



App Layering

Contents

What's new	5
Deprecation	32
Known issues	33
System requirements	43
Install appliance	46
XenServer®	47
Google Cloud	50
MS Azure or Azure Government	58
MS Hyper-V	63
Nutanix AHV	66
VMware vSphere	69
Docker	75
Install the App Layering agent	78
Configure	81
Access management console	82
Change administrator passwords	83
Set up file share	84
Connect to a directory service	85
Assign roles	87
Connector configurations	89
Azure	95
Azure Deployments	96
Authoring ARM templates	102

Starter templates	110
Template parameters	146
Image Portability for Azure	155
XenServer®	157
Citrix Provisioning™ (XenServer, VMware, Hyper-V, Nutanix, Azure)	162
Google Cloud	167
MS Hyper-V	171
Nutanix AHV (Acropolis)	176
VMware vSphere	180
Windows File Share	186
Layering	187
Create and manage layers	189
Prepare the OS for layering	194
Prepare your OS image for layering in XenServer®, Hyper-V, or vSphere	197
Prepare your OS image for layering on Google Cloud	202
Prepare your OS image for layering in Azure	205
Prepare your OS image for layering in Nutanix	209
Create the OS layer	215
Create Platform layer	217
Create or clone an app layer	226
Layer antivirus apps	236
Layer Docker	237
App Layering Recipes	238
Deploy elastic layers	239

App Layering based assignments	244
Studio App Packaging based assignments	249
Deploy user layers	251
Update layer	271
Export and import layers	276
Exclude files from layers (Advanced feature)	281
Publish layered images	284
Create or clone an image template	284
Publish layered images from template	287
Manage image template	289
Image Portability	290
Requirements	293
PowerShell Modules	296
Citrix.AppLayering PowerShell module	297
Citrix.Image.Uploader PowerShell module	299
Manage	301
System settings	302
Storage	306
Appliance settings	309
App Layering services	311
Directory service	313
Migrate Appliance	316
Upgrade	321
Users	325

Users and groups	325
Firewall ports	326

What's new

January 5, 2026

Citrix delivers new features and improvements to Citrix App Layering™ users when they are available. New releases provide more value, so there's no reason to delay updates.

This article covers new and enhanced features, along with the fixed issues in this release.

For the latest App Layering requirements and supported platforms, see [System requirements](#).

App Layering 2511 (This release)

This release contains the following new features:

- **Image Portability feature**

A new feature that allows you to export, prepare, and publish a master catalog image from one compute platform to a different one (e.g. from VMware to Azure). This initial release focuses on exporting an image from Azure, Nutanix, VMwre, and XenServer, and preparing an image in Azure. For details, see the [Image Portability](#) documentation.

- **Studio App Packaging elastic layering**

Users now have the ability to enable app layers for use with Studio App Packages. For more details, see the [Elastic Layering](#) documentation.

- **Support for NVMe disk controllers in Azure**

Citrix App Layering can now produce packaging machines and images that support Azure VM sizes with NVMe disk controllers. To use this feature, ensure your Azure Deployments connector configurations use ARM template specs authored with NVMe support. The provided [starter templates](#) have been updated with NVMe support.

- **Improvements and optimizations for delivering Microsoft Store applications**

- **General improvements for packaging and finalizing app layers**

Fixes

- Explorer.exe may not start on images with user layers or UPL enabled, causing a black screen. [ALHELP-1899]

App Layering 2507

This release contains the following new features:

- **Strict host checking has been enabled by default for all HTTP/HTTPS connections to the App Layering appliance.**

This means the management console can only be accessed via a hostname that is on the allowlist. If you try to access the management console via a hostname not on the allowlist then you will receive a 400 Bad Request error. The allowlist automatically includes the appliance's local IP address and all DNS and IP addresses included as Subject Alternative Names in the [appliance's HTTP certificate](#). The allowlist can be extended by creating/editing the file `/etc/httpd/conf/explicit-alias.conf` on the appliance to have the following content:

```
ServerAlias name1 name2 name3
```

Where each `name*` is a hostname you use to access the management console (the name you use in the browser address bar, e.g. `https://name1/admin-ui`). Strict host checking can be turned off entirely by creating the file `/etc/httpd/conf.d/99-strict-host-check-off.conf` with the following content:

```
StrictHostCheck OFF
```

IMPORTANT:

If you initiate the upgrade from a version prior to 2507 or later when connected via a hostname not on the allowlist, then the upgrade may appear to never finish. This is because the management console will not be able to communicate with the appliance to determine when the upgrade is finished as it will be getting 400 Bad Request errors. To avoid this, initiate the upgrade using a hostname you know will be on the allowlist, such as the IP address of the appliance.

- **Handling WebView2**

The latest version of Microsoft's WebView2 will now be deployed to the images regardless of which layer the highest version is found. All other versions of the WebView2 will not appear in the image. Therefore you don't have to do any special handling of this application while creating application layers, nor are you required to have the OS layer keep it up to date. When using full user layers and UPL the version of the WebView2 found on the image will be used and if there is a later version found in the user's layer, it will not be deleted, but the version from the image will be one that will run, until Microsoft updates the application on its own.

- **The TLS cipher suite `TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA` is disabled by default.**
- **The SSH key exchange (KEX) algorithm `diffie-hellman-group-exchange-sha1` is disabled by default**

- **The App Layering appliance now comes with the required libraries to enable TLS encryption for Rsyslog**

Fixes

- You might see the “Guest OS Version” setting of VMWare VMs created by App Layering incorrectly set to Server 2012. This happens for VMs based on a Windows 11, Server 2019, Server 2022, or Server 2025 OS layer. [UNI-91506]
- You might see “Invalid Credentials or Insufficient Privileges” errors when publishing images, creating layers, and creating or editing connector configurations. This occurs when the connector configuration has offload compositing enabled and is configured to connect to a Nutanix AHV 10.3/AOS 7.3 environment. [UNI-92477]
- You might see the error message “The spool file was not found” when printing with UWP apps. [ALHELP-1817]
- Microsoft Teams might not launch on a published image [ALHELP-1883]
- MSIX packages with ILV and PSF do not work with Elastic Layering enabled [ALHELP-1857]
- The loopback adapter might be missing from published images [ALHELP-1884]
- Sophos installation might fail on the App Layer in windows 11 24H2 MultiSession [ALHELP-1860]
- When you open a file within excel via a sharepoint URL and click on the Sharepoint folder path, the path might redirect to “mydocuments” instead of the sharepoint path. [ALHELP-1827]
- Attempting to publish an image does not automatically shutdown, to complete the process. [ALHELP-1877]

App Layering 2503

This release contains the following new features:

- **Support for Microsoft Office 2024**

Microsoft Office 2024 can now be used with Citrix App Layering.

- **Modify the PVS recipe for Citrix.Xip.ClientService startup behavior**

Starting with this release, the startup type for Citrix.Xip.ClientService in the PVS recipe has been changed to Delayed Start. This update ensures that the Background Intelligent Transfer Service (BITS) starts only after performance counters are rebuilt, preventing issues where BITS counters may not load correctly. Since Citrix.Xip.ClientService is not required during the initial boot process of a PVS image, this change improves system reliability without affecting functionality.

- **By default, the ELM appliance now disables Cipher Block Chaining (CBC) ciphers for SSH.**

Fixes

- You might experience issues with a Nutanix AHV connector configuration when offload compositing is enabled because the system selects the wrong virtual machine (VM) after cloning. [UNI-91714]
- You might see a WebView2 error when launching Teams in App Layering with a full user layer. [UNI-91709]
- Microsoft Teams fails to launch with “the parameter is incorrect” after rebooting a session host with Elastic Layering (EL) enabled and Profile Management. [UNI-91649]
- You may receive frequent high memory usage alerts because of the current calculation method in App Layering. [UNI-91669]
- Microsoft Teams fails to launch when short names are disabled, resulting in the error message ‘The data area passed to a system call is too small.’ [UNI-91520]
- You might see incorrect DNS settings in Oracle Linux 8 (OL8) Enterprise Layer Manager (ELM), including public servers, even after manual removal. [UNI-91636]
- After recent Windows updates, sysprep may fail to clean up the local administrator account’s profile directory and corresponding registry entry during image generalization. [UNI-91361]
- You might face issues in finding Microsoft Teams in the Start search bar after installation in App Layering. [ALHELP-1790]
- Microsoft Teams does not launch on Server 2019 due to reparse point behavior specific to this OS. The [LocalCache](#) folder created for Teams is a reparse point with a specific tag. [UNI-91177]
- In the vSphere Connector UI, the template and folder fields are cleared when the connection information (server, user, password) is updated. This can be inconvenient, as passwords often need to be updated. UNI-90970
- You might face issues with the packaging machine after upgrading an OS layer from Windows 10 to Windows 11 24H2 if caching is enabled. [UNI-91593]
- You might be unable to boot finalized template in Azure Hypervisor as it remains stuck in Hyper-V loop with PXE server response timeout [ALHELP-1819]

App Layering 2409

This release contains the following new features:

- **ELM appliance OS has upgraded**

The ELM appliance OS has transitioned from CentOS 7 to Oracle Linux 8. Direct upgrades from previous releases to version 2409 are not supported. For more information, see the [upgrade documentation](#).

- **ELM migration feature**

A new migration feature allows you to copy or migrate resources between ELMs, even across different hypervisors. For details, see the [migration feature documentation](#).

Fixes

- You may experience issues with versions prior to 24.3.3 while updating the Chrome layer, which may fail with an error message. This problem persists across different OS layers, including Windows 10 and Windows Server 2016. [ALHELP-1766]
- FSLogix Profile Agent may encounter issues with App Layering versions 24.3 and earlier, causing incomplete logins and missing Start Menu items on Windows 11 and other operating systems. [UNI-91138]
- You may experience issues with Teams failing to launch, after upgrading the ELM. As a workaround, you can install Microsoft Teams in the OS layer. This ensures that you will have access to Teams as an NP VDI user, allowing you to perform monthly updates alongside your regular OS updates. [ALHELP-1778]
- When uninstalling the Cortex XDR agent, the packaging machine may experience a hang. This issue is related to a symbolic link in the registry, which causes the hang when deleted. [UNI-91073]
- After upgrading to Citrix App Layering 2403.2, you may encounter intermittent failures when publishing image templates. [ALHELP-1777]
- Unirsd may experience double initialization in full user-layer images. This double initialization can cause the user's registry changes to merge into itself, resulting in delays during the login process. [UNI-91030]
- You may notice that all connectors are missing after upgrading from version 2304 to 24.3.1.13, from the ELM web console. This issue occurs because the `layerDiskCacheSize` is set to null. [UNI-90857]
- When you install an application in an App Layer, it may not create necessary registry keys. Although the installation completes, you might see an error stating that the launching program does not exist. As a workaround, you can install it in an OS layer instead of an App Layer. This method allows the correct registry keys to be created. [ALHELP-1746]
- You may experience performance issues with VS Code extensions when using Elastic layers. Users with multiple Elastic layers attached can see significant delays, causing VS Code to take much longer to start and run extensions. [ALHELP-1774]
- You can only update Microsoft Edge in an OS revision now. If you previously used an application layer to deploy Edge, that layer can't be used in the image anymore, as Edge updates will be ignored. Make sure to keep Edge up to date when updating Windows versions. [UNI-90502]

- When creating an OS layer using the ImportOsLayer.ps1 script, you may encounter a “400 Bad Request” error if any user input (including username and password) contains XML special characters such as &, <, >, ‘, or “. As a workaround, you can manually escape special characters as you enter them:
 - Replace & with &
 - Replace < with <
 - Replace > with >
 - Replace ‘ with '
 - Replace “ with "[UNI-87311]

App Layering 2403

This release contains the following new features:

- Citrix Hypervisor™ rebranded to XenServer**

In line with our latest rebranding strategy, we've updated all instances of Citrix Hypervisor to XenServer.

- OS Machine tools update**

The OS Machine tools are now updated to include automatic cleanups of empty INF folders that can cause performance counter rebuild failures. You must create a new revision of your OS layer and apply the OS Machine Tools to the layer.

Note: Don't run `setup_x64.exe` on an OS revision. The application is only needed before importing an OS into the elm, and running it on an OS revision might cause unknown issues.

- Layer creation tasks display more information about the packaging machine**

The action required state of layer creation tasks now include more information about the packaging machine than just the VM name when using offload compositing. For most connector types, the additional information is the machine's IP address.

- For Nutanix AHV connector types, a link to the packaging machine in the Nutanix Prism UI is displayed.
- Azure Deployment connector types allow custom information to be provided through the ARM templates. For more information, see [Machine Output](#).

- New Connector - Windows File Share**

The **Connectors** UI now includes **Windows File Share**. This connector allows the user to publish a disk to a file share, leveraging an existing offload-enabled connector configuration.

For more information, see [Windows File Share](#).

- **User Personalization Layer - User Layer/UPL space reclamation**

Previously, Windows used the unused blocks on a disk first, requiring a large amount of space on the backend file server. By using the new VHDX space reclamation process, the user layer VHDX files are automatically optimized every time the user logs off.

For more information, see [User Layer/UPL space reclamation](#).

Fixes

- The Builtin AppV apps are not working with the PVS EL-enabled images. [ALHELP-1717]
- The PVS Connector publish fails if the share path has a trailing backslash. [UNI-89881]
- When creating an OS layer from a XenServer gold VM with custom BIOS strings, the `ImportOsLayer.ps1` script fails to determine the hypervisor type and fails with the following error:
`Offload Compositing is not supported on this hypervisor yet`
The hypervisor type can now be explicitly specified through the `-HypervisorType` parameter to avoid this issue.[UNI-90521]

App Layering 2312

This release contains the following new features:

- **New and improved Nutanix AHV connector.** When you create and manage the Nutanix AHV connector in App Layering, you can now experience the new UI with the modernized Hypervisor support. For more information, see [Nutanix AHV connector](#). This release also introduces **Offload Compositing** support for layers and image templates in Nutanix AHV. Modern virtualization technologies such as UEFI, vTPM, and Secure Boot are now available with **Offload Compositing**. For more information, see [About Offload Compositing](#).
- **New approach to override file exclusion.** You can now override file exclusions by updating the Gold Image tool to introduce and accumulate default exclusions for the App Layer file system. For more information, see [Default exclusions](#).

Fixes

- The `HKLM\Software\Wow6432Node\Citrix\PortICA\Policy\Session\PolicyInputValues` and `HKLM\SOFTWARE\Citrix\VirtualDesktopAgent\Policy\Session\PolicyInputValues` keys do not have their values or subkeys persisted between logins. [UNI-89899]

- When you publish an image with App Layering version 2306, you are unable to see the new image tools, like AppRulesCompare, on the image. [UNI-89936]
- A login or logout failure might occur when FSLogix uses Cloud Cache locations for profile or ODFC containers along with App Layering Elastic Layers. A pop-up screen displays login failures and a frozen logout screen displays logout failures. FSLogix log files located at `C:/ProgramData/FSLogix/Logs/Profile/Profile-YYYYMMDD.log` also display the following error message for login failures: "`ERROR:00000005 Access is denied`" errors at various points during the login attempt.

App Layering 2309

This release contains the following new features:

- **Option to keep the MCS deployment machine up and running.** If you want to customize your deployed images from the ELM when the booted machine is being prepared for deployment to MCS, you must update your machine tools in the OS revision to revision 2309 or beyond. For more information, see [Create Platform Layer](#).
- **Support for VHDX format for user layers.** User layers are now created using the VHDX format instead of the VHD format, as in the previous releases. You do not need to convert existing user layers with the VHD format to VHDX. However, if you want to convert the user layers manually, the VHDX format of the layer is used if they both still happen to be in the same folder. If you want to create new user layers using VHD format, you can turn off the feature. For more information, see [Enable User layer](#).
- **OS layer switching option for elastic layers is now enabled by default in the App Layering UI.** You no longer have to manually enable the new OS layer switching option. For more information, see [OS layer switching for elastic layers](#).
- **New image template option - Defragment disk when publishing the image template.** You can now use the new image template option which optimizes the published image created from the App Layering appliance. This option is only available when you select an Offload Compositing connector. For more information, see [Create an image template from scratch](#).
- **New and improved Citrix Hypervisor.** When you create and manage the Citrix Hypervisor connector in App Layering, you can now experience the new UI with the modernized Hypervisor support. For more information, see [Citrix Hypervisor](#). This release also introduces Offload Compositing support for layers and image templates in the Citrix Hypervisor. Modern virtualization technologies such as UEFI, vTPM, and Secure Boot are now available with Offload Compositing. For more information, see [About Offload Compositing](#).

Fixes

- There is a memory leak on images that have elastic layering enabled and Sentinel One installed. [ALHELP-1708]
- Publishing images with many layers can fail due to database size constraints.
- The LSASS process is hung on some files and causes the boot process to fail. [ALHELP-1710]
- VMware vSphere fails with invalid object types.
- Purge Failure Mode waits are causing performance problems. [UNI-89518]
- Session hosts are getting locked on the first login. [ALHELP-1722]
- AppV packages, when preinstalled in an application layer, might either fail to launch or launch but fail to operate correctly. [ALHELP-1717]

App Layering 2306

This release contains the following new features:

- **Role-based access control is now available in the new App Layering UI.** You can now assign roles to App Layering users, defining which features they can access. Any user assigned a role can log in to the management console. For more information, see [Assign roles](#).
- **Removed support for Windows 7 and 32 bit from the ELM.** All Windows 7 bit and 32-bit OS layers have been discontinued and unsupported for several versions of the ELM. If you are still using a Windows 7 bit or 32 bit unsupported OS, do not upgrade to this new version without first backing up your ELM and then contacting Support for options.
- **New App Layering tool - ScanWritableFiles.** ScanWritableFiles is an application where users can identify files and folders that exist only on the writable layer (such as the User-Layer). These file-system objects are therefore not located on any other layer/volume, such as the base image or any elastically assigned layer. Identifying file-system objects of this type can be helpful when:
 - Determining what objects can be safely deleted without affecting elastically-assigned applications
 - Identifying application components directly installed by the user, versus being provided by the image or any elastic layer
 - Identifying application components installed on a packaging machine

For more details, see <C:\Program Files\Unidesk\Tools\ScanWritableFiles\ReadMe.txt>.

- **New App Layering tool - AppRuleCompare.** AppRuleCompare analyzes potential file system and registry conflicts between App layers, including the Platform layer if present. When

run directly on an app-layered machine, AppRuleCompare “processes” the AppRule files associated with layers built into the base image, as well as those that have been elastically attached to the machine. For more details, see <C:\Program Files\Unidesk\Tools\AppRuleCompare\ReadMe.txt>.

Fixes

- On some Provisioning servers, the command to mount a file share with an App Layering agent greater than 2003 hangs indefinitely. [ALHELP-1593]
- When a file or directory is created or renamed, CFS also changes the short-name of the object, which causes unexpected directory notifications. [ALHELP-1682]
- An error in the generatePerfGenRecompileScript routine is causing the BIC to fail. [UNI-89255]
- The rename operation is failing due to a bug in our rename logic. [UNI-89256]
- A driver digital signature error is occurring on 2012R2 after switching to a new Citrix certification. [UNI-89258]
- The new vSphere connector UI is not listing all hosts. [UNI-89314]
- A newly attached virtual disk is not being detected by VDS. [ALHELP-1694]

App Layering 2304

This release contains the following new features:

- **Registration credential support for Azure Deployments.** You can now use registration credentials to configure the Azure Deployments connector, similar to what was available for the Legacy Azure connector. For more information, see [Azure Deployments](#).
- **You can now define directory paths that are not redirected to the user layer with the user exclusions policy in Citrix Studio.** Some things to note:
 - User exclusions do not override AlwaysOnBoot.
 - User exclusions apply to the full user layer and the user personalization layer (UPL), but not to the session host. The session host ignores user exclusions and adds the message to the user layer.
 - Logoff.txt now contains all active user exclusions.
- **Custom Active Directory (AD) attributes support.** You can now use custom AD attributes in a user layer path for user layers and user personalization layers (UPL). AD attributes must be enclosed in hashes (for example, #sAMAccountName#). For more information, see [Deploy user layers](#) and [User personalization layer](#).

- **New UI for the VMware vSphere connector.** The VMware vSphere connector now has a new UI. For more information, see [VMware vSphere](#).
- **Performance counters now work from any layer, including app layers and full user layers.** The gold image tools must be updated in the OS revision to ensure that the performance counters are correctly rebuilt when booting an image. The OS revision needs at least one rebuild of the performance counters. Open an administrator command window in the same OS revision where you're updating the gold image and run the commands `c:\windows\system32\lodctr /r` and `c:\windows\syswow64\lodctr \r`.
- **MSMQ can now be used from any layer included in the image.** The gold image tools must be updated in the OS revision to ensure that MSMQ is properly reset and started, when present. A new log file named `GenRandomQMD.log` is present to indicate the actions of `C:\Windows\setup\scripts\kmsdir\GenRandomQMD.ps1`. This allows MSMQ to start up properly regardless of where it was installed.

Fixes

- Missing Profile key causes a failed mounting of the user layer. [UNI-88890]
- Directory-notification code causing CPU usage spikes when disconnecting an elastic layer from a running machine. [ALHELP-1614]
- The `generate.pl` tool set the key `HKLM\System\WPA` from **Never Virtualize** to **Virtualize Always**, causing the Windows activation key to be lost. [ALHELP-1673]
- FSLogix running with a full user layer failed login on the third and subsequent logins, causing duplicate profiles and other issues. [ALHELP-1672]
- Deleting and recreating volatile keys causes the query handlers to block the caller from seeing the image. [ALHELP-1648]
- SessionHost file probes missing some files, causing conflicts and leading to non-working layers. [ALHELP-1653]
- Improved the performance of reverse-direction enumeration times. [ALHELP-1642]
- The performance counter scanning logic opens files in Read/Share Only Read mode and cannot be turned off, causing some conflicts with files that must be continually updated by the user environment. The performance counter scanning logic now opens files in Read/Share Deny None mode and can be turned off if needed. [ALHELP-1679]
- FSLogix search roaming sometimes does not work correctly on Elastic Layer-only images because of a new function, which is called before user login time before the layer hives are virtualized. [ALHELP-1669]
- The Microsoft MSMQ feature is not working when finalizing the image. [ALHELP-1641]

App Layering 2211

This release contains the following new features:

- **New Azure Deployments connectors.** Two new Azure connectors have been created and implemented into App Layering: Azure Deployments and Machine Creation for Azure Deployments. Both connectors also support Government environments (Azure Government and Machine Creation for Azure Government, respectively). The old Azure connector and Machine Creation for Azure connector (as well as their Government counterparts) are now deprecated, but still available for use for a limited time. For more information, see [Azure deployments](#).
- **Block platform layer finalize if WEM RSA key detected.** If the WEM RSA key exists on the platform layer, issues are caused with Workspace Environment Management (WEM). If the RSA key is detected, you now receive a message before finalizing the platform layer to remove the key. See [Create platform layer](#) for details.
- **Miss-cache mechanism improved.** The default value for entries in the miss-cache mechanism has been increased from 256 entries to 1024 entries. If you have increased the default size using a registry setting, you must remove your custom setting for this enhancement to work correctly.
- **OEM drivers persist in the user layer.** OEM drivers (such as printer drivers) can now be stored in the user layer and don't require re-installation after every user logon. Possible naming conflicts among OEM drivers between the OS image and user layer are resolved automatically. After upgrading, existing user layers are scanned for any unexpected OEM traces.
- **Windows 10 and 11 22H2 support.** You can now use Windows 10 and 11 22H2 as an operating system for layered images. For more information, see [System Requirements](#).

Fixes

- Fixed an issue caused by adding the GAC_MSIL file to the registry AOB list. The GAC_MSIL file was removed. [ALHELP-1612]
- A new ulayer setting (HKLM\Software\Unidesk\ULayer\BasicAADScrubEnabled [dword]) was added which allows you to disable Azure's removal of Windows 10 registry and file locations. By default, this setting is set to true. [UNI-87854]
- Fixed an issue where Purge Mode delays included files opened for read-only access, which caused the system to run slowly. [ALHELP-1621]
- Fixed an issue where the keys in `C:\ProgramData\Microsoft\Crypto\RSA\MachineKeys` were being removed when finalizing. The BIC now includes this folder. [ALHELP-1623]

- Fixed an issue where some local policies were being removed after an OS Machine Tools upgrade. The gposetup.cmd was doing additional initialization and overwriting the gpt.ini file with the defaultgpt.ini file. [ALHELP-1627]
- Fixed an issue causing Access Control Lists to be ordered incorrectly in newly created folders. The issue caused some entries in these folders to be ineffective. [ALHELP-1632]

App Layering 2208

This release contains the following new features:

- **Wildcard characters are now allowed to exclude a directory from a composited layer.** Only one directory can be wildcarded and only one * might be used in a single path. See [Exclude files from layers](#) for more details.

