply the minimum set of 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 *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 select **Next** to continue.
3. Select the updates to install and select **Next** to continue.
4. Select the servers to install updates on and select **Next** to continue.

   Notes:

   - If you are installing a current release, XenCenter also offers to apply the minimum set of updates (hotfixes) after installing the current release.
   - 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.

The **Install Update** wizard performs various prechecks 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 Citrix Hypervisor.

Follow the on-screen recommendations to resolve any update prechecks that have failed. If you prefer XenCenter to automatically resolve all failed prechecks, select **Resolve All**. When the prechecks have been resolved, select **Next** to continue.

If you are installing a current release, 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 halt allows you to fix the issue and resume the update process by clicking the **Retry** button.

See Step 10 to complete the current release installation process.

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 select **Cancel** at this stage, the **Install Update** wizard reverts the changes and removes the update file from the server.

Select **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.

Select **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.

1. From the XenCenter menu, select **Tools** and then **Install Update**.
2. Read the information displayed on the **Before You Start** page and then select **Next**.
3. On the **Select Update** page, select **Browse** to locate the update file and then select **Open**. Select **Next** to continue.
4. Select the pool and servers that you would like to update. Any servers or pools that cannot be updated are grayed out. Select **Next** to continue.
5. Follow the instructions on the Install Update wizard to complete the update installation process.
6. Select **Finish** to exit the wizard.

Live Patching in Citrix Hypervisor

November 16, 2023

Citrix Hypervisor customers who deploy Citrix Hypervisor hosts might often need to reboot their servers after applying hotfixes. This rebooting results in unwanted downtime for the servers while customers have to wait until the system is restarted. Reboots reduce the uptime of the servers and impact business. Live patching enables customers to install some Linux kernel and Xen hypervisor updates without having to reboot the servers. This feature reduces maintenance costs and downtime. Such hotfixes consist of the following components:

- A live patch that is applied to the memory of the server
- 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. Complete these tasks for the update to take effect. In addition, the wizard checks whether you must reboot the servers after applying the update and displays the result. This feature enables customers to know the post-update tasks well in advance and schedule the application of updates accordingly.

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

Live patching scenarios

Hotfixes can be live patched across pools, servers, or on a standalone server. Some updates might need you to reboot the server, 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 the 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 might be required.
- **Updates without live patch** -No change in the behavior here. It works as usual.

> **Note:**
>
> If a server doesn't require a reboot or if the update contains live patches that can apply to the servers, XenCenter displays **No action required** on the **Update Mode** page.

Applying Automated Updates

November 9, 2023

XenCenter allows you to apply automated updates that are needed to bring your servers up-to-date. You can apply these updates simultaneously 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. Where possible, XenCenter limits it to a single reboot at the end.

As a prerequisite, XenCenter requires internet access to fetch the required updates. If your XenCenter is behind a firewall, ensure that it has access to the `updates.ops.xenserver.com` domain and subdomains on the `citrix.com` domain through the firewall.

When you choose to apply automated updates, all the required updates get applied. Automated updates do apply any Cumulative Updates that are available for a host. If a new Current Release version is available as an update, automated updates do not apply this update. In this case, manually select to update 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 is not displayed.

    - **Installed supplemental packs** - lists supplemental packs 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:**
>
> Automated Updates were previously restricted to Citrix Hypervisor Premium Edition customers or Citrix Virtual Apps and Desktops customers. However, in pools with hotfix XS82ECU1053 applied, this feature is available to all users.

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 select **Next**.
3. Select **Automated Updates**. This option is visible only if XenCenter is connected to at least one licensed pool or a licensed standalone server.
4. Select **Next**.
5. Select one or more pools or standalone servers to update and select **Next**. Any server or pool that cannot be updated appears grayed out.
6. The **Install Update** wizard performs several 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, select **Resolve All**.
7. When the prechecks have been resolved, select **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 are uploaded to the default SR of the pool. If the default SR is not a shared SR or has insufficient 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 is uploaded to the local storage of the pool master.
   > - If the update process cannot complete for any reason, XenCenter halts the process. This halt allows you to fix the issue and resume the update process by clicking the **Retry** button.
8. When the updates have been applied, select **Finish** to close the **Install Update** wizard.