Fixes

- Selecting layers for export no longer fails to save after upgrading to 2206. [ALHELP-1605]
- Fixed issues with certain apps that occurred after upgrading to 2206. The `C:\Windows\Microsoft.NET\assembly\GAC_MSIL` directory is no longer excluded from app layers, and the `c:\windows\system32\wbem\repository` directory is only allowed on the platform layer. [UNI-87356]

App Layering 2206

This release contains the following new features:

- **Microsoft Silverlight removed.** App Layering no longer requires or supports MS Silverlight. See [System requirements](#) for details.
- **A few menu items have been moved from the User tab to the System tab in the new UI:**
 - Managing role access for Active Directory users and groups
 - Creating, deleting, and editing directory junctions in the Directory Service
- **More management features have been converted to the new user interface.** The following features are now available in the new interface:
 - Importing and exporting layers
 - **System** - User Layer storage locations

- **Fully implemented new user interface!** With the completion of the above, the new UI is fully implemented, and all documentation now reflects this. Many procedures and sections throughout the documentation have been updated with the new blade workflows, replacing the wizard dialogs from the old UI.
- **Custom firewall settings on user layer.** Administrators can now create a script to automatically run at logon in a user layer to set up firewall options and rules for a user. The script is [PrivilegedLogon.cmd](#) and can be found at `C:\Program Files\Unidesk\Uniservice\UserScripts`.

Fixes

- NVivo no longer crashes when elastically assigned with Full User Layer enabled. [ALHELP-1564]
- Folders can now be deleted in a published image after upgrading. [ALHELP-1582]

App Layering 2204

This release contains the following new features:

- **Support for Microsoft Office 2021.** Office 2021 can now be used with Citrix App Layering.
- **More management features have been converted to the new user interface.** You can access the new interface by way of a unique URL in your web browser. Using the IP address for the appliance, enter the following URL: https://<ip_address_of_new_vm>. The following features are now available in the new interface:
 - System Tab - Manage Appliance
 - System Tab - Settings and Configuration

Fixes

- Setting attributes of a directory no longer fails when located on a read-only volume (elastic app layer). [ALHELP-1500]

App Layering 2202

This release contains the following new features:

- **You can now force the deletion of master key files for app layers.** If you get an Elastic Fit warning for an app layer due to master key file changes, you can force the deletion of master key files by editing the registry setting `DeleteMasterKeys`. See [Deploy App layers as elastic layers](#) for details.

- **The VMware Horizon View connector has been deprecated.** If you enable or use this connector, be aware that it has been deprecated and will be removed in a future release.
- **More management features have been converted to the new user interface.** You can access the new interface by way of a unique URL in your web browser. Using the IP address for the appliance, enter the following URL: https://<ip_address_of_new_vm>.
 - The following features are now available in the new interface:
 - System Tab - Connectors
 - System Tab - Settings and Configuration
 - User Menu - Upgrade appliance

Fixes

- The default setting of the ‘ManageOpenForBackup’ flag has been changed from false to true as a workaround for ACLs being corrupted on C:\windows\syswow64 and C:\windows\system32 when installing an application. (ALHELP-1327)
- Search locations that are removed using the Indexing Options applet are now removed correctly. (ALHELP-1493)
- Setting the attributes of a directory no longer fails when located on a read-only volume (elastic app layer). (ALHELP-1500)
- Connections from VDAs in the process user layer are now evenly distributed across domain controllers. (ALHELP-1535)
- The user personalization layer now installs correctly in XenDesktop. (ALHELP-1545)
- Upgrading the Enterprise Layer Manager (ELM) from 21.06 to 21.12 no longer fails. (ALHELP-1559)
- The user’s session now times out appropriately in the new UI. (UNI-85868)

App Layering 2112

This release introduces a new App Layering management experience and support for additional versions of Windows.

- **Support for additional versions of Windows:** You can use the following versions of Windows as an OS Layer in App Layering 2112 and later.
 - Windows Server 2022
 - Windows 10 version 21H2
 - Windows 11, with the following caveat

Caveat:

To support Windows 11 as an OS Layer, you must upgrade the App Layering appliance to version 2112. That version provides you with the requisite updates to the Optimizer Script builder, the Unattend Script builder, the SetKMS, and the guest installer to set the OS type. The gold image tools used to set up the gold image must be at version 2112 or later.

- **New App Layering management experience.** This version of App Layering introduces Phase 1 of a new, enhanced management experience. Phase 1 of the new user interface will temporarily coexist alongside the existing user interface on the App Layering appliance. You can access each by way of a unique, distinct URL in your web browser. Note that we have not yet ported some administrative activities outside of image templates and layers to the new user interface. For those, you need to continue to use the original interface. When accessing the appliance from a modern web browser such as Chrome, Edge, or Firefox, the new management console appears. If you use Internet Explorer to access the appliance, the legacy management console appears.
 - Using the IP address for the appliance, enter the following URL in a compatible web browser:
`https://<ip_address_of_new_vm>`
 - Features available in the new user interface:
 - Template management
 - App, Platform, OS Layer management
 - Task management
 - Login
 - Elastic Layer user assignment
 - Browser support:
 - Microsoft Edge
 - Google Chrome
 - Mozilla Firefox

Fixes

- Office 365: Excel can become unresponsive on images created in version 2110 of the App Layering appliance. (ALHELP-1537)
- When using Zscaler, a driver error is reported when any elastic layer is attached to the VM. (ALHELP-1528)
- Office 365: User layers are not being attached to the session host. The issue occurs due to a file lock on PEUPTemplate.hive (ALHELP-1525)

- App Layering 2107: VDAs can experience a fatal exception, displaying a blue screen at random. (ALHELP-1436)

App Layering 2110

This release supports:

- VMware Cloud on:
 - Azure
 - Amazon Web Services (AWS)
- The Citrix Provisioning™ connector now lets you specify a hypervisor connector configuration to use for **Offload Compositing**. Selecting **Offload Compositing** in a Citrix Provisioning connector configuration enables support for VHDX disk format, UEFI firmware, and Secure boot.
- If Offload Compositing is enabled in a Citrix Provisioning connector configuration, you can add a custom description to the vDisk in the Publish Image wizard. For example, if you publish an image template using a Citrix Provisioning connector that has Offload Compositing enabled, you can add a comment in the Publish Layered Image wizard and the comment appears as the resulting vDisk's description. If Offload Compositing is not enabled or you leave the Comment field blank, the description defaults to "Layered Image". The maximum length for a vDisk description is 250 characters. If longer, it is truncated.
- (Advanced feature) You can now exclude specific files and folders from a composited layer to prevent files from persisting on a user's desktop. For example, you can exclude antivirus software files and folders that must not persist for a desktop from one login to the next.

Fixes

- An issue where images that include a NetApp layer become stuck at 100% CPU utilization has been fixed in cooperation with NetApp developers. (ALHELP-1508)
- When User layers are enabled and Windows is updated, the Windows **Start Menu** and **Search** features work properly. (ALHELP-1482)
- When you create an OS layer and a platform layer and install the Citrix Virtual Delivery Agent (VDA), the published image no longer results in a blue screen. (ALHELP-1485, ALHELP-1486)
- The issue with MediTech Expanse after upgrading the App Layering appliance has been fixed. (ALHELP-1494)
- The issue that prevented the successful setup and use of Dropbox on a user layer has been fixed. (ALHELP-1416)

- The issue where the Windows search index was corrupted when creating an OS layer or adding a version to it is fixed. (ALHELP-1433, ALHELP-1453)

Labs feature

Labs features are previews of potential functionality. While a feature is in Labs, do not use it in production. There is no guarantee that this feature will be included in the product, nor that it works the same way if it is.

- **You can assign app layers as elastic layers on images that use a different OS layer:** Elastic layer assignments normally require the App layer assigned uses the same OS layer that was used to create the App layer. You can try assigning an App layer as an [Elastic layer](#) on a layered image that uses a different OS layer.

Important:

Issues can result from running an elastic layer on a different OS layer than the one used to create it.

To use a Labs feature, [enable](#) it in **System** settings.

Upgrade path

For the latest fixes and features, including compatibility with other software packages that you use, we encourage you to stay current with the App Layering [upgrades](#).

You can upgrade from any App Layering release from 19.x to present.

App Layering 2107

This release includes the following improvements.

Fixes

- After upgrading from vSphere 6.7 to vSphere 7.0 Update 2 or later, you can now create layers and publish images with VSAN storage. (ALHELP-1410)
- After upgrading to Windows 10 1909 or 20H2, ClickOnce apps now work with existing user layer disks. (ALHELP-1425)
- An app layer that is assigned to a subset of users on a machine and contains certain Windows system files unique to the app no longer causes problems for users that are not assigned to the layer. (ALHELP-1427)

- On an image with EL running, you can select the remote admin share where OneDrive saves files without the machine failing with a blue screen. (ALHELP-1431)
- When Elastic Layering is enabled on an image, a script that installs an application on the image now completes as expected. (ALHELP-1432)

App Layering 2106

This release includes the following new feature:

- An updated version of the OS Machine Tools. We recommend that you update your OS layers with the new tools now, so that you can use any new features that require them in the future.

Fixes

- After you finalize a platform layer, the machine no longer fails with a blue screen. (ALHELP-1177)
- When Offload compositing is enabled and you [set the default size of the elastic layering volume](#), the writable partition size is updated for all published images. (UNI-76795)
- When logging in after adding elastic layers to an image, users no longer receive errors like the following (ALHELP-1445):

“Critical Error: Your Start menu isn’t working. We’ll try to fix it the next time you sign in.”

“Citrix App Layering - System Error: An unexpected system error occurred. Retry the operation or contact technical support.”
- Files on an elastic layer no longer disappear and reappear under certain conditions. (ALHELP-1405)
- When using FSLogix and OneDrive with elastic layers enabled, the VM can now access the user profile folder through the admin share (\\\PCName\CS\Users). (ALHELP-1386, ALHELP-1405, ALHELP-1431)
- FSLogix Profiles now work correctly when the app is installed on an app layer. (UNI-83092)

App Layering 2104

This release includes the following improvements.

Fixes

- Synchronization of Layer version repair data now runs faster. (UNI-82197, ALHELP-1385)

- The OfficeNoReReg.cmd script now updates existing values. (UNI-82088)
- When accessing a published image running Windows 10 1909, you no longer receive the message, “The User Profile Service failed the sign-in. User profile cannot be loaded”. (ALHELP-1307)
- After you install Crowdstrike, Edge launches on the first attempt. (ALHELP-1404)
- When an image is running both elastic layering and Crowdstrike, Chrome and Microsoft Teams now launch on the first attempt. (ALHELP-1392)
- If you install apps that have files with boot-level components and Crowdstrike tags them, the apps no longer fail after you add a version to the layer. (ALHELP-1397)

App Layering 2102

This release includes support for the following enhancements:

- App Layering now runs on Google Cloud! You can:
 - Install an App Layering appliance on Google Cloud.
 - Create connector configurations for Google Cloud and Machine creation for Google Cloud.
 - Create layers on Google Cloud.
 - Move layers from a different platform to Google Cloud, using the Export and Import feature.
 - Publish layered images on Google Cloud, or to Machine creation running on Google Cloud.

Fixes

- Apps published from an image template with elastic layering enabled now launch as expected. (ALHELP-1306, ALHELP-1315, UNI-81247)
- When booting a session host with UEFI and Citrix Provisioning, the session no longer fails with a blue screen on the target device. (UNI-80889)
- Compositing no longer fails because the CE runs out of available drive letters (ALHELP-1286, UNI-80179)
- When an app is installed on a packaging machine, MSIEXEC.EXE no longer removes permissions on syswow64 content. (ALHELP-1327, UNI-81548)
- WebEx no longer fails with Error 1407 when trying to uninstall it from a user layer. (ALHELP-1339, UNI-81434)
- Apps in the Windows Start menu are listed in the correct section rather than in a section with name **ms-resource:AppName** when user layers are enabled. (ALHELP-1323, UNI-81402)

- Logging into the App Layering management console no longer results in a system error. (ALHELP-1332, UNI-81391)
- VDAs no longer fail with a bluescreen on App Layering 20.11. (ALHELP-1337, UNI-82008)
- The VDA is no longer unresponsive when users log in on to the VDA with elastic Layering enabled. (ALHELP-1369, UNI-81777)
- Machine creation for vSphere connector no longer deletes files under UnideskCacheddisks on VMware. (ALHELP-1345, UNI-81662)
- When assigned a user layer, you no longer get an error when accessing the Policy tab in Studio. (ALHELP-1355, UNI-81749)

App Layering 2011

This release includes support for the following enhancements:

- We now support Nutanix version 5.18.
- You can use Windows 10, version 20H2 as an OS Layer in App Layering version 2011 and later. However, the following caveat applies.

Caveat:

If you upgrade the OS layer to Windows 10 20H2 from an earlier release, upgrade directly to Build 19042.630, or above. Upgrading to builds of Windows 10 20H2 released before 11/16/2020 can result in inconsistent image deployments. For example, if you publish images using a template with **Generalize Offline** selected, the published images might not work correctly.

Fixes

- **Windows 10, 2004 login times.** An issue where Windows 10 2004 sometimes took 2–4 minutes to start, and occasionally included black-screens for 30–45 seconds has been fixed. (UNI-80656)
- **kmssetup script not added to the startup scripts folder.** An issue that prevented the kmssetup.cmd script from being added to the startup scripts folder when unzipped has been fixed. (ALHELP-1279, UNI-80410)
- **HP UPD driver not available as a printer driver.** When you include an HP UPD driver in an app layer that becomes part of a published image, the driver is now available as a printer driver. The issue that prevented the driver from being listed has been fixed. (ALHELP-1278, UNI-80426)
- **Sessions hangs for minutes.** An issue that caused sessions to hang for minutes has been fixed. This issue appeared after upgrading to App Layering 2005. (ALHELP-1263, UNI-80262)

- **User cannot reconnect to their desktop in App Layering 2009.** An issue that caused random occasional blue screens when users tried to reconnect to their desktops has been fixed. (ALHELP-1317, UNI-81156)
- **When using the console on a 4K display, dialogue boxes open in top left corner.** The issue that caused dialogue boxes to be displayed off center has been fixed. (ALHELP-1309, UNI-78951, UNI-78952)

App Layering 2009

This release includes support for the following enhancements:

- You can now override the default repository path and layer size for user layers by configuring Citrix Studio Policies.
- The user interface for the App Layering management console has been updated with new Citrix branding images.

App Layering 2008

This release includes support for the following features and enhancements:

- **Windows 10, 2004 support:** Windows 10, 2004 is now supported as the OS for layered images. Using Windows 10, 2004 requires App Layering version 2008 or later. The required changes for this version of Windows 10 are not in previous App Layering releases.
- **Citrix Hypervisor 8.2 support** We now support Citrix Hypervisor version 8.2.

Important:

App Layering 20.8.3 includes an important update to the included drivers. If you are using a secure boot, you must upgrade to 20.8.3!

The 20.8.3 installation and upgrade packages are available for download. For new installations, download the Appliance Installation Package. For upgrades, download the Appliance Upgrade Package.

Fixes

- **VM in vSphere with elastic layers enabled no longer fails after being deployed to a Citrix Provisioning server.** An issue that had caused a VM in vSphere with elastic layers enabled to fail after being deployed to a Citrix Provisioning server was fixed. (ALHELP-1202, UNI-76300)
- An issue that had caused the AutoDesk 2020 installer to fail in an app layer has been fixed. (ALHELP-476)

- Users can now synchronize their OneDrive files without issues when **On Demand syncing** and **Elastic layers** are enabled (ALHELP-468)
- Chrome extensions no longer cause the error message, “FAILED_TO_COPY_EXTENSION_FILE_TO_TEMP_DIR” when elastic layers are enabled and an app layer is assigned. (ALHELP-419)

App Layering 2005

The App Layering 2005 release includes support for the following features and enhancements:

- **Security enhancements:** Security enhancements include cipher suite updates, third-party component upgrades, and runtime environment upgrades.
- **Secure boot support:** Guest drivers are certified via Microsoft’s **WHQL** program. You can use certified guest drivers in secure boot configurations. (UNI-74917)
- **Customer user layer path:** User layer files can now be stored in network shares whose locations are specified using custom paths. The custom paths can include environment variables. (UNI-78291)
- **Applications with services in user layer:** When an application installed in a user layer contains Windows services, such as Google Chrome, those services are started as expected when the user logs in. (UNI-77660)
- **2005 upgrade package:** The App Layering 2005 upgrade package is large enough that older appliances cannot download it automatically. If you are running version 1911 or older, download the package manually from the [downloads site](#). Also, we recommend running the upgrade from a management console in Secure HTTP (HTTPS). If you upgrade while in HTTP, messages do not display in the browser. If that happens, refresh the browser after 20 minutes. (Refreshing won’t cause issues in spite of the message that says not to refresh.) If the upgrade is still running, you get a “service unavailable error.” You can ignore that message and keep refreshing the browser every few minutes until the login page appears.

Fixes

- When using offload compositing with VMware vSphere, you can use layer names that start with brackets [], **as long as you follow the closing bracket with a space**. For example, the name **[OS] Windows 10** works, but the name **[OS]Windows 10** hangs in vSphere studio and times out after about 40 minutes. (UNI-78452)

App Layering 2003

This release includes the following:

- We now support the following hypervisor, provisioning, and connection broker software versions:
 - Citrix Hypervisor 8.1
 - Nutanix AOS 5.16
- When administrators configure a larger default user layer size, the disk will be automatically expanded the next time the user logs in.
- When creating your OS layer in MS Hyper-V or VMware vSphere, you can now import the OS image, by using a script in the OS Machine Tools. The script imports the OS image right from the virtual machine, instead of using the management console and connector configuration. The script uses the Offload Compositing feature, which speeds up the import, and allows you to use a wider variety of virtual machines, including UEFI-based machines.

Fixes

- The Guest Layering service (**ULayer**) has been modified to not depend upon the Server service running on the end users' virtual machines. (UNI-77242)
- When delivered as an elastic layer, the Artiva application no longer fails when users attempt to log in. (UNI-76487)
- An issue that had caused the `StartCCMEXEC.cmd` script to continually grow the log file, `StartCCMExec.txt`, even when `CCMExec` was not installed, has been fixed. To apply the fix, download the new OS Machine Tools folder for 20.3, and replace the `StartCCMEXEC.cmd` file with the new version. (UNI-77471)
- When users install fonts on a user layer or elastic layer, the fonts persist the next time they log in. (UNI-63839)
- When using a connector with **Offload compositing** enabled to publish an image, 8.3 file names are no longer incorrectly modified. (UNI-76961)
- When adding a version to an app layer, you no longer receive the error “The Operation Failed due to a missing File. VMDK was not found.”(UNI-77702)
- The issue with expiring JSON Web Tokens (JWT) when using the Offload Compositing feature in the connector configuration has been fixed. (UNI-76859)
- The Hyper-V connector now reports the disk file size rather than its logical size, which was quickly filling up the cache. (UNI-76692)
- Compatibility with Citrix Studio GPO policies when using images with user layers has been improved. (UNI-76918)

App Layering 2001

This release includes the following:

- We now support the following hypervisor, provisioning, and connection broker software versions:
 - Citrix Provisioning version 1912
 - Citrix Virtual Apps and Desktops™ version 7 1912
- The following Windows 10 versions are now supported as an OS Layer:
 - Windows 10, version 1909
 - Windows 10, Enterprise Virtual Desktop edition (available from Microsoft in Azure only)
- Our VMware vSphere connector now supports the VMware Paravirtual SCSI Controller. (UNI-75620)

Fixes

- In our App Layering appliance (ELM) deployment script for Azure, we have extended the expiration dates of signed Azure URLs pointing to standard repository disks. Make sure that you update your Azure deployment scripts from this release accordingly.
- VDA installations no longer fail on a packaging machine. (UNI-76299)
- When restarted, a layered image with a user layer no longer drops scheduler tasks. (UNI-77084)

App Layering 1911

This release includes the following:

- We now support the following hypervisor, provisioning, and connection broker software versions:
 - XenApp and XenDesktop 7.15 LTSR CU5
- VMware vSphere:
 - Use the new Offload compositing option to significantly reduce the time it takes to package layers and publish images.
 - Select thin provisioned disks and UEFI (no secure boot with elastic layers and user layers) options.
 - Use the vSphere connector to package layers and publish images to VMware Cloud on AWS.

Fixes

- An error when finalizing OS layers with MBR partitions after a major Windows 10 upgrade has been fixed. (UNI-76210)

- Compatibility with the Microsoft System Center Configuration Manager (SCCM) has been improved. (UNI-76198, UNI-76126, UNI-76129)
- Users no longer lose their connection to a session host where FSLogix is running and elastic layering is enabled. (UNI-73793)
- Rapid Reader 8.3 now starts as expected after being installed on a packaging machine. (UNI-76316)
- IntelliJ IDEA Ultimate no longer freezes the desktop when the app is started. (UNI-76075)
- Users no longer get a blue screen when they open SQL Management Studio and attempt to save a Query to OneDrive or a SharePoint folder. (UNI-76427)

App Layering 1910

This release includes the following:

- We support the following new versions of hypervisor, provisioning, and connection broker software:
 - Citrix Virtual Apps and Desktops version 7 1909
 - Citrix Provisioning version 1909
 - Nutanix Acropolis Hypervisor (AHV) version 5.11
 - VMware Horizon View 7.9
- On the Hyper-V platform, you can package layers and publish images using **Offload Compositing** to dramatically improve performance and compatibility. Gen2 VMs and VHDX disk formats are now supported.
- The App Layering Agent has been updated to support the new Offload Compositing feature. If you use the agent in your deployment, we strongly recommend that you install the agent update on any servers where it is installed.
- Windows Search index performance is improved when using full User Layers.
- For administrators publishing layered images, this version adds validation and warnings to give guidance on what disk size to use based on selected app layers.

Fixes

- A PVS Support article that called for removal of the CDF Driver Registry entry has been updated to fix issues troubleshooting Profile Management. (UNI-75720)
- The issue where Defender didn't update correctly has been fixed. (UNI-74918)
- Windows Search Service starts as expected on a packaging machine when using a caching connector. (UNI-75915)

- Windows Defender now starts successfully on an app layer created from a Windows 1809 OS layer. (UNI-74997)
- Users can now sync their OneDrive files when On Demand syncing and elastic layers are enabled. (UNI-74618)
- OneDrive's on-demand feature now streams files correctly when elastic layering is enabled. (UNI-73121)
- User layer repairs now complete successfully when the minimum recommended permissions are set on the Users\ path. (UNI-75552)
- When installing a Chrome extension with user layers enabled and an app layer assigned as an elastic layer, you no longer receive the error `Can not install package: FAILED_TO_COPY_EXTENSION_FILE_TO_TEMP_DIRECTORY` (UNI-75568)
- Windows Search index performance has been improved when using full user layers. (UNI-73049)
- This release adds validation and warnings about what disk size that you can use. The guidance is based on your configuration and the app layers selected. (UNI-54390)

App Layering 1908

This release includes support for these new features and improvements.

- **Windows versions:** Windows 10 version 1903 is now supported as an OS Layer.
- **Connection brokers:** We now support the following connection broker software version:
 - VMware Horizon View 7.9
- **User layer repair:** You can now manually repair user layers so that all files and registry settings coming from a specific set of app layers can be made visible again.

Fixes

- Hyper-V connector PowerShell sessions no longer expire when a file copy takes longer than 30 minutes. (UNI-74283, UNI-74292)
- The issue where System and Display properties did not open correctly when user layers were enabled has been fixed. (UNI-74547)
- Removing store apps from a Windows 10 OS layer no longer causes features, like **Settings** and **Display Properties**, to malfunction. (UNI-74852)
- The Windows Defender update KB4052623, now successfully installs and works for user layer-enabled users. (UNI-74942)

Contact Citrix

We welcome your feedback about this release.