Installing Supplemental Packs

November 16, 2023

Supplemental packs are used to modify and extend the functionality of the Citrix Hypervisor server, by installing software into the control domain (Dom0). Users can add supplemental packs either during the initial Citrix Hypervisor installation, or at any time afterwards. When you upgrade Citrix Hypervisor, previously applied supplemental packs are removed by upgrade and so they must be reapplied during or after the upgrade. Facilities also exist for OEM partners to add their supplemental packs to

Citrix Hypervisor installation repositories to allow automated factory installations. For more information, see 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 Citrix Hypervisor 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, select **Browse** to add the supplemental pack and then click **Next** to continue.
5. On the **Select Servers** page, select the pool or server to apply the supplemental pack to. Click **Next**. This action uploads the supplemental pack to the default SR of the pool or the server.

   > **Note:**
   >
   > If the default SR in a pool is not a shared SR, 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 is 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 is 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 prechecks to determine whether the supplemental pack can be applied onto the selected servers.

   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 select **Cancel** at this stage, the **Install Update** wizard reverts the changes and removes the supplemental pack from the SR.
9. Select **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 is 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.

Install driver disks

October 11, 2023

You can install a driver disk using one of the following methods:

- By using XenCenter (recommended)
- During a clean Citrix Hypervisor installation
- By using the xe CLI

For information on how to install a driver disk during a clean Citrix Hypervisor installation, see Install the Citrix Hypervisor server. For information on how to install a driver disk by using the xe CLI, see Driver disks.

After installing the driver, restart your server for the new version of the driver to take effect. As with any software update, we advise you to back up your data before installing a driver disk.

Install a driver disk by using XenCenter

Perform the following steps to install the driver disk by using XenCenter:

1. Download the driver disk to a known location on a computer that has XenCenter installed.

   You are not required to extract the contents of the zip file. XenCenter can install the driver from the zip file or the iso.

2. In XenCenter, go to **Tools** > **Install Updates**. The **Install Update** dialog opens.

3. On the **Select update** tab, choose **Select an update or supplemental pack from disk** and browse to the location of the zip file or iso.

   Click **Next**.

4. Complete the steps in the dialog to select your servers and install the driver disk.

5. To complete the installation, XenCenter can restart the server now. Alternatively, you can choose to restart manually at a time that is convenient for you.

   > **Note:**
   >
   > The driver does not take effect until after the host is restarted.

Updating XenCenter

May 25, 2023

If automatic update notifications are configured, you might occasionally be notified that a new version of XenCenter is available. New versions of XenCenter are supplied on the Citrix Hypervisor product downloads page.

For more information, see automatic update notification.

To check for new XenCenter versions manually at any time, select **Notifications**, **Updates** and then select **Refresh**.

To provide update notifications, XenCenter requires internet access. If your XenCenter is behind a firewall, ensure that it has access to the `updates.ops.xenserver.com` domain and subdomains on the `citrix.com` domain through the firewall.

To download and install a new version of XenCenter:

1. From the XenCenter navigation pane, select **Notifications** and then **Updates**. This panel displays a list of available updates.
2. Select the required XenCenter update from the list and select **Go to Web Page** from the **Actions** menu. This action opens the Citrix Hypervisor product downloads page in your web browser.
3. Sign in to the website, if necessary.
4. Download the latest version of XenCenter and save the installer to your computer.
5. Exit your current XenCenter session.
6. Browse to the location of your download and double-click the installer `.msi` file to begin installing the new version of XenCenter.

Update Notifications

May 25, 2023

You can configure XenCenter to periodically check for available Citrix Hypervisor and XenCenter updates and new versions.

To configure updates notification:

1. On the **Tools** menu, select **Options** and then select the **Updates** tab.
2. Select **Check for new versions of Citrix Hypervisor** to have XenCenter periodically check and notify you when a new Citrix Hypervisor version is available.
3. Select **Check for Citrix Hypervisor updates** to have XenCenter periodically check and notify you when updates for Citrix Hypervisor 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.

These notifications are displayed in the **Updates** view of the **Notifications** pane.

To provide these update notifications, XenCenter requires internet access. If your XenCenter is behind a firewall, ensure that it has access to the `updates.ops.xenserver.com` domain and subdomains on the `citrix.com` domain through the firewall.

XenCenter Alerts

January 26, 2024

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 can be generated when CPU, memory usage, network, storage throughput, or VM disk activity exceeds a specified threshold on a server, VM, or SR. For information on configuring performance alerts, see Configuring performance alerts.
- **HA (High Availability) status alerts**. Alerts can be generated for changes to a pool's high availability status, such as when a pool becomes over committed.
- **License expiry alerts**. Alerts are generated when Citrix Hypervisor licenses on your managed servers are approaching their expiry dates or have expired.
- **End of life alerts**. (XenCenter 8.1 and later) Alerts are generated when the Citrix Hypervisor versions on your managed servers are approaching or reach their end of life. To ensure that updates to address any future functional and security related issues can be applied, update your environment to a later supported release. Go to the **Updates** view to see what update or upgrade options are available. The end-of-life date for a particular Citrix Hypervisor version can depend on whether your servers are licensed or unlicensed (Express Edition).
- **Certificate alerts**. (XenCenter 8.2 and later) Alerts are generated when the certificate on a Citrix Hypervisor server is approaching its expiry date or has expired. The first alert is generated 30 days before expiry. The severity of the alert increases 14 days and again 7 days before expiry.

Working with alerts

XenCenter is equipped with powerful filtering capabilities. It enables you to filter alerts displayed on the **Alerts** tab. You can view alerts only from specific pools or servers, or only those alerts generated during a specific time period. For some alerts, it might be possible to quickly address the issue that caused the alert to be generated. The following sections list various options available in the **Alerts** view.

Filter by Severity

Filters alerts by their severity

By default, alerts of all severity levels are displayed on the **Alerts** tab. To view alerts of a particular severity, select **Filter by Severity** and then cancel the selection on other severity levels from the list. Select **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 are displayed. To stop displaying alerts from a specific host, select the 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 are displayed. Select the list and select a date range from the list. Alternatively, select **Custom** to define your own date range by specifying the start and end date/time. Select **Show All** to view all the alerts.

Refresh

If new alerts are generated when the **Alerts** tab is open, they might not appear in the list. Select **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, select **Dismiss All**. To dismiss a specific set of alerts, select the required alerts from the list, and select **Dismiss Selected**.

Actions

Enables you to perform specific actions on the alerts displayed. The **Actions** list displays all actions available for the selected alert.

Select an alert from the list and then select the action relevant to the alert to address it. For example, select:

- **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.
- **Install Certificates** to update the certificate on a server. This action opens the **Install Certificates** dialog.

Receiving alert notifications by email

You can configure XenCenter to email notifications when alerts are generated for any servers and VMs in a pool, or for a standalone server and its VMs.

When you turn on the email notification feature, you receive 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 Citrix Hypervisor xe CLI. For more information, see Command line interface.

To turn on email notifications

1. Select a pool or standalone server in the **Infrastructure** view.
2. Select the **General** tab and then **Properties**.
3. Select 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:**
   >
   > Enter the details of an SMTP server which does not require authentication. Emails sent

> through SMTP servers which require authentication are not delivered. For instructions on using authenticated SMTP servers to receive email notifications, see Monitor and manage.

5. Select **OK** to save your changes and close the dialog box.

Troubleshooting

May 25, 2023

- XenCenter Alerts
- XenCenter Event Log
- Creating a Server Status Report
- Resolving SR Connectivity Problems
- VM Recovery Mode

XenCenter Event Log

May 25, 2023

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. You can use this record to troubleshoot any problems that might arise during the XenCenter session.

Viewing events in the current session

To view the events summary for your current XenCenter session, select **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 Citrix Hypervisor environment. To quickly locate the XenCenter log file, from the XenCenter menu, select **Help > 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 the various options available in the **Events** view.

Filter by status

Filters events by their progress

By default, all events for the current XenCenter session are displayed. Select a specific status from the menu 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 are displayed. To stop displaying events from a specific host, select the menu 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 are displayed. Select the menu and select a date range from the list. Alternatively, select **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 select **Dismiss All**. To dismiss a specific set of events, select the required events from the list, and select **Dismiss Selected**.