- Use our online Forum to speak directly with Citrix employees and other organizations deploying Citrix App Layering.
- For product issues and questions, open a Support Case.

We look forward to hearing what you think about App Layering.

Related information

- [Support Knowledge Center](#)
- [Enterprise Architect TechTalks: Citrix App Layering FAQ](#)
- [Citrix App Layering Discussions](#)
- [App Layering Recipes](#)

Deprecation

September 6, 2025

The announcements in this article are advanced notice of the Citrix App Layering features that are being phased out, so that you can make timely business decisions. Announcements can change in subsequent releases and might not include every deprecated feature or functionality. For details about product lifecycle support, see the [Product Lifecycle Support Policy](#) article.

Deprecations and removals

The following list shows the Citrix App Layering™ features that are deprecated or removed.

Deprecated items are not removed immediately. Citrix® continues to support a deprecated item until removing it in a future release.

Removed items are either removed, or are no longer supported, in Citrix App Layering.

Item	Description	Deprecation announced	Alternative
Disabling Offload Compositing	Deprecated support for disabling offload compositing in connector settings	March 2025	None. Users should use offload compositing for connector setup.
Office 365/Office 365 Session User Layer	Deprecated support for Office 365 as a selection for the user layer type.	August 2023	Use Citrix Profile Management for advanced Office 365 use cases.
Support for VMware vSphere 6.7	Removed support for VMware vSphere 6.7.	Use higher versions of VMware vSphere .	
Azure connector	Deprecated support for the legacy Azure connector.	November 2022	Use the Azure Deployments connector , released with App Layering 2211.
VMware Horizon View connector	Deprecated support for the VMware Horizon View connector	March 2022	December 2023

Known issues

January 8, 2026

App Layering upgrade

App Layering 2511

- For the Google Cloud Platform connector there is currently an issue with packaging layers and publishing layered images. The task fails after waiting for the Compositing Engine to start, with an error message that contains *The resource 'citrix-ce/' of type 'Guest Attribute' was not found. [404]*. This issue will be resolved in a future update or release. [UNI-92996]
- Image Portability Service feature currently does not support Azure images configured with NVMe disk controller. Support is planned for a future release.

App Layering 2409

- You may encounter an incorrect message after migration completes with failed agents in Citrix App Layering. The message states, “The migration has completed successfully, but we were unable to register the following App Layering agents with the target appliance: System.Collections.Generic.List‘1[System.String].”[UNI-91339]
- After migration to the target Elastic Layer Manager (ELM), you will not be prompted to set the root and console administrator passwords. If these passwords were not set prior to migration, they will default to the factory settings.

As a workaround:

If the passwords were not set before migration, log into the target ELM console as root and change the passwords manually using the following commands:

- Change the root password:

```
passwd
```

You will be prompted to set a password for root.

- Change the administrator password:

```
passwd administrator
```

You will be prompted to set a password for administrator, use a lower-case ‘a’.[UNI-91472]

- When you connect to a new ELM using a non-root user, you may receive an unclear error message that states, “Cannot convert undefined or null to object.” You should expect a specific message indicating that the user is incorrect. [UNI-91330]
- You should avoid using special characters such as &, <, and > in your default administrator password on the ELM, as this may cause OS layer creation failures via the Import script. [UNI-87311]
- After migration, the management console Administrator password on the target Elastic Layer Manager (ELM) will be overwritten with the password from the source ELM. As a workaround, you can log into the target ELM as a user with Administrator privileges and change the default Administrator password. [UNI-91451]
- After migration, the [Images](#), [Connectors](#), [Directory Services](#), and [User Layer Storage](#) blades on the Target ELM do not refresh automatically. You may notice that these blades appear empty, even though the [Layers](#) blade shows data. To display the correct information, manually refresh your browser or tab.[UNI-91374]
- After performing a Copy migration, you are unable to launch elastically assigned app layers on images published by the target Elastic Layer Manager (ELM) when a new network share is created.

Cause

The new network share does not contain any app layers that were previously assigned and synced to the original network share on the source ELM.

Workarounds

Workaround 1:

Execute the following command:

```
mono /usr/local/unitools/MAService.exe -fire SyncRevisionsToShareEvent
```

Workaround 2:

Manually copy all app layers from the source ELM's network share to the new network share associated with the target ELM.

Workaround 3:

Unassign and then reassign the app layers to you. This action will trigger a sync of the layers to the new network share.[UNI-91377]

App Layering 2312

Upgrading to App Layering 2312 might fail, returning the error “A failure occurred while upgrading the appliance. Try the upgrade again after reverting from a clean snapshot.” We are aware of an issue with some appliances that were upgraded from previous versions. In this scenario, the upgrade might fail due to the presence of legacy upgrader components. We are working on a patched build for App Layering 2306 that contains a permanent fix for this issue.

To work around this issue, it is necessary to remove the legacy module by making the following changes in the Layering appliance:

1. Log in to the appliance console using the root login and password.
2. Run the command “`yum remove mod_http2`”. It might take several minutes for the command to run.
3. After the command completes, log back into the appliance using a web browser and rerun the upgrade.
4. It might take 30 minutes or more for the upgrade to complete, but it must now succeed.

App Layering 2005

- The App Layering 2005 upgrade package is large enough that older appliances cannot download it automatically. If you are running version 2001 or older, download the package manually from the [downloads site](#).

- We recommend running the upgrade from a management console in Secure HTTP (HTTPS). If you upgrade while in HTTP, messages do not display in the browser. If that happens, refresh the browser after 20 minutes. (Refreshing won't cause issues despite the message that says not to refresh.) If the upgrade is still running, you get a "service unavailable error." It is safe to ignore the message and keep refreshing the browser every few minutes until the login page appears.
- If you have two upgrade packages with the same name in different folders in Network File Share, selecting one of those packages causes both packages to be selected. If both packages have the same version, the upgrade succeeds. If they have different versions, the system chooses the lower version number. This occurs with Enterprise Layer Manager (ELM) version 22.2.

App Layering appliance and management console

- In the new UI, platform types cannot be edited. **Edit Platform Types** still functions correctly in the old UI. (UNI-86856)
- When installing the App Layering appliance, you *must* use the default CPU setting of **4 CPUs**.
- If you use roles in a complex Active Directory environment and logins are slow, assign all roles to explicit users rather than to groups.

App Layering documentation links

The documentation links in the management console open as a blank page in Internet Explorer 11. To get around this issue, paste the link into another browser. The documentation displays correctly.

App Layering agent

By default, the Citrix App Layering™ Agent runs under the **Local System** account on the Hyper-V server. If you change the account to anything other than **Local System**, the agent cannot transfer disks to and from the appliance.

App Layering OS Machine Tools

- (Release 19.5 only) After upgrading to release 19.5 (or later) from 19.3 (or earlier), be sure to update KMS Office Activation to use Office 2019. When preparing your OS image for layering, download and run the new App Layering OS Machine Tools.
- (Release 19.1 only) When preparing your OS image for layering, ensure that your KMS Office Activation is triggered at desktop startup. For this release only, download and run the App Layering OS Machine Tools from **Release 18.12**.

Microsoft Teams 2.x

- For details about the Microsoft Teams 2.x, see [New Teams Deployment Guidance for App Layering or User Personalization Layer \(UPL\)](#).

Elastic Layering

- Microsoft Office cannot be *elastically* layered due to the way its licenses are integrated with the Windows Store. The Office app layer must be included in the Layered image.
- When you enable an image with elastic layering, users might be able to view files and directories from other sessions in Windows Explorer. Directories explored in the other session might create folders visible to all sessions that have permission to browse that directory.
- If you use elastic layer assignments with Windows Server 2008 or Windows 7, create your file share with a sector size of 512. For details about this issue and related operating system updates, see the following:
 - [Microsoft support policy for 4K sector hard drives in Windows](#)
 - [Update that improves the compatibility of Win 7 and Win Server 2008 R2 with Advanced Format Disks \(UNI-48984\)](#)

User layers

- **Signing on after upgrade starts the Windows First Sign-in screens:** When you sign in after upgrading to 4.10 or later, the usual *Windows First Sign-in* brings the user layer up-to-date with the OS version. The process preserves user layer files.

Windows 10 support

- **Windows 10, version 20H2 upgrades.** If you upgrade the OS layer to Windows 10 20H2 from an earlier release, upgrade directly to Build 19042.630, or later. Upgrading from builds of Windows 10 20H2 released before 11/16/2020 can result in inconsistent image deployments. For example, if you publish images using a template with **Generalize Offline** selected, the published images might not work correctly.
- **Upgrading requires extra steps when going to a new Windows 10 major release:** During the upgrade, Windows 10 can create a recovery volume on the same disk as the OS layer version. Always delete this volume before you finalize the OS layer version. Otherwise, the recovery volume can cause desktops to fail to start correctly. For more information, see Issue 9 under [Windows 10 v2004, 20H2, 21H1 & 21H2 - Citrix Known Issues](#).

- If you have generated and applied the App Layering `Optimizations.cmd` script to a Windows 10 1909 OS layer, the **Search** option on the **Start** menu might not work as expected. To avoid this issue, add a version to the OS layer and run the program `c:\windows\setup\scripts\Optimize.hta`. To build a new `Optimizations.cmd` script to apply to the new layer version, deselect **Disable tablet input service (Section 6, Option M)** and select **Save File**. Before finalizing the OS layer, run the command `Powershell Set-Service TabletInputService -startuptype manual` to undo the effect of any previous `Optimizations.cmd` that might have disabled the service.

Connectors

- When using the **Azure Deployments** connector, if you delete all templates and edit the connector to choose a new template, a deleted template version appears. Then when you click **Save**, an error appears. As a workaround, reselect a valid template version before clicking **Save**. (UNI-88412)
- When using the Windows mini-boot disk option, you can specify up to four Prerequisite layers for any given App layer. If an app requires more than four other applications to be present during installation, install multiple apps in one layer. (UNI-69524)
- When creating a layer (app, OS, or platform) on Windows 7 64-bit, if you select **Offload Compositing** in the connector configuration, you can have issues adding a version to the layer. An error occurs and the packaging machine is not created. When Offload Compositing is *not* selected in the connector configuration, the packaging machine is created. (UNI-82545)
- Receiving 503 error, “Server Busy issues” from Azure. If you consistently receive this error, follow the steps in [CTX310868](#). This issue happens sporadically, and a solution for it is being tested. (ALHELP-1383)
- If you attempt to enter a name for a connector, and that name already exists, you receive a generic error message from the system, rather than the correct error message. The issue occurs for the **Azure Deployments**, **Machine Creation for Azure Deployments**, **VMware vSphere**, and **Machine Creation for vSphere** connectors. (UNI-89218)

Citrix Provisioning™

- When you create an image template, the target device hardware settings must match the Windows operating system and platform layer settings. Ensure that the hardware settings on the target device match the operating system and platform layer hardware settings, especially the number of CPUs. If the settings don’t match, you can get a “restart required” message when you start the published image. (UNI-50799)
- If you use Provisioning Services, you must disable IPv6 in the OS layer and *not* in the Platform layer. (UNI-53600)

- When importing VHDX files published from App Layering to the PVS disk store, you sometimes receive an invalid disk message. Eliminate the error by changing the period (.) characters in the published file name's date and time. A valid file name contains only one period for the VHDX file name extension. (UNI-75902)
- When **Offload Compositing** is selected in the connector configuration:
 - The path for the Citrix Provisioning Store fails to validate if it contains spaces. Replace the spaces with **%20** to make the name valid. (UNI-84868)
 - Publish jobs fail if the File Share Path ends with a backslash (\). (UNI-85045)
 - Publish jobs fail with a ComponentActivator error message if the Domain User does not have Read and Write permission to the File Share Path. (UNI-85020)
- When you select a Hyper-V connector for the offload compositing configuration, and the OS Layer is Gen 2, you must create another version of the OS layer, and then create the platform layer from that OS version. Otherwise, the target does not boot. (UNI-85044)
- When setting the Compositing File Share Path for the connector configuration, connectivity between the compositing engine and the Citrix Provisioning Store is not verified. If the store path doesn't map to the File Share Path, you receive an error similar to:
 - Error: "An unexpected system error occurred. Retry the operation or contact technical support. Exception Message: Response status code does not indicate success: 404 (Not Found). [Exception Details] (UNI-85045), (UNI-85020)

XenServer®

- When you prepare your operating system image for use in your XenServer, you must open port 5900 to allow console access. (UNI-50846)
- Always set the Citrix App Layering connector configuration to point to the master node. (UNI-52454)
- Prerequisite layers:
 - If a machine hangs at boot and a prerequisite layer is selected, one of the layer disks is probably not attached. Ensure that the Citrix Guest Tools are included in either the OS or platform layers.
 - If you are using Prerequisite layers to create either the OS or platform layer, Citrix Tools must be present. Without the tools, the packaging machine fails and you receive a blue screen. XenServer isn't able to see any devices attached *after* the DVD drive. The DVD drive is always in the third slot. (UNI-67741)

Citrix Virtual Apps and Desktops™ (CVAD)

- When updating Citrix Virtual Apps and Desktops to version 7.15 CU4, you must first install .NET Framework 4.7.1 on a new version of your OS layer, rather than on the platform layer. Installing .NET Framework 4.7.1 on the OS layer ensures that all app layers, platform layers, and images work correctly. The latest Windows updates already include .NET Framework 4.7.1 as part of the updates. (ALHELP-588, UNI-75108)

Google Cloud

- When importing layers from another platform, you must add a version to the OS layer, and switch to using the new layer version from then on. Otherwise, packaging machines and published images are likely to fail with a blue screen.
- Google Cloud Connector Configuration. “Check Credentials” verifies the Service Account User role. If the service account specified in the Google Cloud connector configuration **Service Account JSON key file** is different than the service account associated with your selected **Instance Template**, your service account in the configuration must have the **Service Account User role**. If it does not, then you receive an error when deploying a machine using that connector configuration. (UNI-82082)

Nutanix Acropolis

- The following message during app layer creation indicates that the app layer settings specify a platform layer. Do not use platform layers with the app layer’s performance-enhancing caching feature. (UNI-67742)

Create Application Version Wizard - Confirm and Complete

Verify the version details are correct and click Add Version to create a new version for the App Layer



By selecting a platform layer packaging performance will be degraded

VMware vSphere

- When creating the OS layer using the Create OS Layer Wizard, Unified Extensible Firmware Interface (UEFI) virtual machines are listed. You cannot, however, create UEFI machines using the wizard. Instead, use the new `ImportOsLayer.ps1` script to **import the OS** onto the new OS layer machine.

- When using a vSphere connector configuration with VMware Cloud and a vSAN 7.0 Update 2 (or later) datastore, **Offload Compositing** must be selected. (UNI-85216)
- When using the new VMware vSphere connector in AL 23.4, if you select an opaque network when creating or editing the connector, the connector fails. As a workaround, select a non-opaque network. Existing connector configurations using an opaque network created before AL 23.4 continue to function normally. (UNI-89439)

Microsoft Azure

- App Layering does not support Azure File storage. For storage in Azure, create an SMB file share or a network file share. (UNI-42272)
- Managed disks are only supported for OS imports. Packaging app layers and publishing images only produce unmanaged disks. When creating a virtual machine in Azure, select **No managed disks**.

Microsoft Hyper-V

- When selecting a Hyper-V connector configuration for Offload compositing and your OS layer is Gen 2, the layer must have at least one version besides the original. Also, the platform layer where the Citrix Provisioning target device software is installed must be created using the new OS layer version. (UNI-85044)
- When a Hyper-V connector configuration is set for Offload Compositing with Gen 2 (UEFI) and VHDX, choosing VHD as the Disk Format in the Citrix Provisioning connector configuration is allowed, but this configuration is not supported.
- When you configure Elastic Layering in Hyper-V, you must use *unmanaged* RDS pools (UNI-53545)
- When creating an app layer, if a platform layer is specified in the app layer settings, you receive an error. Do not use platform layers with App Layering's caching feature. (UNI-71868, UNI-67743)

Create Application Version Wizard - **Confirm and Complete**

Verify the version details are correct and click Add Version to create a new version for the App Layer



By selecting a platform layer packaging performance will be degraded

- Creating an OS layer on Hyper-V Server 2019 can result in this error:

‘Failed to create VHD. Make sure that there is enough space on the share specified in the connector configuration.’

This error is due to an issue with the Microsoft PowerShell New-VHD cmdlet. We are keeping our eye out for a fix from Microsoft. In the meantime, use the following workaround for this error:

1. Make sure that the Gold VM has no checkpoints.
2. Make sure the Gold VMs disk is in the same directory path that is configured in the connector config. Example:
Local path is D:\Brock
Gold VM disk is stored in D:\Brock\Win10Gold\Win10GoldDisk.vhdx

- App Layering fails to create an app/platform layer if the path to storage in the Hyper-V connector configuration contains a backtick (`). For example:

```
1 mystoragename`
```

Printing

With App Layering images configured for Full User Layer, you can install your own printer devices directly. However, when you log out and log in, **Printers & Scanners** no longer show the self-installed printer devices.

As a workaround, you can access or select the printers from within the applications.

Network

Applications that include network components may not operate correctly when installed in an App Layer and then included in a published image. Applications of this type need to be installed into either the OS Layer or the Platform Layer to ensure a proper merge of the network-related registry information in the image.

Related information

- [Support Knowledge Center](#)
- [Enterprise Architect TechTalks: Citrix App Layering FAQ](#)
- [Citrix App Layering Discussions](#)
- [Application Layering Recipes](#)

System requirements

January 6, 2026

The App Layering virtual appliance runs on the supported hypervisors listed in this article. The appliance is where you deploy the Enterprise Layer Manager (ELM) during installation.

You can create layers to use on virtually any hypervisor or provisioning software. For the best user experience, publish images to the supported platforms.

Hypervisors for the appliance

The Enterprise Layer Manager (ELM) runs on a virtual appliance deployed to a supported hypervisor.

- XenServer, versions 6.5, 7.0–7.6, 8.0, 8.1, 8.2, 8.4
- Azure Resource Manager
- Google Cloud
- Microsoft Hyper-V running on Windows Server 2016, Windows Server 2019, Windows Server 2022, or Windows Server 2025.
 - Generation 2 Virtual machines are supported when you use the Offload compositing feature.
- Nutanix AOS 6.8, 6.10, 7.3
- vSphere vCenter, version 7.0 or 8.0 (including subsequent updates)
 - Generation 2 Virtual machines are supported on vSphere vCenter version 7.0 or 8.0 (and subsequent updates) when you use the Offload compositing feature

Network file share protocol

- Server Message Block (SMB)

Network Connection

- Citrix® recommends a 10 Gbps connection between the appliance and the file share.

Directory service

- Microsoft Active Directory

Internet browser for management console

The management console supports the following web browsers:

- Edge version 94 or later (version 102 officially tested)
- Chrome (version 90 and 102 officially tested)
- Firefox (version 45 - 52.9 officially tested)

Operating system for layered images

You can layer the following versions of the Windows operating system. Windows Store apps work on all supported Windows versions.

- Windows Server operating systems - The following session host versions are supported:
 - Windows Server 2025, 64-bit (Standard and Datacenter Editions)
 - Windows Server 2022, 64-bit (Standard and Datacenter Editions)
 - Windows Server 2019, 64-bit (Standard and Datacenter Editions)
 - Windows Server 2016, 64-bit (Standard and Datacenter Editions)
- Desktop operating systems - The following desktop versions are supported:
 - Windows 11, version 25H2, supported in App Layering version 2511 and later.
 - Windows 11, version 24H2, supported in App Layering version 2409 and later.
 - Windows 11, version 23H2, supported in App Layering version 2312 and later.
 - Windows 11, version 22H2, supported in App Layering version 2211 and later.
 - Windows 10, version 22H2, supported in App Layering version 2211 and later.
 - Windows 11, version 21H2. Supported in App Layering version 2112 and later.
 - Windows 10, version 21H2. Supported in App Layering version 2110 and later.
 - Windows 10, version 21H1. Supported in App Layering version 2107 and later.
 - Windows 10, version 20H2. Supported in App Layering version 2011 and later.

Caveat:

If you upgrade the OS layer to Windows 10 20H2 from an earlier release, upgrade directly to Build 19042.630, or later. Upgrading to builds of Windows 10 20H2 released before 11/16/2020 can result in inconsistent image deployments.

- Windows 10, 64-bit, versions 2004, 1909, 1903, 1809, 1803, 1709, 1703, and 1607 (Education and Enterprise Editions)

- Windows 10, 64-bit, versions 2004, 1909, 1903, 1809, 1803 (Professional edition)

App Layering supports single-byte language packs for the base US English Windows operating system. It supports multi-byte language packs when the OS layer is deployed on supported versions of:

- XenServer®
- VMware vSphere
- Microsoft Hyper-V

User layers

Full User layers are supported on the following platforms:

- Operating systems:
 - Windows 10, 64-bit
 - Windows 11, 64-bit (only if deployed to a platform enabled for offload compositing)
- Publishing platforms:
 - Citrix Virtual Desktops™

Layered images

Layered Images are bootable images composited from Layers. Each Layered Image contains an OS Layer, a Platform Layer, and any number of App Layers. You can publish layered images to these platforms:

- Machine Creation for XenServer (formerly Citrix MCS for XenServer)
- Machine Creation for Azure and Azure Government
- Machine Creation for vSphere
- Machine Creation for Nutanix AHV
- Citrix Provisioning, versions 2203, 2106, 2012, 2009, 2006, 2003, 1912, 1909, 7.15 CU6, 1912 CU3, and Provisioning Service (PVS), versions 7.15 LTSR (any version)

Notes:

- Citrix recommends network speeds of 10 Gbps to the Provisioning store.
- Provisioning Service (PVS) versions 7.1, 7.6–7.9, 7.11–7.18, and 1808 are no longer supported.

The appliance and connectors run in the following environments:

- Citrix Virtual Apps and Desktops™, versions 7: 1808, 1811, 1903, 1906, 1909, 1912 (LTSR, CU1-CU5), 2003, 2006, 2009, 2012, 2103, 2106, 2112, 2203 (LTSR), 2206, 2209, 2212
- Citrix XenApp and XenDesktop, versions 6.5, 7.0–7.18, and 7.15 (LTSR, CU3 - CU7)
- Citrix Virtual Apps and Desktops Essentials for Azure
- XenServer
- Google Cloud
- Microsoft Azure
- Microsoft Hyper-V
- Nutanix Acropolis
- VMware vSphere

Desktop provisioning and application delivery

You can use layered images for Persistent desktops, as long as you do *not* enable Elastic layering or User layers (Full or Office 365).

To enable User layers or Elastic layers, you *must use Non-persistent machines*.

Note:

Elastic layering does not support View Persona Management.

App Layering features by edition

Per the [Citrix DaaS and Citrix Virtual Apps and Desktops \(CVAD\) Feature Matrix](#), Citrix App Layering is available in all editions. This means that you can do the following across any number of Citrix DaaS and CVAD sites in your environment:

- Create an unlimited number of OS, platform, and application layers
- Create an unlimited number of layered images
- Create an unlimited number of elastic layers

Install appliance

January 6, 2026

The articles in this section explain how to install the App Layering appliance and agent, as well as additional considerations when installing.

Installing the App Layering appliance

First, you install and deploy the App layering appliance, which is powered by the Enterprise Layer Manager (ELM) technology.

Once installed, you can configure the time zone, NTP servers, and network settings.

For installation details, select your hypervisor:

- [XenServer](#)
- [Google Cloud](#)
- [MS Azure](#)
- [MS Hyper-V](#)
- [Nutanix AHV](#)
- [VMware vSphere](#)

Once you have installed the appliance, you can continue on to [configuring](#) the App Layering appliance.

Installing the App Layering agent

If you need the App Layering agent, install it on the servers where it is required.

- [Install the App Layering agent](#)

Additional considerations

You can install an App Layering appliance on a hypervisor, and use it to publish images to the same hypervisor or to a different one.

- **Management console:** The appliance hosts a web-based management console where you can [manage](#) the system, [layers](#), [image templates](#), and [Image Portability](#) tasks.
- **Backups:** We recommend backing up each appliance so you don't lose the layers you spend time creating. You need a full backup of each appliance to guarantee that you can recover all information from it. Although you can export and import layers, this feature is not designed for failure recovery. For more detailed information on availability, backup, and recovery, see [this article](#) in the Citrix Tech Zone.

XenServer®

September 6, 2025

To install the App Layering service, you deploy the appliance to a virtual machine in XenServer. If the appliance requires a static IP address, you configure one using the appliance configuration utility. The last section describes the next steps.

Requirements

To start installing App Layering, you need the following:

- [A supported version of XenServer](#)
- Storage
- **XenServer account and privileges**

A XenServer account to use for App Layering service.

The account must have XenServer privileges to:

- Create and remove virtual disks.
- Copy and delete layers on virtual disks by using XenServer file APIs.

- **XenServer software and settings**

Access to the XenServer Tools to install on the layer.

- **XenServer resource information**

The XenServer info listed in [XenServer Connector Configuration](#).

Download the installation package

Note:

As of release 2409, XenServer installations will automatically create a 300GB repository disk. You can increase the size of the repository disk if needed.

Download the installation package, `citrix_app_layering_citrix-hypervisor_xx`, from the [Citrix download site](#). The zip file includes:

File	Description
<code>xenserver_xx.x.x.x.zip</code>	Contains the OVA file for the appliance virtual machine
<code>citrix_app_layering_agent_installer.exe</code>	App Layering Agent installer
<code>citrix_app_layering_os_machine_toolsOSxmachine tools.x.x.exe</code>	

Deploy the App Layering appliance

To deploy the appliance:

1. Download the installation zip file [citrix_app_layering_xenserver_xx.x.x.zip](#)
2. Extract the [xenserver_xx.x.x.ova](#) file to a folder on your local drive.
3. In your XenCenter® client, select **File > Import**.
4. In the wizard that opens, select the following values:
 - **Import Source** - Browse to the source on your local drive.
 - **Location** - Choose the XenServer where you want to deploy the appliance.
 - **Storage** - Use the default value to put storage disks on the Local XenServer Storage.
 - **Networking** - Select the correct network for your XenServer configuration.
 - **Security** - This tab is not available. Click **Next**.
 - **OS Fixup Settings** - Select **Don't use Operating System Fixup**.
 - **Transfer VM Settings** - Choose the correct network and choose **DHCP**.
 - **Finish** - Review your settings and then click **Finish**.
5. Switch to **Notification view** and wait for the deployment to complete. The process can take 20–35 minutes.
6. Switch to **Infrastructure** view.
7. Select your new appliance, which has the name:
Citrix Enterprise Layer Manager.
8. Click **Properties**, and enter a new name and description for your new appliance. The new virtual machine has 8 GB of memory and 4 CPUs. Take note of the IP address assigned in the Networking tab.
9. Start Internet Explorer, type the IP address for the new appliance into the address bar and log in as “administrator” with the password “Unidesk1”.

The first time you log on to the appliance that you are required to change the App Layering administrator passwords. For details, see [Change administrative passwords](#).

Be sure to install the App Layering Agent, if your environment requires it. See the next section for details.

If you want the appliance to use a Static IP address

You can change the appliance’s IP address and/or its DNS servers. When the appliance is first deployed, the DNS settings are retrieved through DHCP. If DHCP is not available and you are using sta-

tic IP addresses, once you select **Static**, you are prompted to enter the **IP addresses** for your DNS servers.

1. Log in to the Appliance Configuration utility. Using either your hypervisor console or SSH, log in to the appliance as an administrator.
The first time you log in, use the default password, [Unidesk1](#).
2. At the Action prompt, enter C (for Configure Networking), and press Return.
3. At the next prompt, type D for Dynamic (DHCP) or S for Static.
If you choose **Static**, you are prompted for the IP address and Subnet mask, along with the default addresses for the Gateway and DNS addresses.
4. When prompted, enter Y to save settings.
5. At the Action prompt, enter Q to quit.
6. Restart the appliance.

Next steps

Once the appliance is installed and the IP address is correctly configured, you need to:

- [Install the App Layering agent](#)
- [Configure the App Layering service](#)

Google Cloud

September 6, 2025

Note:

Version 2409 of the App Layering appliance for Google Cloud is not yet available but will be released soon.

To install the App Layering service on Google Cloud, you deploy the appliance to a virtual machine on Google Cloud. You can configure a [static IP address](#), if you need one, using the appliance configuration utility.

Google Cloud project configuration

Configure a Google Cloud project.

Enable APIs

Enable the following Citrix Cloud™ APIs:

- Compute Engine API
- Cloud Storage API
- Cloud Resource Manager API
- Identity and Access Management (IAM) API
- Cloud Build API

See the Citrix Virtual Apps and Desktops instructions for [enabling the Google Cloud APIs](#).

Create a Service Account

The service account needs the following three roles:

- Service Account User
- Compute Admin
- Storage Admin

See [Create a Service Account](#) for details.

Note:

The account for App Layering does not need all of the roles & permissions described in the article above.

Open required firewall ports

Open the [firewall ports for Google Cloud](#) so the appliance can communicate with the Google Cloud.

Virtual machine requirements and settings

Ensure that the virtual machine where you install the appliance on Google Cloud is connected to a Google Cloud virtual network.

The virtual machine that you use for the appliance must be configured as follows:

Name	Value
Virtual Machine Name	App Layering Appliance (Enterprise Layer Manager)
Virtual Machine Generation	Generation 1

Name	Value
Memory	8192 MB
CPUs	4
Boot Disk	unidesk_gcp-system
Additional Disk	unidesk_gcp-repository

Download the installation package

Download the installation package, `citrix_app_layering_gcp_x.x`, from the [Citrix download site](#). The zip file includes:

File	Description
<code>unidesk_gcp-system.tar.gz</code>	Tarball containing system disk for the appliance VM
<code>citrix_app_layering_agent_installer.exe</code>	App Layering agent installer
<code>citrix_app_layering_os_machine_tools.exe</code>	OS Machine Tools

Upload the system disk to Google Cloud

1. Extract the download package, **`citrix_app_layering_gcp_x.x.zip`**.
2. Using the **Navigation** menu on the top left of the Google Cloud UI, go to the **STORAGE** section and select **Storage > Browser**. You can either create a new bucket or add a personal folder to an existing bucket.
3. Select the **upload-disks** bucket link to upload the disk.
4. Select the **UPLOAD FILES** link, and navigate to the directory where you extracted the **`unidesk_gcp-system.tar.gz`** file.

elm-upload-bucket

OBJECTS CONFIGURATION PERMISSIONS RETENTION LIFECYCLE

Buckets > elm-upload-bucket 

UPLOAD FILES UPLOAD FOLDER CREATE FOLDER MANAGE HOLDS DOWNLOAD DELETE

Filter by name prefix only   Filter Filter objects and folders

<input type="checkbox"/>	Name	Size	Type	Created time 	Storage class	Last modified 
No rows to display						

5. Select **unidesk_gcp-system.tar.gz** and click the **Open** button. See the status message on the lower right side of the window to track the progress of the upload.

elm-upload-bucket

OBJECTS CONFIGURATION PERMISSIONS RETENTION LIFECYCLE

Buckets > elm-upload-bucket 

UPLOAD FILES UPLOAD FOLDER CREATE FOLDER MANAGE HOLDS DOWNLOAD DELETE

Filter by name prefix only   Filter Filter objects and folders

<input type="checkbox"/>	Name	Size	Type	Created time 	Storage class	Last modified 
<input type="checkbox"/>	unidesk_gcp-system.tar.gz	4.1 GB	application/x-gzip	Feb 10, 2021, 9:16:24 PM	Standard	

Create an image for the system disk

1. Select the **Navigation** menu on the top left of the Google Cloud window, go to the **STORAGE** section, and select **Compute Engine > Images**.
2. Select **[+] CREATE IMAGE**.
3. In the Name field, enter a permanent name for the image. Google Cloud warns that the name is permanent.
4. Under Source, select **Cloud Storage file**.
5. Under **Cloud Storage file**, use the **Browse** button to select the **unidesk_gcp-system.tar.gz** file.
6. Under **Location**, select Multi-regional or Regional, and the Location.
7. Select **Create**. The **Images** page tracks the image creation as it progresses. A green checkmark appears when the image is created. If creation fails, the image name no longer appears on the Images screen.

Create a VM instance

Next, create a VM instance, attach the system disk, and create repository disks:

1. Select the **Navigation** menu on the top left of the Google Cloud window, go to the **COMPUTE** section, and select **Compute Engine > VM instances**.
2. Select **[+] CREATE INSTANCE**.
3. Enter data, choosing the desired Region and Machine configuration.
4. Expand the **CPU platform and GPU** section, and check **Turn on display device**.
5. In the **Boot disk** section, select the **Change** button. The Boot disk window opens.

Boot disk

Select an image or snapshot to create a boot disk; or attach an existing disk. Can't find what you're looking for? [Search](#)

[Public images](#) [Custom images](#) [Snapshots](#) [Existing disks](#)

Show images from

Show deprecated images

Image

Created on Oct 20, 2020, 11:19:41 PM

Boot disk type

Size (GB)

30

6. Select the **Custom images** tab and the system image previously created. Your data populates the project's **Show images from** and **Boot disk type and Size (GB)** dropdowns.
7. Click **Select**. The **Create an instance** page appears with the boot disk image set to the image you just selected.
8. Select **Allow HTTPS traffic**.
9. Expand the **Management, security, disks, networking, sole tenancy** section. Select the **Disks** tab.

App Layering

Additional disks [?](#) (Optional)

New disk (gcp-elm-repo-disk, Blank, 300 GB)

Name (Optional) [?](#)
Name is permanent
gcp-elm-repo-disk

Description (Optional)

Type [?](#)
Standard persistent disk

Snapshot schedule
Use snapshot schedules to automate disk backups. [Scheduled snapshots](#)

No schedule

Create snapshot schedules to automatically back up your data. [Learn more about creating snapshot schedules](#) Dismiss

Source type [?](#)
 Blank disk Image Snapshot

Mode
 Read/write Read only

Deletion rule
When deleting instance
 Keep disk Delete disk

Size (GB) [?](#)
300

Estimated performance [?](#)

Operation type	Read	Write
Sustained random IOPS limit	225.00	450.00
Sustained throughput limit (MB/s)	36.00	36.00

10. In the **Additional disks** section, select **+ Add new disk** button.
11. Enter a descriptive name.
12. For **Source type**, select the **Blank disk** tab and enter the **Size (in GB)** for the repository disk.
13. Select **Done** to finish adding an additional disk.
14. Still in the Management, security, disks, networking, sole tenancy section, select the **Networking** tab.

Management Security Disks **Networking** Sole Tenancy

Network tags ? (Optional)

Hostname ?
Set a custom hostname for this instance or leave it default. Choice is permanent
gcp-elm-host-name

Network interfaces ?
Network interface is permanent

Network interface

Network ?
network-name-demo

Subnetwork ?
subnetwork-name-demo (---.0.0/20)

Primary internal IP ?
Ephemeral (Automatic)

External IP ?
None

IP forwarding ?
Off

Done **Cancel**

15. In the **Network interface** section, select a **Network**. Verify that a value appears in the **Subnet**.

work field, and select a value for it.

Important:

We recommend that you do not create an external (public) IP address.

16. (Recommended) From the **External IP** dropdown, select **None**.
17. Select **Done** to finish editing the network interface.
18. Select **Create** to create the VM instance. The VM is created and the appliance powers on. Once the appliance is created, you no longer need the system disk download.
19. To clean up after installation, remove:
 - The .zip file that you downloaded.
 - The folder containing the uncompressed files.
 - The **unidesk_gcp-system.tar.gz** file from the folder where you uploaded it.

If you want the appliance to use a Static IP address

You can change the appliance's IP address and its DNS servers. When the appliance is first deployed, the DNS settings are retrieved through DHCP. If DHCP is not available and you are using static IP addresses, once you select **Static**, you are prompted to enter the IP addresses for your DNS servers.

1. Log in to the Appliance Configuration utility, using the steps and default password described in [Appliance settings](#).
2. At the Action prompt, enter **C** (for Configure Networking), and press Return.
3. At the next prompt, type **D** for Dynamic (DHCP) or **S** for Static.
If you choose Static, you are prompted for the IP address and Subnet mask, along with the default addresses for the Gateway and DNS addresses.
4. When prompted, enter **Y** to save settings.
5. At the Action prompt, enter **Q** to quit.
6. Restart the appliance.

Next steps

Once the appliance is installed and the IP address is correctly configured, you need to:

- [Install the App Layering agent](#)
- [Configure the App Layering service](#)
- [Access the App Layering management console](#)

MS Azure or Azure Government

January 5, 2026

To install the App Layering service, you deploy the appliance to a virtual machine using the Azure (or Azure Government) resource manager.

- **Requirements.** - Make sure you have the Azure (or Azure Government) resources you must install the appliance.
- **Install the App Layering appliance in Azure (or Azure Government).** - To install the appliance, you download the installation package and deploy the appliance.
- **Next steps** - Links to steps for installing the App Layering agent and configuring the App Layering service.

Note: If you have not already set up a connection to an Azure Virtual Network, see [Get started with Azure](#) or [Azure Government](#) for more information.

The steps for installing the appliance in Azure and Azure Government are the same, aside from the installation script that you run.

Requirements

Before installing App Layering in Azure or Azure Government, be sure you have the following.

- **An Azure account and subscription**

To deploy and configure the App Layering appliance, you need the credentials for an account that has administrative access to your Azure subscription. For more information, refer to the [Microsoft Azure Sign-in page](#).

- **A Virtual Network in Azure (or Azure Government)**

Your deployment in Azure (or Azure Government) can operate in a point-to-site or site-to-site Virtual Network. The appliance and its network file share must have network connectivity. However, the appliance does not require network connectivity to the layered images you publish. A site-to-site connection between your corporate and Azure (or Azure Government) networks is recommended for accessing the management console on the appliance. For more information, refer to the [Microsoft Azure Virtual Network page](#).

Note: If you have not already set up a connection to an Azure Virtual Network, see [Get started with Azure](#) for more information.

- **A Network File Share (Azure or Azure Government specifics)**

A file share server in Azure (or Azure Government) performs better than an on-premises file share. Even though the Azure (or Azure Government) File Share feature is not supported, you

can use an existing network file share or create a file share in the Azure (or Azure Government) environment.

Important: Using Premium Storage is recommended.

- **Azure (or Azure Government) Resource Manager**

App Layering works with Azure's Resource Management (ARM) model. We do not support Azure's Classic deployment model or Bicep. All resources such as virtual network, file shares and OS machines that App Layering will interact with must be created with Azure Resource Manager. For more information, refer to the [Azure Resource Manager overview page](#).

- **Azure Powershell v7**

Azure Powershell v7 must be installed on the Windows system that will be used to install the appliance in Azure.

- **Assigned managed identity**

The App Layering appliance must be [assigned](#) a managed identity to support deployment on Azure.

Install the App Layering appliance

To deploy the App Layering appliance to Azure or Azure Government:

1. Check the contents of the installation package.
2. Learn what's included in the installation script.
3. Deploy the appliance.

Installation package

The installation package, is named `citrix_app_layering_azure_yy.m.b.zip`, where:

- *yy* is the year
- *m* (or *mm*) is the month
- *b* is the build
- *x* is the sub-version
- For example, `citrix_app_layering_azure_24.3.0.zip` contains `azure_24.3.0.3.zip`

File	Description
<code>azure_yy.m.b.x.zip</code>	Contains the VHD file (<code>unidesk_azure-system.vhd</code>) for the appliance VM
<code>New-AzAppLayeringAppliance.ps1</code>	Installation Script for Azure and Azure Government
<code>citrix_app_layering_agent_installer.exe</code>	App Layering agent installer
<code>citrix_app_layering_os_machine_tools.exe</code>	OS Machine Tools
<code>DeployAzureRmVm.template.json</code>	App Layering template

Installation script

The installation script included in the installation package does the following.

1. Copies the included VHD to the Azure location you specify.
2. Creates a virtual machine in Azure using the VHD.
3. Creates and Attaches the repository disk.
4. Boots the Azure appliance.

When you run the script:

- **IMPORTANT:** Be sure to note the Resource group location you select, as you will need this information later. For more information about resource groups, refer to [Using the Azure Portal to manage your Azure resources](#).
- When selecting a virtual machine size, it is strongly recommended that you select a machine with 4 CPUs, and at least 14 GB of memory (script default).
- The name you specify for the new virtual machine must comply with Azure naming conventions.
- Select a Virtual Network in which HTTP port: 80 is accessible (Public IP can be disabled).

Using the installation script `New-AzAppLayeringAppliance.ps1` is the new installation script which replaces the previous script, `AzureELMDeploymentV7.ps1`.

The new installation script now accepts parameters and no longer requires the user to manually provide input to the command line prompt.

You may still run the script without parameters and provide the inputs manually.

Requirements

- The new script requires the latest version of `DeployAzureRmVm.template.json` included in the package.
- The new script requires the current version of `AzCopy`. The binary for `AzCopy` must either be in the same directory as `New-AzAppLayeringAppliance.ps1` or added to your path. For more information on installing `AzCopy` see [Get started with AzCopy](#)

Parameters

Parameter	Description	Data Type	Default Value	Required? (Y/N)
DeploymentName	The base name used to derive names for all resources created during the deployment.	string	-	Y
ResourceGroupName	The name of the Azure resource group where resources will be deployed. If it doesn't exist, it will be created.	string	<code>"{DeploymentName} -group"</code>	N
LocationId	The Azure region where all resources will be created or reused.	string	-	Y
SubnetId	The ID of the subnet that the appliance will be connected to.	string	-	Y
ParamIpAddress	The IP address to assign to the appliance.	string	<code>"Dynamic"</code>	N
RepositoryDiskSizeInGb	The size of the repository disk attached to the appliance (in gigabytes).	int	300	N

App Layering

Parameter	Description	Data Type	Default Value	Required? (Y/N)
VhdPath	The path to the appliance VHD file to upload.	string	<code>"{ current_dir } /azure_ */unidesk\\ _azure- system.vhd"</code>	N
AzureCopyPath	The path to the <code>azurerecopy.exe</code> utility.	string	<code>"{ current_dir } /azurerecopy .exe"</code>	N
TrustedLaunch	Enables trusted launch for the appliance VM	-	-	N
ForceUpload	Forces the disk to be uploaded even if it already exists in Azure	-	-	N
DefaultProfile	The Azure credentials, account, tenant, and subscription used for communication with Azure.	string	-	N

Example Usages

```
1 .\New-AzAppLayeringAppliance.ps1
```

```
1 .\New-AzAppLayeringAppliance.ps1 -DeploymentName "MyDeploymentName" -  
LocationId "East US" -SubnetId /subscriptions/<subscription-id>/  
resourceGroups/<resource-group-name>/providers/Microsoft.Network/  
virtualNetworks/<virtual-network-name>/subnets/default
```

Deploy the appliance in Azure or Azure Government

The App Layering ZIP download requires 31GB of space when uncompressed.

1. Extract the download package, `citrix_app_layering_azure_yy.mm.x.zip`. The files

included are listed above.

2. Extract the ZIP file to a folder on your local drive.
3. Open an Azure Powershell window.
4. Execute the installation script (included in the installation package) with the appropriate parameters. Running the script:
 - Copies the VHD to the Azure location of your choice.
 - Creates and attaches the repository disk.
 - Boots the appliance.
 - If the script fails, check the values to make sure that the values are correct for your environment.

The first time you log onto the appliance you are required to change the App Layering administrator passwords. For details, see [Change administrative passwords](#).

MS Hyper-V

October 23, 2024

To install the App Layering service, you deploy the appliance to a virtual machine in Hyper-V. If the appliance requires a static IP address, you configure one using the appliance configuration utility, as described below. The last section describes next steps.

Requirements

Ensure that the Hyper-V virtual machine from where you are installing the appliance meets the following prerequisites:

- Windows Server 2016, Windows Server 2012 R2
- Virtual network in Hyper-V

Download the installation package

Note:

As of release 2409, we no longer include a repository disk with the installation package. You must create a new disk and attach it to the appliance VM.

Download the installation package, `citrix_app_layering_hyperv_xx`, from the [Citrix download site](#). The zip file includes:

File	Description
<code>hyperv_xx.x.x.x.vhdx</code>	VHDX file for the appliance VM
<code>citrix_app_layering_agent_installer.exe</code>	App Layering agent installer
<code>citrix_app_layering_os_machine_tools.exe</code>	OS Machine Tools

Virtual machine settings for the appliance

Note:

As of release 2409, we no longer support Generation 1 VMs or booting with BIOS.

The virtual machine that you use for the appliance must be configured as follows:

Name	Value
Virtual Machine Name	App Layering Appliance (Enterprise Layer Manager)
Virtual Machine Generation	Generation 2
Secure Boot Template	Microsoft UEFI Certificate Authority
Memory	8192 MB
CPUs	4
Disk 1	<code>unidesk_hyperv-system.vhdx</code>
Disk 2	<code><your-new-repository-disk-name>.vhdx</code>

Deploy the appliance in Hyper-V

To deploy the appliance:

1. Extract the download package, `citrix_app_layering_hyperv_xx.x.x.zip`. One disk (`vhdx`) file is included.
2. Copy the disk to a storage location that the Hyper-V server can access.
3. Open the Hyper-V Manager, right-click the Hyper-V server where you want to deploy the appliance, and select **New Virtual Machine**.

4. On the first wizard tab that opens, click **Next** to begin configuring the virtual machine.
5. On the Specify Name and Location tab, set the **Name** and **Location** of the new VM. Ideally, use the location where you extracted the disk in step 2.
6. On the Specify Generation tab, ensure that *Generation 2 VMs* is selected. Only Generation 2 VMs are supported.
7. On the Assign Memory tab, set the VM to use 8 GB of RAM. Make sure the **Use Dynamic Memory for this virtual machine** check box is *not* selected.
8. On the Configure Networking tab, specify the NIC for the network adapter to use to connect to the network.
9. On the Connect Virtual Hard Disk tab, attach the system disk ([unidesk_hyperv-system.vhdx](#)), the disk that you extracted in step 2.
10. On the Summary tab, verify your choices and click **Finish**.
11. Back in the Hyper-V Manager, select the VM and click **Settings** from the VM panel.
12. Select **Hardware > Processor**, and set the Number of Virtual Processors to 4.
13. Select **SCSI Controller > Hard Drive**, click **Add**. (This MUST be done prior to first boot)
14. Select the **Virtual hard disk** radio button, click **New**.
15. This will open up a window for the New Virtual Hard Disk Wizard. Click **Next** to begin configuring the new disk.
16. On the Choose Disk Type tab, select **Fixed size** and click **Next**.
17. On the Specify Name and Location tab, set the **Name** and **Location** of the new disk. Ideally, use the location where you extracted the disk in step 2.
18. On the Configure Disk tab, select the **Create a new blank virtual hard disk** radio button, set the **Size** to 300 GB and click **Next**.
19. On the Summary tab, verify your choices and click **Finish**.
20. Select **Security > Secure Boot**, select **Enable Secure Boot** and set the Secure Boot template to **Microsoft UEFI Certificate Authority**.
21. Power on the VM.
22. Type the IP address and logon for an administrator with permission to access the App Layering management console.

The first time you log on to the appliance you are required to change the App Layering administrator passwords. For details, see [Change administrative passwords](#).

Configure a Static IP address for the appliance, if necessary

You can change the appliance's IP address and its DNS servers. When the appliance is first deployed, the DNS settings are retrieved through DHCP. If DHCP is not available, and you are using static IP addresses, once you select **Static** you are prompted to enter the IP addresses for your DNS servers.

1. Log in to the Appliance Configuration utility, using the steps and default password described in

[Appliance settings.](#)

2. At the Action prompt, enter C (for Configure Networking), and press Return.

3. At the next prompt, type **D** for Dynamic (DHCP) or **S** for Static.

If you choose Static, you are prompted for the IP address and Subnet mask, along with the default addresses for the Gateway and DNS addresses.

4. When prompted, enter Y to save settings.

5. At the Action prompt, enter Q to quit.

6. Restart the appliance.

Next steps

Once the appliance is installed and the IP address is correctly configured, you need to:

- [Install the App Layering agent](#)
- [Configure the App Layering service](#)

Nutanix AHV

September 6, 2025

To install the App Layering service, you deploy the appliance to a virtual machine on your hypervisor.

- **Requirements.** Make sure you have the required Nutanix accounts, settings, and resources.
- **Install the App Layering appliance in Nutanix.** - To install the appliance, you download the installation package and deploy the appliance in Nutanix.
- **If the appliance needs a Static IP address.** - You can set a Static IP address by accessing the appliance configuration utility.
- **Next steps.** - Links to steps for installing the App Layering agent and configuring the App Layering service.