Actions

Enables you to perform specific actions on the events displayed

Select an event from the list and then select:

- **Dismiss** to dismiss the event
- **Go To** to navigate to the host from which the event originated. Selecting this action takes you to the **Infrastructure** view
- **Copy** to copy information about the event to the clipboard

Creating a Server Status Report

May 25, 2023

The **Server Status Report** wizard provides a convenient way to collect and package a comprehensive snapshot of a specific Citrix Hypervisor 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 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

By default, the files gathered for a server status report can be limited in size. If you need log files that are larger than the default, you can run the command `xenserver-status-report -u` in the Citrix Hypervisor server console.

To generate a server status report

On the **Tools** menu, select **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 might 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 select **Next**.
3. **Report Destination**. Browse to locate the folder to save the report in and then select **Next**.

4. **Compile Report**. This page shows the progress of the report compilation and reports any problems with the data collection. When the report compilation is complete, select **Finish** to save the report files to your specified folder and then close the wizard.

Resolving SR Connectivity Problems

May 25, 2023

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 for the SR to be available. If a PBD is unplugged for some reason, the SR is no longer available to the server and appears with a broken storage icon in the **Resources** pane.

You might 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 select **Repair Storage Repository** on the shortcut menu.

Alternatively, on the **Storage** menu, select **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 available.
- **Unplugged**. The storage is unavailable because the PBD is unplugged.
- **Connection missing**. The storage is unavailable because the PBD cannot be found.

Select **Repair** to have XenCenter attempt to repair the storage. Progress and results are displayed within the **Repair Storage Repository** dialog box.

VM Recovery Mode

May 25, 2023

If you experience serious problems with a paravirtualized Linux VM, you can try starting it up in Recovery Mode. This process turns HVM mode on temporarily and sets the CD drive as the first boot. You can 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 want to start in recovery mode.
- From the main menu, choose **VM** > **Start/Shut down** > **Start in Recovery Mode**.

> **Note:**
>
> Attach your usual operating system rescue CD, boot the VM from this CD, and then fix the VM from the rescue CD.

See VMs and Templates to find out more about HVM and paravirtualized modes.

XenCenter Plug-in Specification Guide

December 9, 2023

This document explains how to write a plug-in for XenCenter, the GUI for Citrix Hypervisor. Using the plug-in mechanism third-parties can:

- Create menu entries in the XenCenter menus linked to an executable file or PowerShell script, including full use of the Citrix Hypervisor PowerShell Module (XenServerPSModule) cmdlets.
- Cause a URL to be loaded into a tab in XenCenter.

The XenCenter plug-in mechanism is context aware, allowing you to use XenSearch to specify complicated queries. Also, plug-ins can take advantage of contextual information passed as arguments to executables or as replaceable parameters in URLs.

A XenCenter plug-in consists of the following components:

- An XML configuration file.
- A resource DLL for each supported locale. Currently XenCenter exists in English and Japanese versions only.
- The application and any resources it requires.

Put these components of a plug-in in a subdirectory of the XenCenter installation directory. For example, a default installation of XenCenter requires that a plug-in reside in `C:\Program Files (x86)\Citrix\XenCenter\Plugins\<organization_name>\<plug-in_name>`

XenCenter loads all valid plug-ins found in subdirectories of the plug-ins directory when it starts:

- The plug-in name (<plug-in_name>) must be the same as the directory in which it is placed.
- The resource DLL and the XML configuration file must follow these naming conventions:
  - `<plug-in_name>.resources.dll`
  - `<plug-in_name>.xcplugin.xml`

For example, if your organization is called Citrix and you write a plug-in called Example which runs a batch file called `do_something.bat`, the following files must exist:

- `C:\Program Files (x86)\Citrix\XenCenter\Plugins\Citrix\Example\Example.resources.dll`
- `C:\Program Files (x86)\Citrix\XenCenter\Plugins\Citrix\Example\example.xcplugin.xml`
- `C:\Program Files (x86)\Citrix\XenCenter\Plugins\Citrix\Example\do_something.bat`

These paths assume that you use the default XenCenter installation directory.