Requirements

If you are installing the appliance and building your Layers on Nutanix VMs, or you are publishing Layered Images that will be used in a Nutanix environment, you need the following settings and resources:

- **Nutanix account and privileges**

- An existing or new Nutanix AHV account to use for App Layering.

- The account must have Nutanix AHV privileges to:
 - Create and remove virtual disks.
 - Copy and delete layers on virtual disks using Nutanix file APIs.
- **Nutanix software and settings** - Access to the Nutanix Tools to install on the layer.
- **Nutanix resource information** - The info listed in the [Nutanix AHV Connector Configuration](#).

Install the App Layering appliance in Nutanix AHV

To deploy the App Layering appliance to Nutanix:

- Verify the contents of the installation package.
- Deploy the appliance.

Installation package

Note:

As of release 2409, we no longer include a repository disk with the installation package. You must create a new disk and attach it to the appliance VM.

The installation package, [citrix_app_layering_nutanix_xx](#), includes:

File	Description
nutanix_xx.x.x.x.zip	Contains the IMG file for the appliance VM
citrix_app_layering_agent_installer.exe	App Layering Agent installer
citrix_app_layering_os_machine_toolsOSxMachine Tools.x.x.exe	

Deploy the appliance in Nutanix

1. Extract the download package, [citrix_app_layering_nutanix_xx.x.x.zip](#). The files included are listed above.
2. Next, unzip the [nutanix_xx.x.x.zip](#) file, containing one [img](#) file.
3. In the Nutanix Prism console, select the **Tools** menu in the top right corner of the UI, and choose **Image Configuration**.
4. Click the **Upload Image** button, and name the disk.

5. Select the **DISK** Image Type.
6. Select the **Upload a File** option, browse to your file share, and choose the App Layering Boot Disk ([citrix_applayering_nutanix-system.img](#)). Wait for the upload to complete.
7. Select **Tasks** and make sure *both* the Image Create and Image Update tasks are complete. Once this is done, you can create the virtual machine.
8. Select the **VM** (Virtual Machine) tab on the top left drop-down menu, and click the **Create VM** button.
9. In the **General Configuration** section, complete the **Name** and **Description** of the new VM.
10. In the **Compute Details** section, set VCPU(S) to **1**.
11. Set the **Number of Cores per vCPU** to **4**.
12. Set Memory to **8 GB**.
13. Add the system disk. In the **Disk**s section, click **Add new disk** and choose type **Disk**.
14. In the **Operation** drop-down menu, choose **Clone from Image Service**.
15. In the **Bus Type** drop-down menu select **SCSI**.
16. In the Image Box select the Boot disk that you uploaded during step (4), and click **Add**.
17. Add the repository disk. In the **Disk**s section, click **Add new disk** and choose type **Disk**.
18. In the **Operation** drop-down menu, choose **Allocate on Storage Container**.
19. In the **Bus Type** drop-down menu select **SCSI**.
20. In the **Size (GiB)** box, enter **300** GB (or an appropriate size of your choosing) and click **Add**.
21. Remove the CD-ROM (IDE) disk from the Disks.
22. In the **Boot Configuration** section, select the **UEFI** radio button
23. Select the **Secure Boot** checkbox.
24. Add NIC by clicking on **Add new NIC**.
25. Click **Save** to create the VM.
26. Power on the VM.

Add the Citrix App Layering™ appliance to the Nutanix allow list

Be sure to add the appliance to the Nutanix allow list, so that the appliance is allowed to connect to Nutanix.

If the appliance needs a Static IP address

You can change the appliance's IP address and its DNS servers. When the appliance is first deployed, the DNS settings are retrieved through DHCP. If DHCP is not available, and you are using Static IP addresses, once you select **Static** you are prompted to enter the **IP addresses** for your DNS servers.

1. Log in to the Appliance Configuration utility, using the steps and default password described in [Appliance settings](#).
2. At the Action prompt, enter **C** (for Configure Networking), and press Return.

3. At the next prompt, type **D** for Dynamic (DHCP) or **S** for Static. If you choose Static, you are prompted for the IP address and Subnet mask, along with the default addresses for the Gateway and DNS addresses.
4. When prompted, enter **Y** to save settings.
5. At the Action prompt, enter **Q** to quit.
6. Restart the appliance.

Next steps

Once the appliance is installed and the IP address is configured as you want it to be, you can install the App Layering agent, if necessary, and configure the App Layering service. See [When the agent is required](#).

- [Install the App Layering agent](#)
- [Configure the App Layering service](#)

VMware vSphere

October 23, 2024

To install the App Layering service, you deploy the appliance to a virtual machine on your hypervisor.

Requirements

The App Layering appliance requires the following virtual machine settings, vSphere requirements, and vCenter permissions.

Virtual machine settings

When you create the appliance's virtual machine, it requires:

- 4 CPUs
- 8 GB RAM

VMware vSphere requirements

To install the App Layering appliance in a VMware vSphere environment, you need the following:

- A supported version of VMware vSphere.
- A virtual network in vSphere.
- vCenter account and privileges.

You need a vCenter account with permissions on a data center for:

- Creating and removing virtual machines.
- Creating, copying, and removing virtual machine disks.

Also, the account needs this permission at the vCenter level:

- Removing virtual machines from inventory.

For details, see the list of [vCenter permissions](#) in the next section.

- The Role for App Layering that would be applied to the data center.

To set up a new role:

1. In the vSphere Client, navigate to **Home > Administration > Roles**.
2. Click **Add Role**.
3. Enter a name. For example: *CALAdmin*.
4. Set the privileges for the account.

- Privileges defined for the vCenter role that you’re using for the App Layering service. You apply the new role to the Data Center you plan to use for App Layering. For details, see the list of [vCenter permissions](#) in the next section.
 1. Open the **Assign Permissions** window.
 2. In the vSphere Client, navigate to **Home > Inventory > Hosts and Clusters**.
 3. Select your **DataCenter** and then right-click, and select **Add permission**.
 4. In the **Assign Permissions** window, under **Assigned Role**, expand **All Privileges**.
 5. Select the required [vCenter permissions](#).
 6. Select the **Propagate to Child Objects** check box, and click **OK**.
- The App Layering role must be assigned to the administrator account, as follows:
 1. Add the administrator account and then assign the App Layering role to it.
 2. Allow the permissions to propagate to the entire data center.

Note: If you want to restrict this user from accessing specific folders in the data center, grant the user more restrictive permissions for those folders.

- Because the **Virtual Machine > Inventory > Remove** permission must be assigned at the vCenter level, you must create a second role.

1. In the vSphere Client, navigate to **Home > Administration > Roles**.
2. Click **Add Role**, and enter a name, for example: CALAdmin-vmremove.
3. Add **Virtual Machine > Inventory > Remove**, and leave everything else set to *read-only*.
4. In the vSphere Client, navigate to **Home > Inventory > Hosts and Clusters**.
5. Select the **vCenter Permissions** tab, right-click, and select **Add permission** (or modify the permissions on an existing account).

Note:

Make sure to use the account that has the data center permissions set.

6. Select the new CALAdmin-vmremove role you defined. Make sure the Propagate to Child Objects check box is selected and then click OK.

Note:

If your security policy allows, you can set all permissions at the vCenter level instead.

vCenter permissions

Expand **All Privileges**, then each of the following permissions categories, and select the required permissions. Permission names might differ depending on your release of VMware VSphere.

Note:

As of release 19.11, the Host.Configuration.System Management permission is no longer required.

Datastore

- Allocate space
- Browse datastore
- Low level file operations

Folder

- Create folder

Global

- Cancel task

Network

- Assign network

Resource

- Assign virtual machine to resource pool

vApp

- Export
- Import

Virtual machine > Configuration

- Add existing disk
- Add new disk
- Add or remove device
- Advanced
- Change CPU count
- Change resource
- Memory
- Modify device settings
- Remove disk
- Rename
- Set annotation
- Settings
- Upgrade virtual machine compatibility

Virtual machine > Interaction

- Configure CD media
- Console interaction
- Connect devices
- Power off
- Power on
- Reset
- VMware Tools Install

Virtual machine > Inventory

- Create from existing
- Create new
- Remove

Virtual machine > Provisioning

- Clone template (optional, but required to use a vSphere template as the source virtual machine)
- Clone virtual machine

Virtual machine > Snapshot management

- Create snapshot
- Revert to snapshot
- Remove snapshot

Install the App Layering appliance in VMware vSphere

To deploy the App Layering appliance to vSphere:

- Make sure that you have the vSphere requirements.
- Familiarize yourself with the contents of the installation package.
- Deploy the appliance.

Installation package

The installation package, `citrix_app_layering_vsphere_xx`, includes:

File	Description
<code>vmware_xx.x.x.x.zip</code>	Contains the OVA file for the appliance VM
<code>citrix_app_layering_agent_installer.exe</code>	App Layering Agent installer
<code>citrix_app_layering_os_machine_toolsOSMachine Tools.x.x.exe</code>	Machine Tools

Deploy the appliance

1. Extract the download package, `citrix_app_layering_vsphere_xx.x.x.zip`. The files included are listed at the beginning of this topic.
2. Extract the `vmware_xx.x.x.x.ova` to a folder on your local drive.
3. In the vSphere Web Client you are using, navigate to the **VMs and Templates** page.
4. Right-click the folder in vSphere where you want to deploy the template and select **Deploy OVF Template**. The “Deploy OVF Template” wizard appears.
5. In the Deploy OVF Template wizard, do the following:
 - a) On the **Select source** page, select the **Local file** option, and browse to the `vmware_xx.x.x.x.ova` file to select it.
 - b) On the **Select name and folder** page, designate a name and location for the deployed OVF template.

- c) On the **Select a resource** page, select a location to run the deployed OVF template.
- d) On the **Select storage** page, select the **Thick Provision Lazy Zeroed** setting of the **Select virtual disk format** option, select a storage policy, and specify a storage location.
- e) On the **Setup networks** page, select your vSphere virtual network in the **Destination** column and select the **IPv4** setting of the **IP protocol** option.
- f) On the **Customize template** page, input your Layer Repository Disk Size (GB). The default is 300 GB.
- g) On the **Ready to complete** page, review the template settings and then click **Finish** when you are satisfied with the settings.

If the appliance needs a Static IP address

You can change the appliance's IP address and its DNS servers. When the appliance is first deployed, the **DNS** settings are retrieved using the Dynamic Host Configuration Protocol (DHCP). If DHCP is not available, you can use static IP addresses. You select **Static** and enter the **IP addresses** for your DNS servers.

1. Log in to the Appliance Configuration utility, using the steps and default password described in [Appliance settings](#).
2. At the Action prompt, enter **C** (for Configure Networking), and press Return.
3. At the next prompt, type **D** for DHCP, or **S** for Static. If you choose **Static**, you are prompted for the following:
 - **IP address**
 - **Subnet mask**
 - **Default Gateway address**
 - **Default DNS address**
4. At the next prompt, enter **Y** to save settings.
5. At the Action prompt, enter **Q** to quit.
6. Restart the appliance.

Next steps

Once the appliance is installed and the IP address is configured, proceed to:

- [Install the App Layering agent](#)
- [Configure the App Layering service](#)

Docker

March 31, 2025

To install the App Layering service, you deploy the appliance to a virtual machine on your hypervisor.

How To Install WSL and Docker

These instructions apply to all App Layering image types but are primarily intended for use with full UserLayer deployments of non-persistent VDI desktops. See related guidance at [Manual installation steps for older versions of WSL](#). For App Layering, these instructions will provide WSL and Docker as a base for domain users in a deployed VDI pool of VMs.

Domain users attempting to use WSL and/or Docker must be members of the domain admin group. Using a separate admin account on behalf of a domain user during the user portion of these steps will not work.

Other pre-requisites exist that are not covered below, such as ensuring the virtual machines have nested virtualization enabled (if you plan to use WSL2), and having a computer GPO present to create and populate a ‘docker-users’ group to which users need to be a member. Windows 10/11 versions should be current with the version of WSL/Docker being used.

Step-by-step guide - Docker / WSLv2 (Win 10 or 11)

1. In a new app layer (preferred), or os-layer revision, as admin:

a. `dism.exe /online /enable-feature /featurename:Microsoft-Windows-Subsystem-Linux /all /norestart`

b. `dism.exe /online /enable-feature /featurename:VirtualMachinePlatform /all /norestart`

c. Disable wsl auto updates

```
1 i. reg add hklm\software\Microsoft\Windows\CurrentVersion\Lxss /t  
REG_DWORD /v AutoUpdate /d 0 /f
```

d. Reboot to install features, then log back in

e. Download and install the latest stable WSL release from [Microsoft/WSL](#)

```
1 i. Move the download to a place where users can access it later (do this BEFORE installing it)
```

2

```
3 ii. Install of WSL must be done from the same folder location
4     that users will be using
5
6     1. Do not use the wsl --install command
7
8     2. Sample install command: 'msiexec -i wsl.2.3.26.0.x64.msi /
9         L*V wsl-install.log'
10
11    3. Suggestion: create a "UserInstall.cmd" script containing a
12        similar msiexec command to be used in step 2.a
13
14    iii. If an older WSL installation is present it should be
15        uninstalled first (using 'msiexec /x' to fully remove it)
16
17    iv. It is not recommended to install a specific distro into the
18        app layer. Users can choose a distro later if they wish to.
19
20    v. Do not remove the WSL .msi installer from this folder. Users
21        will need to apply this exact .msi during their setup later.
```

f. Use 'wsl –version' to verify wsl is working

g. (Optional) add a .wslconfig to the Default user profile

```
1 i. Existing users would need to copy this to their own profile
2
3 ii. The default network operation for wsl2 should be established
4     in a file named .wslconfig and placed in \Users\Default.
5
6     1. .wslconfig contains two lines: '[wsl2]' and '
7         networkingMode=mirrored' (without quotes, case as shown)
8
9     2. Only new users-layer users will automatically have this
10        file copied to their profile area
11
12    3. Place a copy of the .wslconfig file in the WSL installer
13        location for existing users to copy
```

h. Download and install Docker

```
1 i. See [Install Docker Desktop on Windows] (https://docs.docker.com/desktop/install/windows-install/)
2
3 ii. Docker will log out
4
5 iii. Log back in (Docker will start itself, but this can be
6     ignored now)
7
8 iv. 'Accept' terms and Close the Docker Desktop sign-in window
9
10 v. In the latest releases of Docker, an error popup may appear at
11     this time (see exact text of this popup in Info section below)
12     . Docker now seems to demand the presence of an installed
13     distro. This popup can be dismissed and ignored, and Docker
```

will exit due to **this**. That too can be ignored at **this** point in the installation steps. i. Finalize the layer (**for** app layer, must be included in the image; not delivered elastically)

2. For each user logging in:

a. Locate and run the WSL installer (saved from step 1.e.i)

```
1 i. This will be a fast and minimal install that only sets up  
     reparse-points for the user  
2  
3 ii. This is a one-time install for each user, but can be used  
      again if wsl problems arise  
4  
5 iii. Do not copy the wsl .msi to an alternate location before  
       installing it (e.g., 'cd \WSL' and repeat the original/similar  
       msiexec command)
```

b. Install a private distro using powershell: (example using the ubuntu-1604 distro)

```
1 i. Invoke-WebRequest -Uri [https://aka.ms/wsl-ubuntu-1604] (https  
     ://aka.ms/wsl-ubuntu-1604) -OutFile ~/Ubuntu.zip -  
     UseBasicParsing  
2  
3 ii. mkdir ~/distros  
4  
5 iii. Expand-Archive ~/Ubuntu.zip ~/distros/Ubuntu  
6  
7 iv. Run the installer as admin ~/distros/Ubuntu/ubuntu*.exe (   
       exact name is version specific)  
8  
9 v. Provide a username and password you will use to log into wsl,  
     and complete the install  
10  
11 vi. When this completes, update the distro in the distro-specific  
      fashion ( sudo apt-get update; sudo apt-get upgrade )  
12  
13 vii. (Optional) create a desktop shortcut to the installed distro  
      's main application  
14  
15 viii. Leave your distro window running for the remaining steps
```

c. Launch Docker Desktop from start menu

```
1 i. Allow Docker to initialize and start the docker engine  
2  
3 ii. Important: Docker needs to be made aware of your distro  
4  
5   1. In Docker, go to 'Setting->Resources->WSL Integration'  
6  
7     a. Enable the 'additional distro' that matches yours  
8  
9     b. Disable the 'default WSL distro' selection  
10
```

```
11 iii. Also, you can control Docker auto-startup in 'Settings->  
12 General'  
13 iv. Remember that your distro/WSL session needs to be started  
before the Docker engine can start successfully
```

Upgrading

Any time an upgraded WSL/Docker App Layering is rolled out, each affected user will need a repair operation using the prior WSL/Docker layer.

Install the App Layering agent

September 6, 2025

The App Layering agent enables the appliance or a packaging machine to run PowerShell commands locally. If you supply the proper credentials, the agent can run PowerShell commands as a specific user.

When the agent is required

The App Layering agent is required if you plan to:

- Launch scripts using your connector configurations.
- Run the App Layering appliance in Microsoft Hyper-V.
- Publish layered images to Citrix Provisioning.

You install the agent in the following locations:

- On all Hyper-V servers you plan to use for layer creation or image publishing.
- On any Citrix Provisioning servers where you plan to publish layered images.
- In locations where your connector configurations run PowerShell scripts.

Registering the agent with the appliance

The App Layering agent installer prompts you to register the agent with an App Layering appliance. If you do not register the agent during installation, you can manually register it later. However, the PowerShell scripts do not run until the agent is registered with the appliance.

Prerequisites

Before you install the App Layering agent, make sure you that the system where you are installing the agent meets the following requirements:

- An account with administrator privileges
- .NET 4.5
- PowerShell 3.0 or later
- PowerShell Snap-in. Before using the App Layering agent on a Citrix Provisioning Services server, you must ensure that the PowerShell Snap-in is installed. See the steps to [manually register the agent](#) with the appliance.

If you previously installed the App Layering agent, you can download updated versions from the Citrix® downloads page. The agent installation package is included in the App Layering ZIP download.

To download the App Layering agent

1. Go to the [Citrix home page](#) and then click **Sign in**.
2. Click **Downloads** and select **Citrix App Layering™** from the list.
3. On the Citrix App Layering page, under **Citrix App Layering > Product Software**, click Citrix App Layering.
4. At the bottom of the page, click **Tools**.
5. Click the Citrix App Layering agent **Download File**.

The file citrix_app_layering_agent_installer.exe downloads to your computer.

Install the App Layering agent

The App Layering agent installer prompts you to register the agent with an App Layering appliance. If you do not register the agent during installation, you can manually register it later. However, keep in mind that the agent must be registered with the appliance before the PowerShell scripts can run.

1. Using an account with administrator privileges, log into the system where you are installing the agent.
2. Copy the Citrix_app_layering_agent_installer.exe file to a convenient location on the server.
3. Run the Citrix_app_layering_agent_installer.exe as an administrator, and when prompted, enter the path to the directory where you want to install the App Layering agent. The default location is C:\Program Files (x86)\Citrix\Agent.

The agent installer checks to see that all prerequisites are present. If any prerequisites are missing, the installer reports this and exits without installing.

4. The installer prompts you for an Agent Port number. You can accept the default port number (8016) or specify a different one if the default port is already in use.
5. The installer prompts you for the credentials (address, user name, and password) for your App Layering appliance. Register the App Layering agent with the appliance by entering the IP address and logon credentials for a management console user on the appliance with administrator privileges, (for example, the credentials you use to log on to the management console).

Note:

If the App Layering appliance is not available or you choose not to register with it now, you can manually register later using the procedure described in Register with the App Layering appliance manually.

6. Click **Finish** to exit the wizard.

Manually register the App Layering agent with the appliance

Register the App Layering agent with the appliance.

If the App Layering Agent is not registered with an App Layering appliance during installation, you can register it later by using the following procedure.

1. As an administrator, log in to the server where you installed the App Layering agent.
2. Open a command window (cmd.exe) as administrator and navigate to the directory where the App Layering Agent is installed. The default location is C:\Program Files (x86)\Citrix\Agent.
3. Run the following PowerShell command:

`Enable-PSRemoting`

4. Verify that port 8016 is open by running this command:

`netstat -a`

5. Open a command window (cmd.exe) as administrator and navigate to the directory where the App Layering Agent is installed. The default location is:

`C:\Program Files (x86)\Citrix\Agent`

6. Run the following command, using the IP address of the appliance where indicated:

```
1 Citrix.AppLayering.Agent.Service.exe register /i /e:  
IP_address_of_appliance /u:Administrator
```

7. When prompted, enter the password for a user who has administrator privileges in the management console on the appliance.

- When the registration process finishes, a message appears informing you of the successful outcome. The App Layering agent registration appears in the audit log for the management service.

If the process does not succeed, examine the agent log file in the installation directory:

`C:\Program Files (x86)\Citrix\agent\Logs\applayering.agent.log`

You can also view Help for the App Layering agent command line options by running the following command:

`Citrix.AppLayering.Agent.Service.exe /?`

- Run PowerShell command to load the appropriate DLL files on the Citrix Provisioning Server.

For PVS 7.7 and later:

- Navigate to:

`C:\program files\citrix\provisioning services console`

- Run the command:

`C:\Windows\Microsoft.NET\Framework64\v4.0.30319\InstallUtil.exe Citrix.PVS.snapin.dll`

For PVS 7.1–7.6:

- Navigate to:

`C:\program files\citrix\provisioning services console`

- Run the command:

`C:\Windows\Microsoft.NET\Framework64\v4.0.30319\InstallUtil.exe McIPIPSSnapIn.dll`

In case some Citrix Provisioning and Broker snap-ins are unregistered, for instance after a major Windows update, find out which ones.

- Run the following command to list the registered snap-ins:

`get-pssnapin -registered`

- See this [article](#) for instructions to reregister all snap-ins.

Configure

June 13, 2023

Once you've installed the App Layering software and deployed the appliance, you can [access the management console](#) that runs on the appliance.

To secure the appliance, you are required to [set your own administrator passwords](#) the first time you log into the console. Then, you can proceed to:

- [Set up a file share](#)
- [Connect to a directory service](#)
- [Assign roles to users](#)

You also have the opportunity to [enable Labs features](#) currently in testing.

Access management console

September 6, 2025

You can access the App Layering appliance by entering the appliance's IP address in a web browser.

The first time you log into the management console, you:

- Enter the default user name (**administrator**) and password (**Unidesk1**).
- Agree to the Citrix License Agreement.
- Change all administrative passwords on the appliance.

Access the management console

To log directly into the management console hosted on the App Layering appliance:

1. In your hypervisor, locate the virtual machine you created for the appliance, and determine its IP address.
2. Using the IP address for the appliance, enter the following URL in a compatible web browser:
`http://<ip_address_of_new_vm>`
3. Log into the management console. Enter the default user name and password listed in the introduction.

The first time anyone logs into the management console on a new appliance, the system forces a password change for each of the appliance's administrative accounts.

Important:

Be sure to record the new passwords in a safe place, according to your company's security guidelines.

Change administrator passwords

June 14, 2022

The appliance has three accounts that you can use to manage its features and settings.

- **Management console “administrator”account** - Lets you access the management console hosted on the appliance. There you can create and manage layers, and publish layered images. The default password is **Unidesk1**.
- **Appliance “administrator”account** - Lets you access the appliance's configuration utility where you can change the network settings, date, time, NTP server, and time zone. The default password is **Unidesk1**.
- **Appliance “root”user account** - The appliance's default Linux superuser account. The password for this account is required if you ever reset your other administrative accounts. The root user has access to all commands and files on the appliance's Linux OS. The default password is **v9Yx*6uj**.

The administrator account for the management console is the most used. You can easily configure and use the App Layering service without ever accessing the other two accounts.

Important:

Keep the password for the root user in a safe place. If you need it to reset the passwords for the other two administrator accounts.

To secure your appliance, you must change the passwords for these accounts the first time that you access the management console after installation.

The first time you access the appliance after installation

When the appliance is installed and you log in for the first time, a tab is displayed where you *must* change the passwords for the administrator accounts that you use to manage the appliance.

1. For each account, enter the new password and then reenter it in the Confirm Password field.
2. On the Confirm and Complete tab, click **Change Credentials**.
3. Safely store these passwords in case you need them.

WARNING:

You *must* keep the **root** user password in a safe location. Without it you cannot reset your other administrator accounts.

To change the password for an administrative account

1. Log into the management console.
2. Go to **System > Administrators > Default Administrator**.
3. Click the **Edit** button.
4. Enter and confirm the new password, and click **Confirm and Complete**.

Set up file share

January 5, 2026

The App Layering appliance must be connected to a file share. If you haven't yet configured a file share, use these instructions.

You can add storage locations for users' persistent data and settings. You can also add space to storage disks already in use. For more about managing storage, see the article about managing [storage](#).

If you are using more than one appliance, each appliance needs to use a different file share, or a unique folder on the same share.

Requirements

When setting up the appliance's file share:

- The file share must be configured using Server Message Block (SMB) technology.
- The Service account that the App Layering appliance uses to connect to the file share must have **full permissions** for that file share.
- Users require **read-only** access to the file share. If you plan to [enable user layers](#) on the images you publish, also set the file share permissions detailed in [Configure security on user layer folders](#).
- Ensure that you have the minimum storage space requirement of 40-100 GB for your file share.

Note: Storage space is expandable. You can add space to a disk, or add other disks to the appliance.

Create the network file share

Configure a file share that uses the SMB protocol.

- Follow the vendor's instructions for setting up a file share using the SMB protocol.

Configure the App Layering appliance to access the file share

Once you have created a file share, configure the App Layering appliance to attach to it. You can configure the appliance via the App Layering management console.

1. In the Management Console, select **System > Network File Shares** and click **Edit**.
2. Specify an SMB file share path, User name, and Password for the file share.
3. Click **Confirm and Complete** to see if you can connect to the file share. The File Share is saved if the connection succeeds, or displays an error if the connection fails.

Connect to a directory service

September 6, 2025

You can configure the Citrix App Layering™ appliance to connect to Active Directory. When you connect to your directory service, you create one or more Directory Junctions to access specific domains or organizational units (OUs).

The appliance does *not* modify the directory service to which you connect. The software caches the attributes for each directory service entry. If the connection to the directory service is lost temporarily, the software can use the cached information for management tasks.

When creating a Directory Junction, use the following industry standard acronyms:

- OU - Organizational Unit
- DC - Domain Component

About connecting the appliance to a directory service

What happens when you add Directory Junctions

Each Directory Junction that you create specifies a starting node in the directory tree. A new directory junction cannot include users who are already members of another junction. You can't nest junctions.

If you're creating several Distinguished Names

The system compares the Domain Component first; the portions of the Distinguished Name that start with “DC=”.

In Distinguished Names, order matters. For example, DC=A,DC=B is different from DC=B,DC=A.

The system adds Directory Junctions In the following instances:

- The domain components differ.
- Their domain components match and the remaining components do not overlap.

Directory Junctions merge if their domain components match and their other components are related.

User attributes are imported from the directory service

The App Layering software imports and caches user and group attributes from your directory service when:

- You assign administrator privileges to a user.
- The values of the attributes change in the directory service.

The attributes that the software caches are read-only. All changes to the attributes for directory service users come from the directory server.

Imported attributes are synchronized regularly

The software synchronizes the information it caches for directory service users with the directory service every 12 hours. If a user is no longer an object in the directory service, the user is considered abandoned. You can view this information in the Information view for the user.

To create a directory junction

1. Click **System > Directory Services**.
2. Click **Add Directory Junction**.
3. Specify the details for the directory server:
 - **Server address** - The name for the server that you use for the directory service (IP Address or DNS Name).
 - **Port** - Specify the port number for communicating with the directory server.

- **Use SSL** - Click to enable Secure Sockets Layer (SSL) communication. If certificate errors occur, the wizard displays a list of these errors. If you are sure it is safe to ignore them, click **Ignore Certificate Errors**.
- **Connect** - Click to verify so the appliance can connect to the directory service.
- **Bind Distinguished Name (DN)** - To determine the correct syntax for the Bind DN or user name, see the documentation for your directory. The following examples show some of the ways you can specify a user for the directory service:
 - domain\username
 - username@domain.com.
- **Bind Password** - Type the password.
- **Connect** - Click to verify so the appliance can connect to the directory service.
- **Base Distinguished Name** - Specify where the software starts searching for users and groups in the remote directory service.
- **Directory Junction Name** - The name of the folder that you see in the tree view. You can use any name, including the name of a domain in your directory service tree.

4. Click **Confirm and Complete**.

Assign roles

January 5, 2026

App Layering roles define which App Layering modules (features) a user can manage. Any users assigned one or more roles can log into the management console. These users are listed in the **System > Access Control** tab.

App Layering users, rights, and roles

The App Layering service supports two types of users:

- **App Layering administrator.** This account is unique to App Layering. You receive it when you first install the App Layering appliance and log on to the management console. You can use it to get started. This “built-in”administrator account has the rights to perform *all* App Layering operations. You can edit this administrator’s properties, including the name, password, and contact info. The first time you log into the appliance, you are required to change the password for this administrator and agree to the EULA.
- **Active Directory (AD) users.** Other than the built-in administrator account, all App Layering users are AD users imported via one or more directory junctions. Once your directory junctions

have been created, you can grant other rights to users. For more information, see [Connect to a directory service](#).

Rights and roles

You can select all the rights and roles for each user by selecting the following check boxes. If you select **Administrator** or **Reader**, all other check boxes are grayed out.

- **Reader:** Can log into the App Layering management console, but cannot modify anything.
- **App Layer Contributor:** Can add, edit, and delete application layers and versions.
- **OS Layer Contributor:** Can add, edit, and delete OS layers and versions.
- **Platform Layer Contributor:** Can add, edit, and delete platform layers and versions.
- **Image Template Contributor:** Can add, edit, and delete image templates.
- **Image Publisher:** Can publish images.
- **Elastic Layer Assignment Contributor:** Can add, edit, and delete Elastic Layer assignments.
- **Administrator:** Can do everything.
- **Layer Importer:** Can import OS layers, application layers, and platform layers from external sources.
- **Layer Exporter:** Can export OS layers, application layers, and platform layers to an external location.
- **Platform Connector Contributor:** Can add, edit, and delete platform connector configurations.
- **Log Exporter:** Can export and download the App Layering log package.
- **Image Porter:** Can export images from a platform and prepare images for use on a different platform.

User credentials for logging into the management console

When you assign roles to Directory Service users, they can use their Directory Service credentials to log into the management console.

Who can assign App Layering roles?

You can change a user's role if you are logged into the management console as a user assigned the Administrator role.

Assign App Layering roles to users

1. Log in to the management console.

2. Select **System > Access Control**.
3. Select a user and click **Edit**. The Edit User blade opens.
4. Choose the role(s) for the selected user.
5. Click **Confirm and Complete**, then click **Save**.

Connector configurations

January 5, 2026

The first time you use the App Layering service, plan to create one or more “connector configurations”. App Layering connector configurations are stored sets of credentials that the appliance uses to access locations on your hypervisor or provisioning server. A connector configuration can specify where to build a layer, or where to publish images.

Why use connector configurations

Connector configurations let the App Layering appliance access locations in your environment for creating layers or publishing layered images. Using a connector configuration automates the file transfer process and saves significant amounts of time.

You can use a connector configuration to:

- Create an app layer or platform layer, or add a version to an OS layer, and if you choose, also run a script.
- Publish layered images to your hypervisor or provisioning server, and if you choose, also run a script.

Create layers

Layer creation is easier when you use a connector configuration. The connector configuration includes the credentials for the location where you plan to install the software for a layer. This connector configuration includes a Packaging Cache option, which is enabled by default to afford you the best layering performance.

Publish layered images

You need a connector configuration for each location to which you are publishing layered images. Once published, you use layered images to provision systems for specific groups of users.

Requirements

To create a connector configuration, you need to meet the following requirements.

Credentials

Valid account credentials that the appliance can use to access a location in your environment. For details about the values you need, select the hypervisor or Provisioning Service later in this section.

About offload compositing

When creating a layer or publishing a layered image, you can use the built-in App Layering *compositing engine* to create layers. The compositing engine is enabled by selecting the **Offload Compositing** connector configuration option. The `ImportOsLayer.ps1` script is used when creating an OS layer from scratch. Be sure to open the firewall ports required for the [compositing engine](#).

Note:

An OS layer is created (initially) by importing the OS from an existing virtual machine. Once created, you update the OS layer just like any other layer.

Offload Compositing:

- Reduces the layering task processing time.
- Enables support for creating UEFI images with vTPM and GPT partitions. These features are required for modern Windows OS versions such as Windows 11.
- Enables support for creating VHDX disks.
- Can automatically eliminate some file system issues with native Windows tools

To create a layer using Offload Compositing:

- **App layer, platform layer:** While creating a layer, choose a connector configuration with **Offload Compositing** enabled.
- **OS layer:** Run a script. For details, see [Create OS layer](#).

Disk space for cache

On all supported hypervisors, except for Azure, the default *Packaging cache size* is set to the recommended starting level. Allow enough disk space to increase the cache size, if needed.

- Creates the layer or image as a UEFI or Generation 2 machine.

- Uses VHDX disk format for either BIOS and Generation 1 format, or UEFI and Generation 2 images.

To create an app or platform layer using Offload Compositing, you select a connector configuration with **Offload Compositing** enabled. To create an OS layer, you run a script instead. For details, see [Create OS layer](#).

Packaging layers

When using **Offload Compositing** in the connector configuration, packaging a layer begins when you shut down the machine for finalization. The packaging is done automatically. You do not have to manually select **Finalize** in the management console.

Important:

You must run the script to create an OS layer. You can also build the OS on a BIOS or Generation 1 machine, then add a version to the layer with **Offload Compositing** selected.

When you select **Offload Compositing**, choose either **UEFI** or **Generation 2** for the new layer version. Select **VHDX** for the disk format. VHDX format is supported on BIOS (Generation 1) and UEFI (Generation 2) machines.

Publish images

When you publish an image, the publishing machine builds the image on the hypervisor (Hyper-V, VMware vSphere) server.

Connectivity

When you use Offload Compositing, a temporary worker VM is created in your environment, which is called a compositing engine. The compositing engine requires direct connectivity to the App Layering appliance over SSL (port 443) and iSCSI (port 3260). Ensure that this traffic is permitted within your App Layering environment. Otherwise, tasks created from the appliance are not completed successfully.

About “Packaging cache size” and “Select rate”

Connector configurations for all hypervisors except Azure let you configure space for a *Packaging cache*. The appliance uses this cache on your hypervisor to speed up layer packaging.

Packaging cache size

We recommend using the default packaging cache starting size for your hypervisor:

- vSphere: 250 GB
- XenServer®: 480 GB
- Hyper-V: 200 GB
- Nutanix: 480 GB

The more app layers you create, the larger the cache you need. Increase the [Packaging cache size](#), if needed.

The Hit rate

The [Hit](#) rate is the percentage of times that the appliance has found a disk in the cache. A low value in this field indicates that the cache is not providing much value. Increase the cache size to accommodate more unique layers.

You can improve a low rate by increasing the Packaging cache size. Increasing the cache size allows more disks to be stored in the cache. It also increases the likelihood of finding a disk for packaging in the cache. The result is a higher [hit](#) rate value.

View connector configurations

The **Connectors** tab lists the connector configurations that you have created.

Sort connector configurations

By default, connector configurations are listed in alphabetical order by name.

View connector configuration details

When caching is enabled for a connector configuration, the **Connectors** tab displays the cache size for that configuration.

To see the values for any given connector configuration:

1. Click the **View Details** button or the info icon on the right side of the screen. Details about the connector configuration are displayed.

View Cache size and Hit rate

The Packaging Cache Size and [Hit](#) rate are displayed when the Packaging Cache Size is set to a value greater than zero (0).

Add a connector configuration

This section explains how to add and manage a connector configuration on the **Connectors** tab.

A connector configuration speeds up the process of creating a layer or image template. If you need one to access a specific location, you can create one by selecting **Add Connector Configuration** on the **Connectors** tab. From there, use the detailed steps later in this section.

To add a connector configuration, follow these steps:

1. Access the management console, and select **Connectors** to see your connector configurations, if you have created any.
2. Click **Add Connector Configuration** in the Action bar. A dialog box opens where you can select the type of connector configuration to create.
3. Choose the connector type from the drop-down menu.
4. Complete the fields on the connector configuration page.
5. Click the **Test** button to verify that the connector can access the location specified using the credentials supplied.
6. Click **Save**. The new connector configuration appears on the **Connector** tab.

Name	Platform	Cache Size	Cache Hit Rate
(i) BIG Cache	VMware vSphere	2147483647 GB	<div style="width: 0%;"></div>
(i) ESXP14 Slow	Citrix MCS for vSphere		<div style="width: 0%;"></div>
(i) Fitx App Layering Dev	Azure RDSH		<div style="width: 0%;"></div>
(i) SSDESPX13	VMware vSphere	200 GB	<div style="width: 5%;"></div>
(i) UD-Hyperv2	Microsoft Hyper-V		<div style="width: 0%;"></div>
(i) vcenter6 - QA	VMware vSphere		<div style="width: 0%;"></div>
(i) XenServer2	XenServer		<div style="width: 0%;"></div>

Edit a connector configuration

To edit a connector configuration:

1. Access the management console, and select **Connectors**.
2. Select the connector and click **Edit** on the action panel or the right side of the screen. The connector configuration is displayed.
3. Update the configuration, as needed.
4. Click the **Test** button to verify that the connector can access the location specified using the credentials supplied.

5. Click **Save**. The connector configuration is updated.

Increase the Cache size and Hit rate

To improve the utility of the cache and therefore increase the [hit](#) rate by editing the Packaging Cache Size:

1. Select the connector and click **Edit** on the action panel or on the right side of the screen.
2. Scroll to **Layer Disk Cache Size in GB** and enter the amount of space that the cache can occupy.
3. Click the **Test** button to verify that the connector can access the location specified by using the credentials supplied.
4. Click **Save**. The connector configuration is updated.

Continue adjusting the cache size until you achieve the layering performance you are looking for.

Disable or re-enable caching

Caching is *enabled* and set to a recommended starting size by default. We strongly recommend that you use caching.

If you disable caching for a connector configuration, set the cache size to zero (0). You can re-enable it by increasing the Packaging cache size.

Delete a connector configuration

To delete a connector configuration:

1. In the management console, select **Connectors**.
2. Make sure that the connector configuration is not in use.
3. Click **Delete** on the Action bar or the right side of the screen.
4. In the pop-up window that opens, select **Delete**. The connector configuration is deleted.

The **Connectors** tab takes a few minutes to update after another administrator deletes a connector configuration.

Messages when deleting connector configurations

If a connector configuration is in use when you attempt to delete it, you receive a message similar to the following example:

“Validation Error: Unable to delete connector configuration ‘Citrix Provisioning™ - ConnectorExample’ as it is in use. If you receive this error message, remove the connector configuration from the layer or image template where it is still in use. Then delete the configuration.”

Azure

January 5, 2026

The articles in this section outline the different connectors available for your Azure environment, as well as required fields for all Azure connector configurations.

Connectors

Package layers

For creating and versioning layers in an Azure environment, use the [Azure Deployments](#) connector.

Publish layered images

For publishing layered image into an Azure environment, the [Azure Deployments](#) connector and [Machine Creation for Azure Deployments](#) connector are available.

For publishing images to PVS for use with Azure targets, use the [Citrix Provisioning](#) connector configured with [Azure Deployments](#) as its offloading connector configuration.

Image Portability

For performing IPS tasks in Azure, use the [Image Portability for Azure](#) connector.

Required fields

The following is information on required fields that are shared across all available Azure connectors.

Identity

Identity specifies how to authenticate with your Azure environment.

There are two ways to configure Azure connectors to authenticate with Azure: using a **managed identity** or a **registered application**. While using a managed identity is a convenient way to grant rights to the appliance, using app registration credentials enables the connector to access resources across tenants and to be configured on appliances that don't reside in Azure.

Managed identity With the Azure connectors, you are able to authenticate to Azure with the [managed identity](#) assigned to the App Layering appliance in Azure.

Since a managed identity must be [assigned](#), this method is only supported on an App Layering appliance deployed in Azure.

Registered application A [registered application](#) establishes a trust relationship between your App layering appliance and your Azure environment. To authenticate using your registration credentials, the Azure connector configurations require the following information:

- **Azure Environment** - The environment being used, whether Azure Public Cloud or Azure Government.
- **Tenant ID** - An Azure Active Directory instance, this GUID identifies your organization's dedicated instance of Azure Active Directory (AD).
- **Client ID** - An identifier for the App Registration, which your organization has created for App Layering.
- **Client Secret** - The password for the Client ID you are using. If you have forgotten the Client Secret, you can create a new one.

Azure Deployments

January 5, 2026

When creating layers or publishing images in an Azure environment, use the **Azure Deployments** or **Machine Creation for Azure Deployments** connector configuration. This article describes the connector configuration settings. For more information about connector configurations and how to add new ones, see [Connector configurations](#).

Note:

Azure Deployments also supports Azure Government.

Overview

Azure Deployments in App Layering refer to the creation of Azure deployments using Azure Resource Manager (ARM) templates. ARM templates are Azure-specific JSON documents that define infrastructure and configuration as code. For more information on ARM templates, refer to the Azure documentation [here](#).

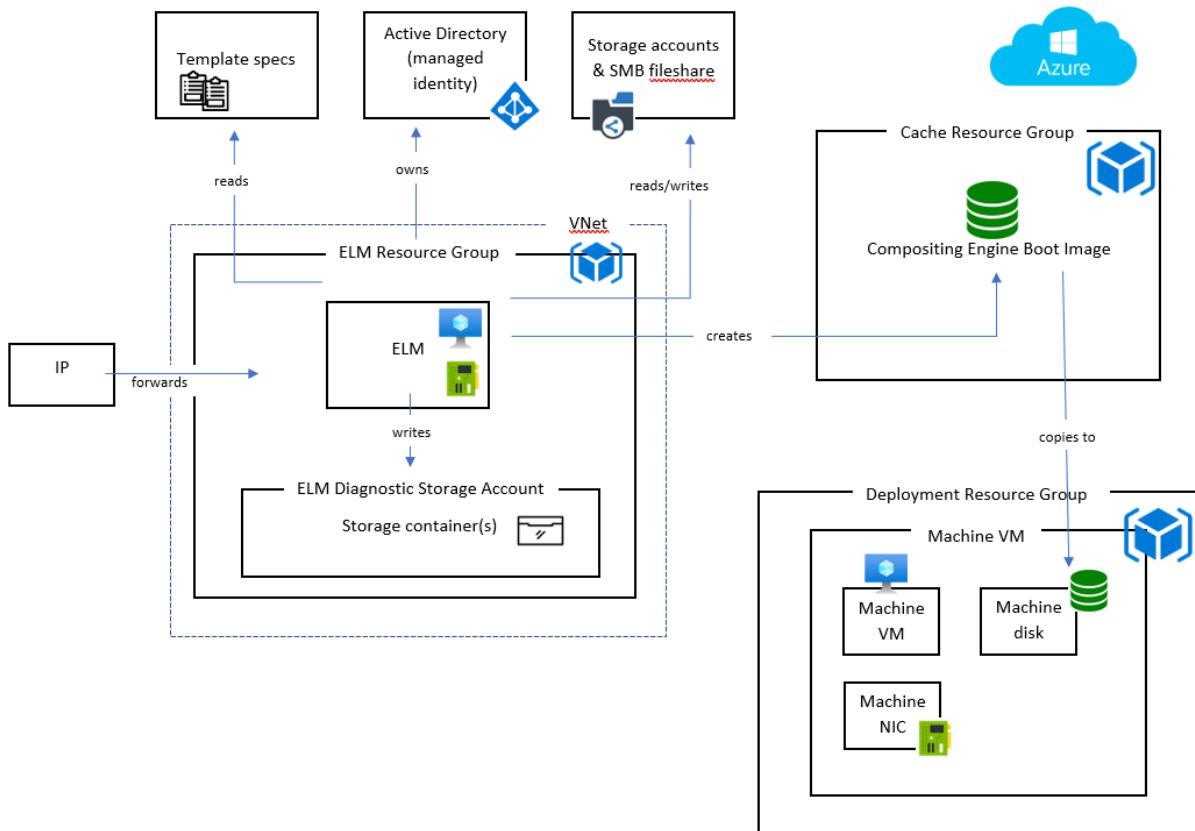
All Azure resources created by the App Layering Azure Deployments connector are created using the deployment of a user-specified ARM template. These templates allow an administrator to extensively customize what resources are created and how they're configured.

Identity

Information on the **Identity** field can be found on our [Azure](#) page.

Azure template specs

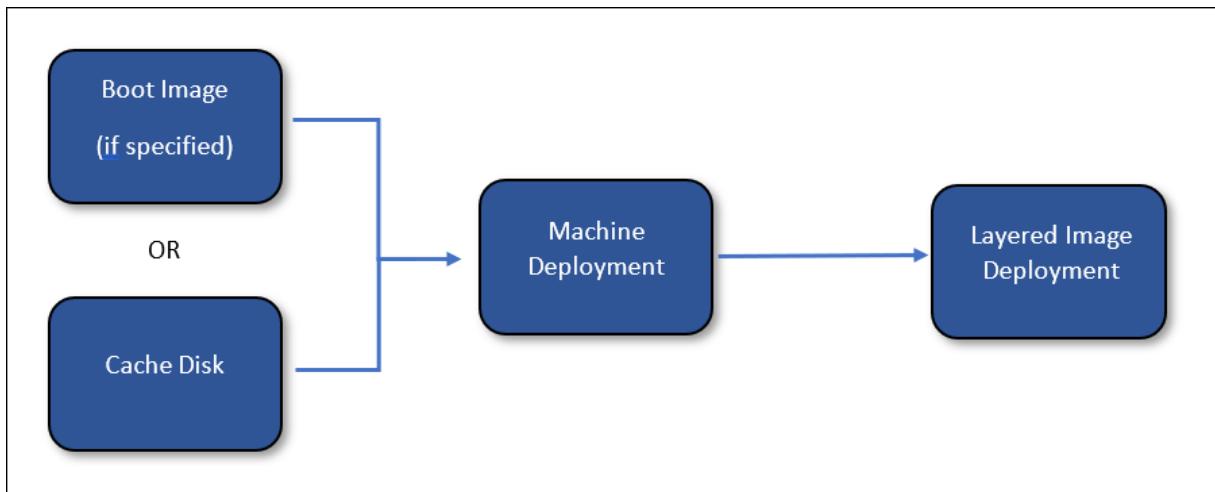
[Azure template specs](#) are a type of Azure Resource that store and version an ARM Template for later use in an ARM template deployment. You must specify between two to four template specs for each Azure Deployments connector configuration. Each deployment type in an Azure Deployments connector requires a corresponding version of a template spec. The **Cache Disk** and **Machine** deployments are required, but the **Boot Image** and **Layered Image** deployments are optional.



Deployment types

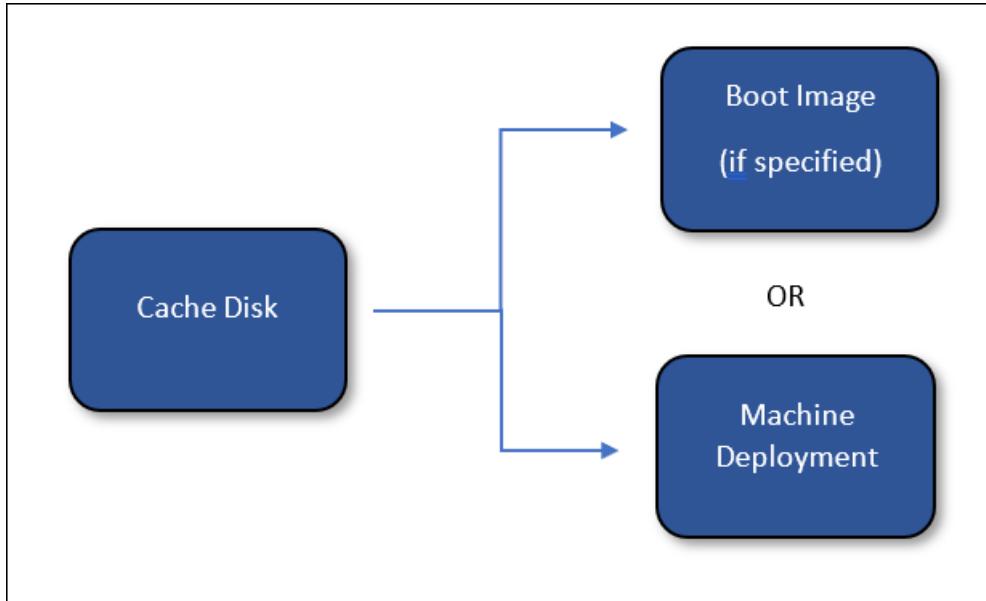
There are four deployment types, each requiring its template spec. The deployment types differ in the type of resources that they create, the inputs they receive, and the outputs they produce to override the default behavior. For more information on these concepts, see [Authoring ARM templates](#).

Machine The **Machine** deployment creates a virtual machine (VM). VMs created by **Machine** deployments can composite layered images and package layers. If the optional **Layered Image** deployment isn't specified, then a VM is the final result of publishing an image. In this case, the VM can be used as-is or as a Machine Creation Services™ (MCS) master image.



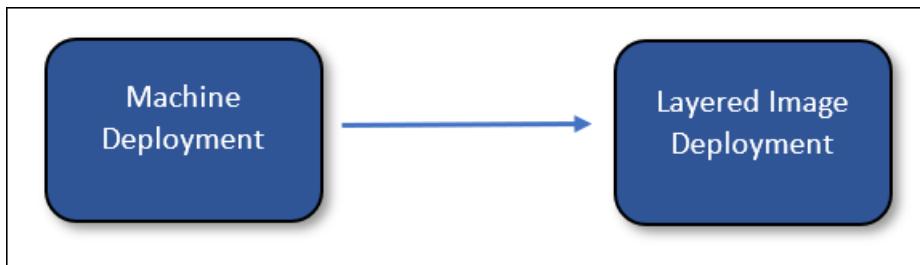
If a **Layered Image** deployment is specified, then the resources created by the **Machine** deployment are deleted after the **Layered Image** deployment completes. Otherwise, App Layering does not delete the resources (unless the deployment fails).

Cache disk The **Cache Disk** deployment creates an Azure-managed disk. This disk is used to contain the Compositing Engine boot image. The App Layering appliance uploads the contents to the disk after it's created.



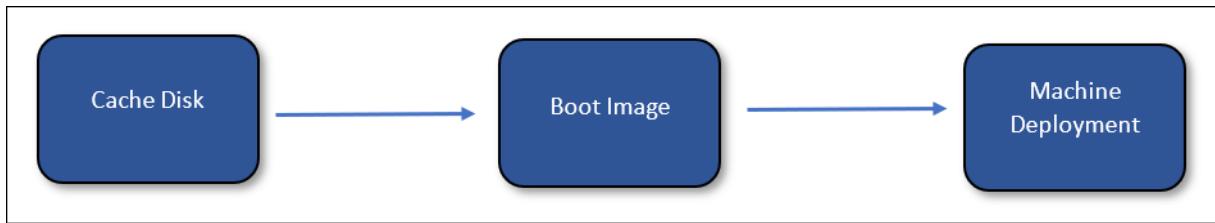
If a **Boot Image** deployment is specified, then the resources created by the **Cache Disk** deployment are deleted after the **Boot Image** deployment completes. Otherwise, App Layering deletes the resources during cache cleanup.

Layered image (optional) The **Layered Image** deployment is an optional deployment type. The resulting resources are the final result of publishing a layered image. No particular resource type is required to be created. The **Layered Image** deployment can be used to produce a [compute gallery image](#), a managed disk, or any other type of resource.



App Layering doesn't delete the resources created by the layered image deployment (unless the deployment fails).

Boot image (optional) The **Boot Image** deployment is an optional deployment type. The resulting resources are used to create the OS disks of the VMs created by **Machine** deployments. It isn't required to create any particular type of resource, however it must create a resource that can be used to create an OS disk for a VM. This deployment can be used to produce a [compute gallery image](#), or any other type of resource that can be used as the source of a disk.



App Layering deletes the resources created by the boot image deployment during cache cleanup.

Add a connector configuration

Refer to the following descriptions for information on each field in the connector configuration screen.

Defaults (optional)

Defaults are optional and can be used to apply **Tags** and **Custom Data** to all deployment types in the connector configuration. The data specified in the defaults are merged with the data specified in the corresponding fields of each deployment type.

Deployments

Each deployment type contains the following fields.

Template The template spec used for the deployment. The user creates and manages template specs in Azure. The managed identity or the registration credentials of the appliance must have read permission on the resource group containing the template spec.

Version The version of the template spec to use for deployment. The newest version is selected by default when the **Template** selection is changed.

Resource Group The Azure resource group to which to deploy. All resources created by the deployment are created in this resource group.

The managed identity or registration credentials of the appliance must have permission to:

- Deploy templates to the resource group
- Create each type of resource in the template
- Delete each type of resource in the template

Assigning the [general Contributor role](#) to the managed identity or registration credentials on the resource group grants the required permissions. More granular permissions or roles can be specified instead, but the permissions required depend on the resources specified in the template.

Note:

The managed identity of the appliance must have permission to connect a device to the virtual network specified for the VM created by the **Machine** deployment. If the virtual network isn't located in any of the resource groups specified for the deployments, then roles assigned to the managed identity or registration credentials for those resource groups won't apply to the virtual network and a role must be assigned directly to the virtual network.

Tags (advanced) The tags to apply to the Azure deployment artifact created by a deployment. You can include data from the input parameter (see [Authoring ARM templates](#)). To do so, put the JSON path of the field you want to reference between braces. For example, `{ context.user }` evaluates to the name of the App Layering user who created the task that caused the deployment. This works for both the tag name and the tag value fields. If you want to use literal braces in your tags, you can escape them by doubling them. For example, evaluates to `{}{}`.

Important:

These tags are *only* applied to the deployment resource itself. They aren't applied to the resources created by the deployment. To apply tags to the created resources, specify those tags in the ARM template.

Custom Data (advanced) Arbitrary data specified in JSON format. This data can be referenced in the ARM template associated with the deployment. The data is accessed using the input parameter's `context.config.custom` object. See [Authoring ARM templates](#) for more details.

Considerations

- Network connectivity is required between the App Layering appliance and the VMs created by the Machine deployment.
 - From the VMs created by the Machine deployment, IP traffic must be routable to the App Layering appliance's IP address on ports 443 (HTTPS) and 3260 (iSCSI). Also, App Layering appliance traffic must be routable to these VM's (created by the Machine deployment) IP address on port 443 (HTTPS).
 - App Layering appliances deployed on-premises must be connected to the Azure virtual network specified for VMs created by the Machine deployment. You can connect on-premises virtual networks to Azure through [Azure ExpressRoute](#) and [Azure VPN Gateway](#).

- The Legacy Azure connectors are deprecated but are still available for a limited time. There's no upgrade nor migration path from any Legacy Azure connectors to the new Azure Deployments connector types.

For more information, see [Citrix App Layering in Azure](#).

Authoring ARM templates

January 5, 2026

This section is intended for users familiar with ARM templates. It provides detailed information about authoring templates for the App Layering Azure Deployments connector. For general ARM template authoring information, see the [Microsoft documentation](#).

Input

Every deployment type is passed one parameter, an object named `al`. This object has two properties, `input` and `context`. The `input` property is an object specific to each deployment type and its properties change depending on the deployment type. The `context` property is the same for all deployment types. It contains data about the App Layering task, item (layer or image template), and connector configuration associated with the current deployment. For detailed information about the parameter object, see [Azure Deployments Template Parameter](#).

Each template must declare the `al` parameter in its parameters section, as follows:

```
1  {
2
3  ...
4  "parameters": {
5
6      "al": {
7
8          "type": "object"
9      }
10
11  }
12  ,
13  ...
14 }
```

A template can declare more parameters, but the parameters must all have default values. Otherwise, App Layering doesn't provide value to them. This can be useful for using functions that can only be used in the default value section of a parameter, for example `utcNow`.

Output

All ARM templates can have outputs. With the Azure Deployments connector, template outputs can be used to pass information to the next deployment. They can also be used to override some default behaviors.

The outputs of a deployment are passed to the next deployment using the `input` property of the [all template parameter](#).

For example, when a **Cache Disk** deployment has these outputs:

```
1  {
2
3  ...
4  "outputs": {
5
6      "generation": {
7
8          "type": "string",
9          "value": "[variables('generation')]"
10         }
11     ,
12     "name": {
13
14         "type": "string",
15         "value": "[variables('name')]"
16         }
17     ,
18     }
19
20     ...
21 }
```

The **Boot Image** deployment receives this input:

```
1  {
2
3  "input": {
4
5      "type": "BootImage",
6      "source": {
7
8          "generation": "V2",
9          "name": "MyCoolDiskName"
10         }
11     }
12
13     }
14     ,
15     "context": {
16
17     ...
18 }
```

```
18     }
19
20 }
```

Notice the `source` property of the `input` object has a property for each output specified by the **Cache Disk** deployment template. The origins of each output depend on the type of deployment.

Deployment type details

Each deployment type has a different set of inputs and outputs that can change the behavior of the App Layering operation. These deployment-specific details are described in this section.

For real-world examples that use all of these concepts, see [Starter Templates](#).

Cache Disk

The Cache Disk deployment must create a managed disk resource. You can optionally create other resources in addition to the disk. The App Layering appliance must have permission to write to the disk using a SAS token (generated by the appliance). A boot image containing the App Layering compositing engine is uploaded to the disk after it's created.

Cache Disk Requirements

- Must create a managed disk resource
- The managed disk's `createOption` must be set to `"Upload"`
- The managed disk's `uploadSizeBytes` must be set to the `uploadSize` specified by the `input`, such as `"[parameters('al').input.uploadSize]"`
- The App Layering appliance must be able to write to the managed disk using a SAS token

```
1 {
2
3 ...
4     "resources": [
5         {
6             "type": "Microsoft.Compute/disks",
7             ...
8             "properties": {
9                 ...
10                "creationData": {
11                    "createOption": "Upload",
12                    "uploadSizeBytes": "[parameters('al').input.
13                                  uploadSize]"
14                }
15            }
16        }
17 }
```

```
17
18
19     }
20
21     }
22
23   ]
24   ...
25 }
```

Cache Disk Input The [input object](#) includes the `size` and `uploadSize` properties. This object does not include output from another deployment.

Cache Disk Output The output of the deployment is passed to the **Boot Image** deployment if one is specified. Otherwise, it's passed to the **Machine** deployment.

An output named `diskId` can be specified to explicitly tell App Layering which disk to use. If no `diskId` output is specified, App Layering automatically adds one and sets it to the resource ID of the first managed disk resource created by the deployment. The disk specified by `diskId` has the App Layering compositing engine boot image uploaded to it.

Boot Image

This deployment creates a resource from the managed disk created by the **Cache Disk** deployment. There are no hard requirements as to what type of resources are created. However, it must create a resource that can be used as the source of an OS disk when creating a VM, such as a compute gallery image version.

Boot Image Disk Requirements

- Must create a resource that can be used to create the OS disk of a VM with the same contents as the disk with the ID passed in as input.

As an example, a compute gallery image version using the input `diskId` as a source:

```
1  {
2
3     ...
4     "resources": [
5         {
6
7             "type": "Microsoft.Compute/galleries/images/versions",
8             ...
9             "properties": {
```

```
11          ...
12          "storageProfile": {
13
14              "osDiskImage": {
15
16                  "source": {
17
18                      "id": "[parameters('al').input.source.
19                          diskId]"
20
21                  }
22
23              }
24
25          ...
26      }
27
28  }
29
30  ]
31  ...
32 }
```

Boot Image Input The [input object](#) includes the [source](#) property. The [source](#) represents the outputs of the **Cache Disk** deployment, with a property for each output specified. Use the [diskId](#) property for the source of the resource being created.

Boot Image Output The output of the **Boot Image** deployment is passed to the **Machine** deployment. There are no special or required outputs. However, you must include the data required to create a VM from the created resource, like a resource ID.

Machine

The **Machine** deployment must create a virtual machine resource. The virtual machine must be attached a network on which it can reach the App Layering appliance and the other way around, as per [Firewall ports internal connections for Compositing machine](#).

Important:

Don't attach the disk created by the **Cache Disk** deployment to the virtual machine. The **Cache Disk** is a shared resource and considered read only. Create a copy of the disk and attach that instead when not using the **Boot Image** deployment.

Machine Requirements

- Must create a virtual machine resource
- The virtual machine must be attached to a network that allows communication to and from the App Layering appliance
- The virtual machine's OS disk must be created using **Boot Image** or **Cache Disk** resource as its source
- The virtual machine's OS disk size must be set to "[parameters('al').input.disk.size]"
- The virtual machine's `userData` property must be set to "[parameters('al').input.vm.userData]"

```
1  {
2
3  ...
4  "resources": [
5  {
6
7      "type": "Microsoft.Compute/disks",
8      "name": "[variables('diskName')]",
9      ...
10     "properties": {
11
12         ...
13         "creationData": {
14
15             "createOption": "Copy",
16             "sourceResourceId": "[parameters('al').input.disk.
17                 image.diskId]"
18         }
19         ,
20         "diskSizeGB": "[parameters('al').input.disk.size]",
21         ...
22     }
23     ,
24     {
25
26         "type": "Microsoft.Compute/virtualMachines",
27         ...
28
29         "dependsOn": [
30             "[resourceId('Microsoft.Compute/disks', variables('
31                 diskName'))]"
32         ],
33         ...
34
35         "properties": {
36
37             ...
38             "storageProfile": {
39                 "osDisk": {
```

```
40          ...
41          "createOption": "Attach",
42          "managedDisk": {
43
44              "id": "[resourceId('Microsoft.Compute/disks
45                  ', variables('diskName'))]"
46          }
47      }
48      ,
49      "dataDisks": []
50  }
51  ,
52  ...
53  "userData": "[parameters('a1').input.vm.userData]"
54  ...
55  }
56
57  }
58  ]
59  ...
60  ...
61 }
```

Machine Input The [input object](#) includes the `disk` and `vm` properties.

The `disk.image` property contains the output from the **Boot Image** deployment if one was specified. Otherwise, it contains the output of the **Cache Disk** deployment. The `disk.size` property contains the size of the disk in GB.

The `vm.userData` property contains the user data that must be assigned to the created virtual machine.

Machine Output The output of the **Machine** deployment is passed to the **Layered Image** deployment if one is specified. If you're using a **Layered Image** deployment, then you must include the ID of the VM or OS disk in the output so it can be referenced by the **Layered Image** deployment.

An output named `machineId` can be specified to explicitly tell App Layering which virtual machine to use. If no `machineId` output is specified, App Layering automatically adds one and sets it to the resource ID of the first virtual machine resource created by the deployment.

An output named `ipAddress` can be specified to explicitly tell App Layering which IP address to use to communicate with the machine. If no `ipAddress` output is specified, App Layering uses the primary private address of the primary network card attached to the virtual machine resource.

An output named `message` can be specified to provide a message that is appended to the final status of a publish image task and the action required status of a layer creation task in the App Layering UI.

This message is only used in the final status of the image publish task if a **Layered Image** deployment isn't specified.

- The [Machine Starter Template](#) sets the message output parameter to a link to the machine in the Azure portal.

Layered Image

The **Layered Image** deployment creates a resource from the virtual machine or other resources created by the **Machine** deployment. There are no hard requirements as to what type of resources are created. However, it creates a resource that can be used as an input to a Provisioning Service such as Machine Creation Services™ (MCS). A compute gallery image resource is a good example.

Layered Image Disk Requirements

- Creates a resource that can be used by a Provisioning Service to create virtual machines.

As an example, the following code block creates a compute gallery image version using the input `diskId` as a source. This assumes that the **Machine** deployment included an output named `diskId` that is set to the ID of the machine's OS disk:

```
1  {
2
3  ...
4  "resources": [
5      {
6
7          "type": "Microsoft.Compute/galleries/images",
8          "name": "[format('{
9  0 })",
10         '/',
11         1 }",
12         ', variables('galleryName'), variables('name'))]",
13         ...
14         "resources": [
15             {
16
17                 "type": "versions",
18                 ...
19                 "dependsOn": [
20                     "[resourceId('Microsoft.Compute/galleries/
21                         images', variables('galleryName'), variables
22                         ('name'))]"
23                 ],
24                 ...
25                 "properties": {
26
27                     ...
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
59
60
61
62
63
64
65
66
67
68
69
69
70
71
72
73
74
75
76
77
78
79
79
80
81
82
83
84
85
86
87
87
88
89
89
90
91
91
92
92
93
93
94
94
95
95
96
96
97
97
98
98
99
99
100
100
101
101
102
102
103
103
104
104
105
105
106
106
107
107
108
108
109
109
110
110
111
111
112
112
113
113
114
114
115
115
116
116
117
117
118
118
119
119
120
120
121
121
122
122
123
123
124
124
125
125
126
126
127
127
128
128
129
129
130
130
131
131
132
132
133
133
134
134
135
135
136
136
137
137
138
138
139
139
140
140
141
141
142
142
143
143
144
144
145
145
146
146
147
147
148
148
149
149
150
150
151
151
152
152
153
153
154
154
155
155
156
156
157
157
158
158
159
159
160
160
161
161
162
162
163
163
164
164
165
165
166
166
167
167
168
168
169
169
170
170
171
171
172
172
173
173
174
174
175
175
176
176
177
177
178
178
179
179
180
180
181
181
182
182
183
183
184
184
185
185
186
186
187
187
188
188
189
189
190
190
191
191
192
192
193
193
194
194
195
195
196
196
197
197
198
198
199
199
200
200
201
201
202
202
203
203
204
204
205
205
206
206
207
207
208
208
209
209
210
210
211
211
212
212
213
213
214
214
215
215
216
216
217
217
218
218
219
219
220
220
221
221
222
222
223
223
224
224
225
225
226
226
227
227
228
228
229
229
230
230
231
231
232
232
233
233
234
234
235
235
236
236
237
237
238
238
239
239
240
240
241
241
242
242
243
243
244
244
245
245
246
246
247
247
248
248
249
249
250
250
251
251
252
252
253
253
254
254
255
255
256
256
257
257
258
258
259
259
260
260
261
261
262
262
263
263
264
264
265
265
266
266
267
267
268
268
269
269
270
270
271
271
272
272
273
273
274
274
275
275
276
276
277
277
278
278
279
279
280
280
281
281
282
282
283
283
284
284
285
285
286
286
287
287
288
288
289
289
290
290
291
291
292
292
293
293
294
294
295
295
296
296
297
297
298
298
299
299
300
300
301
301
302
302
303
303
304
304
305
305
306
306
307
307
308
308
309
309
310
310
311
311
312
312
313
313
314
314
315
315
316
316
317
317
318
318
319
319
320
320
321
321
322
322
323
323
324
324
325
325
326
326
327
327
328
328
329
329
330
330
331
331
332
332
333
333
334
334
335
335
336
336
337
337
338
338
339
339
340
340
341
341
342
342
343
343
344
344
345
345
346
346
347
347
348
348
349
349
350
350
351
351
352
352
353
353
354
354
355
355
356
356
357
357
358
358
359
359
360
360
361
361
362
362
363
363
364
364
365
365
366
366
367
367
368
368
369
369
370
370
371
371
372
372
373
373
374
374
375
375
376
376
377
377
378
378
379
379
380
380
381
381
382
382
383
383
384
384
385
385
386
386
387
387
388
388
389
389
390
390
391
391
392
392
393
393
394
394
395
395
396
396
397
397
398
398
399
399
400
400
401
401
402
402
403
403
404
404
405
405
406
406
407
407
408
408
409
409
410
410
411
411
412
412
413
413
414
414
415
415
416
416
417
417
418
418
419
419
420
420
421
421
422
422
423
423
424
424
425
425
426
426
427
427
428
428
429
429
430
430
431
431
432
432
433
433
434
434
435
435
436
436
437
437
438
438
439
439
440
440
441
441
442
442
443
443
444
444
445
445
446
446
447
447
448
448
449
449
450
450
451
451
452
452
453
453
454
454
455
455
456
456
457
457
458
458
459
459
460
460
461
461
462
462
463
463
464
464
465
465
466
466
467
467
468
468
469
469
470
470
471
471
472
472
473
473
474
474
475
475
476
476
477
477
478
478
479
479
480
480
481
481
482
482
483
483
484
484
485
485
486
486
487
487
488
488
489
489
490
490
491
491
492
492
493
493
494
494
495
495
496
496
497
497
498
498
499
499
500
500
501
501
502
502
503
503
504
504
505
505
506
506
507
507
508
508
509
509
510
510
511
511
512
512
513
513
514
514
515
515
516
516
517
517
518
518
519
519
520
520
521
521
522
522
523
523
524
524
525
525
526
526
527
527
528
528
529
529
530
530
531
531
532
532
533
533
534
534
535
535
536
536
537
537
538
538
539
539
540
540
541
541
542
542
543
543
544
544
545
545
546
546
547
547
548
548
549
549
550
550
551
551
552
552
553
553
554
554
555
555
556
556
557
557
558
558
559
559
560
560
561
561
562
562
563
563
564
564
565
565
566
566
567
567
568
568
569
569
570
570
571
571
572
572
573
573
574
574
575
575
576
576
577
577
578
578
579
579
580
580
581
581
582
582
583
583
584
584
585
585
586
586
587
587
588
588
589
589
590
590
591
591
592
592
593
593
594
594
595
595
596
596
597
597
598
598
599
599
600
600
601
601
602
602
603
603
604
604
605
605
606
606
607
607
608
608
609
609
610
610
611
611
612
612
613
613
614
614
615
615
616
616
617
617
618
618
619
619
620
620
621
621
622
622
623
623
624
624
625
625
626
626
627
627
628
628
629
629
630
630
631
631
632
632
633
633
634
634
635
635
636
636
637
637
638
638
639
639
640
640
641
641
642
642
643
643
644
644
645
645
646
646
647
647
648
648
649
649
650
650
651
651
652
652
653
653
654
654
655
655
656
656
657
657
658
658
659
659
660
660
661
661
662
662
663
663
664
664
665
665
666
666
667
667
668
668
669
669
670
670
671
671
672
672
673
673
674
674
675
675
676
676
677
677
678
678
679
679
680
680
681
681
682
682
683
683
684
684
685
685
686
686
687
687
688
688
689
689
690
690
691
691
692
692
693
693
694
694
695
695
696
696
697
697
698
698
699
699
700
700
701
701
702
702
703
703
704
704
705
705
706
706
707
707
708
708
709
709
710
710
711
711
712
712
713
713
714
714
715
715
716
716
717
717
718
718
719
719
720
720
721
721
722
722
723
723
724
724
725
725
726
726
727
727
728
728
729
729
730
730
731
731
732
732
733
733
734
734
735
735
736
736
737
737
738
738
739
739
740
740
741
741
742
742
743
743
744
744
745
745
746
746
747
747
748
748
749
749
750
750
751
751
752
752
753
753
754
754
755
755
756
756
757
757
758
758
759
759
760
760
761
761
762
762
763
763
764
764
765
765
766
766
767
767
768
768
769
769
770
770
771
771
772
772
773
773
774
774
775
775
776
776
777
777
778
778
779
779
780
780
781
781
782
782
783
783
784
784
785
785
786
786
787
787
788
788
789
789
790
790
791
791
792
792
793
793
794
794
795
795
796
796
797
797
798
798
799
799
800
800
801
801
802
802
803
803
804
804
805
805
806
806
807
807
808
808
809
809
810
810
811
811
812
812
813
813
814
814
815
815
816
816
817
817
818
818
819
819
820
820
821
821
822
822
823
823
824
824
825
825
826
826
827
827
828
828
829
829
830
830
831
831
832
832
833
833
834
834
835
835
836
836
837
837
838
838
839
839
840
840
841
841
842
842
843
843
844
844
845
845
846
846
847
847
848
848
849
849
850
850
851
851
852
852
853
853
854
854
855
855
856
856
857
857
858
858
859
859
860
860
861
861
862
862
863
863
864
864
865
865
866
866
867
867
868
868
869
869
870
870
871
871
872
872
873
873
874
874
875
875
876
876
877
877
878
878
879
879
880
880
881
881
882
882
883
883
884
884
885
885
886
886
887
887
888
888
889
889
890
890
891
891
892
892
893
893
894
894
895
895
896
896
897
897
898
898
899
899
900
900
901
901
902
902
903
903
904
904
905
905
906
906
907
907
908
908
909
909
910
910
911
911
912
912
913
913
914
914
915
915
916
916
917
917
918
918
919
919
920
920
921
921
922
922
923
923
924
924
925
925
926
926
927
927
928
928
929
929
930
930
931
931
932
932
933
933
934
934
935
935
936
936
937
937
938
938
939
939
940
940
941
941
942
942
943
943
944
944
945
945
946
946
947
947
948
948
949
949
950
950
951
951
952
952
953
953
954
954
955
955
956
956
957
957
958
958
959
959
960
960
961
961
962
962
963
963
964
964
965
965
966
966
967
967
968
968
969
969
970
970
971
971
972
972
973
973
974
974
975
975
976
976
977
977
978
978
979
979
980
980
981
981
982
982
983
983
984
984
985
985
986
986
987
987
988
988
989
989
990
990
991
991
992
992
993
993
994
994
995
995
996
996
997
997
998
998
999
999
1000
1000
1001
1001
1002
1002
1003
1003
1004
1004
1005
1005
1006
1006
1007
1007
1008
1008
1009
1009
1010
1010
1011
1011
1012
1012
1013
1013
1014
1014
1015
1015
1016
1016
1017
1017
1018
1018
1019
1019
1020
1020
1021
1021
1022
1022
1023
1023
1024
1024
1025
1025
1026
1026
1027
1027
1028
1028
1029
1029
1030
1030
1031
1031
1032
1032
1033
1033
1034
1034
1035
1035
1036
1036
1037
1037
1038
1038
1039
1039
1040
1040
1041
1041
1042
1042
1043
1043
1044
1044
1045
1045
1046
1046
1047
1047
1048
1048
1049
1049
1050
1050
1051
1051
1052
1052
1053
1053
1054
1054
1055
1055
1056
1056
1057
1057
1058
1058
1059
1059
1060
1060
1061
1061
1062
1062
1063
1063
1064
1064
1065
1065
1066
1066
1067
1067
1068
1068
1069
1069
1070
1070
1071
1071
1072
1072
1073
1073
1074
1074
1075
1075
1076
1076
1077
1077
1078
1078
1079
1079
1080
1080
1081
1081
1082
1082
1083
1083
1084
1084
1085
1085
1086
1086
1087
1087
1088
1088
1089
1089
1090
1090
1091
1091
1092
1092
1093
1093
1094
1094
1095
1095
1096
1096
1097
1097
1098
1098
1099
1099
1100
1100
1101
1101
1102
1102
1103
1103
1104
1104
1105
1105
1106
1106
1107
1107
1108
1108
1109
1109
1110
1110
1111
1111
1112
1112
1113
1113
1114
1114
1115
1115
1116
1116
1117
1117
1118
1118
1119
1119
1120
1120
1121
1121
1122
1122
1123
1123
1124
1124
1125
1125
1126
1126
1127
1127
1128
1128
1129
1129
1130
1130
1131
1131
1132
1132
1133
1133
1134
1134
1135
1135
1136
1136
1137
1137
1138
1138
1139
1139
1140
1140
1141
1141
1142
1142
1143
1143
1144
1144
1145
1145
1146
1146
1147
1147
1148
1148
1149
1149
1150
1150
1151
1151
1152
1152
1153
1153
1154
1154
1155
1155
1156
1156
1157
1157
1158
1158
1159
1159
1160
1160
1161
1161
1162
1162
1163
1163
1164
1164
1165
1165
1166
1166
1167
1167
1168
1168
1169
1169
1170
1170
1171
1171
1172
1172
1173
1173
1174
1174
1175
1175
1176
1176
1177
1177
1178
1178
1179
1179
1180
1180
1181
1181
1182
1182
1183
1183
1184
1184
1185
1185
1186
1186
1187
1187
1188
1188
1189
1189
1190
1190
1191
1191
1192
1192
1193
1193
1194
1194
1195
1195
1196
1196
1197
1197
1198
1198
1199
1199
1200
1200
1201
1201
1202
1202
1203
1203
1204
1204
1205
1205
1206
1206
1207
1207
1208
1208
1209
1209
1210
1210
1211
1211
1212
1212
1213
1213
1214
1214
1215
1215
1216
1216
1217
1217
1218
1218
1219
1219
1220
1220
1221
1221
1222
1222
1223
1223
1224
1224
1225
1225
1226
1226
1227
1227
1228
1228
1229
1229
1230
1230
1231
1231
1232
1232
1233
1233
1234
1234
1235
1235
1236
1236
1237
1237
1238
1238
1239
1239
1240
1240
1241
1241
1242
1242
1243
1243
1244
1244
1245
1245
1246
1246
1247
1247
1248
1248
1249
1249
1250
1250
1251
1251
1252
1252
1253
1253
1254
1254
1255
1255
1256
1256
1257
1257
1258
1258
1259
1259
1260
1260
1261
1261
1262
1262
1263
1263
1264
1264
1265
1265
1266
1266
1267
1267
1268
1268
1269
1269
1270
1270
1271
1271
1272
1272
1273
1273
1274
1274
1275
1275
1276
1276
1277
1277
1278
1278
1279
1279
1280
1280
1281
1281
1282
1282
1283
1283
1284
1284
1285
1285
1286
1286
1287
1287
1288
1288
1289
1289
1290
1290
1291
1291
1292
1292
1293
1293
1294
1294
1295
1295
1296
1296
1297
1297
1298
1298
1299
1299
1300
1300
1301
```

```
26         "storageProfile": {
27
28             "osDiskImage": {
29
30                 "source": {
31
32                     "id": "[parameters('al').input.
33                         source.diskId]"
34
35                 }
36
37             }
38
39             ...
40         }
41
42     }
43
44     ]
45   }
46
47 ],
48 ...
49 }
```

Layered Image Input The `input object` includes the `source` and `diskName` properties. The `source` represents the outputs of the **Machine** deployment, with a property for each output specified. The `diskName` property is the name of the disk specified in the App Layering image template.

Layered Image Output The output of the deployment isn't passed to any other deployments. However, an output named `message` can be specified to provide a message that is appended to the final status of a publish image task in the App Layering UI.

Starter templates

January 5, 2026

This section contains a full set of ARM templates that can be used with the Azure Deployments connector. These templates can be used as-is, or they can be modified to meet specific needs.

Each resource created by these templates is tagged with the same set of tags. These tags include useful information about the context of the deployment, such as the name of the user who started the task and the comment they entered.

The templates make extensive use of custom data in the [connector configuration](#). Custom data allows the user to define common parameters such as `location`, `vmSize`, `generation`, and other parameters without needing to modify the template.

The following table lists all the custom data properties used by these templates. It specifies which templates apply to each property and whether the property is required.

R = Required, O = Optional, - = Not used

Name	Type	Description	Default	Cache Disk	Boot Image	Machine	Layered Image
location	string	The region where resources are created.	Same region as the target resource group.	O	O	O	O
gallery	string	The name of the compute gallery to create images in. The gallery must be in the target resource group.	-	-	R	-	R
storageSku	string	The name of the SKU to use for managed disks.	“Standard-SSD_LRS”	O	-	O	-

App Layering

Name	Type	Description	Default	Cache Disk	Boot Image	Machine	Layered Image
NVMe	boolean	<p>true to use NVMe disk controllers, false otherwise. This must be the same value for all deployment types.</p>	false	0	0	0	0
generation	string	The generation of virtual machine.	“V2”	-	0	0	0
trustedLaunch	boolean	<p>true to enable Trusted Launch, false otherwise. This must be the same value for all deployment types.</p>	false	0	0	0	0

App Layering

Name	Type	Description	Default	Cache Disk	Boot Image	Machine	Layered Image
diskAccessIdstring		The resource ID of the disk access used when up-loading the disk contents. If this is specified, the disk is created with public network access disabled. For more information, see Azure documentation .	null	0	-	-	-
vmSize	string	The size of the VM to create.	“Standard_D2s_v3”	-	-	0	-

App Layering

Name	Type	Description	Default	Cache Disk	Boot Image	Machine	Layered Image
prefix	string	A prefix to prepend to the name of resources created by the machine deployment	“al-“	-	-	O	-
licenseType	string	The on-premises license type to apply to created virtual machines.	null	-	-	O	-
subnetId	string	The resource ID of the subnet to which to attach the VM's network card.	-	-	-	R	-

App Layering

Name	Type	Description	Default	Cache Disk	Boot Image	Machine	Layered Image
acceleratedNetworking	boolean	<p>true to enable accelerated networking, false otherwise. This must be the same value for all deployment types.</p>	false	0	0	0	0
replicaCount	number	The default number of replicas per region of the gallery image version	1	-	-	-	0
imageStorageStringType	string	Specifies the storage account type to be used to store the image.	Specifies the storage account type to be used to store the image.	“Standard_LRS”	-	-	0

Name	Type	Description	Default	Cache Disk	Boot Image	Machine	Layered Image
targetRegionarray	<p>The target regions of the gallery image version. This is specified as an array of region name strings and/or TargetRegion objects.</p> <p>The array must contain the region of the source disk (which is specified by the location custom data).</p>	<p>The region specified by the location custom data</p>	-	-	-	-	0

Name	Type	Description	Default	Cache Disk	Boot Image	Machine	Layered Image
publishAs	string or array	The type of resource(s) to publish images as. Specified as an array or string consisting of ‘gallery-Image’ and/or ‘managedDisk’	[“gallery-Image”]	-	-	-	0
diskNameFormat	string	Controls the name of images published as managed disks. “version” to append the version to the image name, “time-stamp” to append the time of publish.	“version”	-	-	-	0

Example custom data:

```
1  {
2
3      "gallery": "MyGallery",
4      "subnetId": "/subscriptions/ab3d1259-f5a9-407f-bbdd-bfd5701e2e94/
5          resourceGroups/PDGTPB/providers/Microsoft.Network/
6          virtualNetworks/MyVnet/subnets/mysubnet"
7  }
```

Another example of custom data:

```
1  {
2
3      "location": "eastus",
4      "gallery": "MyGallery",
5      "storageSku": "Premium_LRS",
6      "trustedLaunch": true,
7      "diskAccessId": "/subscriptions/ab3d1259-f5a9-407f-bbdd-
8          bfd5701e2e94/resourceGroups/MyResourceGroup/providers/Microsoft.
9              Compute/diskAccesses/MyDiskAccess",
10     "vmSize": "Standard_D4s_v3",
11     "subnetId": "/subscriptions/ab3d1259-f5a9-407f-bbdd-bfd5701e2e94/
12         resourceGroups/MyResourceGroup/providers/Microsoft.Network/
13         virtualNetworks/MyVnet/subnets/mysubnet",
14     "replicaCount": 2,
15     "targetRegions": [
16         "eastus",
17         {
18             "name": "eastus2",
19             "regionalReplicaCount": 5,
20             "storageAccountType": "Premium_LRS"
21         }
22     ]
23 }
```

Cache Disk

Creates a managed disk.

Cache Disk custom data

App Layering

Name	Type	Description	Default	Required
location	string	The region in which resources are created.	Same region as the target resource group	no
storageSku	string	The name of the SKU to use.	“Standard-SSD_LRS”	no
NVMe	boolean	true to enable NVMe disk controllers, false otherwise. This must be the same value for all deployment types.	false	no
acceleratedNetworking	boolean	true to enable accelerated networking, false otherwise. This must be the same value for all deployment types.	false	no
trustedLaunch	boolean	true to enable Trusted Launch, false otherwise. This must be the same value for all deployment types.	false	no

Name	Type	Description	Default	Required
diskAccessId	string	The resource ID of the disk access to use when uploading the disk contents. If specified, the disk is created with public network access disabled.	null	no

Cache Disk template

- 1.2.0.0 - Added support for accelerated networking and NVMe disk controllers
- 1.1.0.0 - Added support for Trusted Launch
- 1.0.0.0 - Initial version

```
1  {
2
3      "$schema": "https://schema.management.azure.com/schemas/2019-04-01/
4          deploymentTemplate.json#",
5      "contentVersion": "1.2.0.0",
6      "parameters": {
7
8          "al": {
9              "type": "object"
10         }
11     }
12   ,
13   ,
14   "variables": {
15
16       "custom": "[parameters('al').context.config.custom]",
17       "location": "[if(contains(variables('custom'), 'location'),
18           variables('custom').location, resourceGroup().location)]",
19       "name": "[concat(parameters('al').context.item.name, '-',
20           parameters('al').context.item.id)]",
21       "tags": {
22
23           "alTaskId": "[parameters('al').context.taskId]",
24           "alUser": "[parameters('al').context.user]",
25           "alComment": "[parameters('al').context.comment]",
26           "alItemType": "[parameters('al').context.item.type]",
27           "alItemId": "[parameters('al').context.item.id]"
28       }
29   }
```

```
26     "alItemName": "[parameters('al').context.item.name]",  
27     "alItemVersion": "[parameters('al').context.item.version.  
28         name]",  
29     "alConfigId": "[parameters('al').context.config.id]",  
30     "alConfigName": "[parameters('al').context.config.name]"  
31   }  
32   ,  
33   "hasDiskAccess": "[contains(variables('custom'), 'diskAccessId  
34     ')]",  
35   "storageSku": "[if(contains(variables('custom'), 'storageSku'),  
36     variables('custom').storageSku, 'StandardSSD_LRS')]",  
37   "trustedLaunch": "[if(contains(variables('custom'), '  
38     trustedLaunch'), variables('custom').trustedLaunch, false())  
39     ]",  
40   "acceleratedNetworking": "[if(contains(variables('custom'), '  
41     acceleratedNetworking'), variables('custom').  
42     acceleratedNetworking, false())]",  
43   "NVMe": "[if(contains(variables('custom'), 'NVMe'), variables('custom').NVMe, false())]"  
44 }  
45 ,  
46   "resources": [  
47     {  
48       "type": "Microsoft.Compute/disks",  
49       "apiVersion": "2024-03-02",  
50       "name": "[variables('name')]",  
51       "location": "[variables('location')]",  
52       "tags": "[variables('tags')]",  
53       "sku": {  
54         "name": "[variables('storageSku')]"  
55       }  
56     },  
57     {  
58       "type": "Microsoft.Compute/virtualMachines",  
59       "apiVersion": "2024-03-02",  
60       "name": "[variables('vmName')]",  
61       "location": "[variables('location')]",  
62       "tags": "[variables('tags')]",  
63       "properties": {  
64         "creationData": {  
65           "createOption": "Upload",  
66           "uploadSizeBytes": "[parameters('al').input.  
67             uploadSize]"  
68         }  
69         ,  
70           "incremental": "false",  
71           "diskAccessId": "[if(variables('hasDiskAccess'),  
72             variables('custom').diskAccessId, null())]",  
73           "networkAccessPolicy": "[if(variables('hasDiskAccess'),  
74             'AllowPrivate', 'AllowAll')]",  
75           "publicNetworkAccess": "[if(variables('hasDiskAccess'),  
76             'Disabled', 'Enabled')]",  
77           "securityProfile": "[if(variables('trustedLaunch'),  
78             createObject('securityType', 'TrustedLaunch'), null  
79             ())]"  
80         }  
81       }  
82     }  
83   ]  
84 }
```

```
65         "supportedCapabilities": {
66
67             "diskControllerTypes": "[if(variables('NVMe'), 'NVMe, SCSI', 'SCSI')]",
68             "acceleratedNetwork": "[variables('acceleratedNetworking')]"
69         }
70
71     }
72
73 }
74
75 ]
76 }
```

Boot Image

The Boot Image deployment creates a gallery image and image version in the gallery specified by the custom data. It outputs the ID of the created compute gallery image version for use by the **Machine** template.

Boot Image custom data

Name	Type	Description	Default	Required
location	string	The region in which resources are created.	Same region as the target resource group.	no
generation	string	The generation of the virtual machine used in the disk.	“V2”	no
NVMe	boolean	true to enable NVMe disk controllers, false otherwise. This must be the same value for all deployment types.	false	no

Name	Type	Description	Default	Required
acceleratedNetworking	boolean	true to enable accelerated networking, false otherwise. This must be the same value for all deployment types.	false	no
trustedLaunch	boolean	true to enable Trusted Launch, false otherwise. This must be the same value for all deployment types.	false	no
gallery	string	The name of the compute gallery to create the image in. The gallery must be in the target resource group.	-	yes

Boot Image template

- 1.2.0.0 - Added support for accelerated networking and NVMe disk controllers
- 1.1.0.0 - Added support for Trusted Launch
- 1.0.0.0 - Initial version

```
1  {
2
3  "$schema": "https://schema.management.azure.com/schemas/2019-04-01/
4    deploymentTemplate.json#",
5  "contentVersion": "1.2.0.0",
6  "parameters": {
7    "al": {
8      "type": "object"
9    }
10 }
```

```
10         }
11     }
12   }
13   ,
14   "variables": {
15     "custom": "[parameters('al').context.config.custom]",
16     "location": "[if(contains(variables('custom'), 'location'),
17                   variables('custom').location, resourceGroup().location)]",
18     "tags": {
19       "alTaskId": "[parameters('al').context.taskId]",
20       "alUser": "[parameters('al').context.user]",
21       "alComment": "[parameters('al').context.comment]",
22       "alItemType": "[parameters('al').context.item.type]",
23       "alItemId": "[parameters('al').context.item.id]",
24       "alItemName": "[parameters('al').context.item.name]",
25       "alItemVersion": "[parameters('al').context.item.version.
26                     name]",
27       "alConfigId": "[parameters('al').context.config.id]",
28       "alConfigName": "[parameters('al').context.config.name]"
29     }
30   ,
31   "name": "[concat(parameters('al').context.item.name, '.', 
32                   replace(parameters('al').context.config.id, '-', ''), '.', 
33                   parameters('al').context.item.id)]",
34   "version": "[parameters('al').context.item.version.name]",
35   "galleryName": "[variables('custom').gallery]",
36   "generation": "[if(contains(variables('custom'), 'generation'),
37                   variables('custom').generation, 'V2')]",
38   "acceleratedNetworking": "[if(contains(variables('custom'), ' 
39                   acceleratedNetworking'), variables('custom').
40                   acceleratedNetworking, false())]",
41   "NVMe": "[if(contains(variables('custom'), 'NVMe'), variables('
42                   custom').NVMe, false())]",
43   "trustedLaunch": "[if(contains(variables('custom'), ' 
44                   trustedLaunch'), variables('custom').trustedLaunch, false())
45                   ]",
46   "possibleFeatures": [
47     {
48       "include": "[variables('trustedLaunch')]",
49       "feature": {
50         "name": "SecurityType",
51         "value": "TrustedLaunch"
52       }
53     }
54   ,
55   {
56     "include": "[variables('acceleratedNetworking')]",
57   }
58 }
```

```
53         "feature": {
54             "name": "IsAcceleratedNetworkSupported",
55             "value": "True"
56         }
57     }
58     ,
59     {
60         "include": "[variables('NVMe')]",
61         "feature": {
62             "name": "DiskControllerTypes",
63             "value": "NVMe, SCSI"
64         }
65     }
66     ]
67 }
68 ,
69 "resources": [
70     {
71         "type": "Microsoft.Compute/galleries/images",
72         "name": "[concat(variables('galleryName'), '/', variables('
73             name'))]",
74         "apiVersion": "2021-07-01",
75         "location": "[variables('location')]",
76         "tags": "[variables('tags')]",
77         "properties": {
78             "description": "[parameters('al').context.item.
79                 description]",
80             "features": "[map(filter(variables('possibleFeatures'),
81                 lambda('x', lambdaVariables('x').include)), lambda
82                 ('x', lambdaVariables('x').feature))]",
83             "hyperVGeneration": "[variables('generation')]",
84             "osType": "Windows",
85             "osState": "Specialized",
86             "endOfLifeDate": "2030-01-01T00:00:00Z",
87             "identifier": {
88                 "publisher": "Citrix",
89                 "offer": "[parameters('al').context.config.id]",
90                 "sku": "[parameters('al').context.item.id]"
91             }
92         }
93     }
94     ,
95     {
96         "resources": [
97             {
98                 "type": "Microsoft.Compute/galleries/images",
99                 "name": "[concat(variables('galleryName'), '/', variables('
100             name'))]",
101                 "apiVersion": "2021-07-01",
102                 "location": "[variables('location')]",
103                 "tags": "[variables('tags')]",
104                 "properties": {
105                     "description": "[parameters('al').context.item.
106                         description]",
107                     "features": "[map(filter(variables('possibleFeatures'),
108                         lambda('x', lambdaVariables('x').include)), lambda
109                         ('x', lambdaVariables('x').feature))]",
110                     "hyperVGeneration": "[variables('generation')]",
111                     "osType": "Windows",
112                     "osState": "Specialized",
113                     "endOfLifeDate": "2030-01-01T00:00:00Z",
114                     "identifier": {
115                         "publisher": "Citrix",
116                         "offer": "[parameters('al').context.config.id]",
117                         "sku": "[parameters('al').context.item.id]"
118                     }
119                 }
120             }
121         ]
122     }
123 }
```

```
102
103     "type": "versions",
104     "apiVersion": "2021-07-01",
105     "name": "[variables('version')]",
106     "location": "[variables('location')]",
107     "dependsOn": [
108         "[resourceId('Microsoft.Compute/galleries/
109             images', variables('galleryName'), variables
110             ('name'))]"
111     ],
112     "tags": "[variables('tags')]",
113     "properties": {
114
115         "publishingProfile": {
116
117             "replicaCount": 1,
118             "targetRegions": [
119                 {
120                     "name": "[variables('location')]"
121                 }
122             ]
123         }
124     },
125     "storageProfile": {
126
127         "osDiskImage": {
128
129             "source": {
130
131                 "id": "[parameters('al').input.
132                     source.diskId]"
133             }
134         }
135     }
136 }
137 }
138 }
139 }
140 }
141 ]
142 }
143 ]
144 ],
145 ],
146 "outputs": {
147
148     "id": {
149
150         "type": "string",
151         "value": "[resourceId('Microsoft.Compute/galleries/images/
```

```
        versions', variables('galleryName'), variables('name'),
        variables('version'))]"
152    }
153
154    }
155
156 }
```

Machine

The Machine deployment creates a virtual machine, NIC, and managed disk. This template works with or without the **Boot Image** deployment being specified.

Machine custom data

Name	Type	Description	Default	Required
location	string	The region in which resources are created.	Same region as the target resource group.	no
storageSku	string	The name of the SKU for the created disk.	“Standard-SSD_LRS”	no
generation	string	The generation for the virtual machine.	“V2”	no
NVMe	boolean	true to use NVMe disk controllers, false otherwise. This must be the same value for all deployment types.	false	no

Name	Type	Description	Default	Required
acceleratedNetworking	boolean	true to enable accelerated networking, false otherwise. This must be the same value for all deployment types.	false	no
trustedLaunch	boolean	true to enable Trusted Launch, false otherwise. This must be the same value for all deployment types.	false	no
secureBoot	boolean	true to enable Secure Boot, false otherwise. This is only applied if trustedLaunch is set to true .	true	no
vTpm	boolean	true to enable the vTPM, false otherwise. This is only applied if trustedLaunch is set to true .	true	no
vmSize	string	The size of VM to create.	“Standard_D2s_v3”	no
licenseType	string	The on-premises license type to apply to the virtual machine.	null	no

Name	Type	Description	Default	Required
subnetId	string	The resource ID of the subnet to which to attach the NIC.	-	yes
prefix	string	A prefix to prepend to the name of created resources	“al-“	no

Machine template

- 1.2.0.0
 - Added support for accelerated networking
 - Added support for NVMe disk controllers
 - Added support for IPS export
 - Added ability to customize the prefix prepended to resource names
 - Added the IP address to the message displayed in the App Layering subtask
 - Disabled network access to the OS disk
- 1.1.0.0 - Added support for Trusted Launch
- 1.0.0.0 - Initial version

```

19      "name": "[concat(variables('prefix'), parameters('al').context.
20          taskId)]",
21      "vmName": "[concat(variables('name'), '-vm')]",
22      "nicName": "[concat(variables('name'), '-nic')]",
23      "diskName": "[concat(variables('name'), '-disk')]",
24      "diskId": "[resourceId('Microsoft.Compute/disks', variables('
25          diskName'))]",
26      "tags": {
27          "alTaskId": "[parameters('al').context.taskId]",
28          "alUser": "[parameters('al').context.user]",
29          "alComment": "[parameters('al').context.comment]",
30          "alItemType": "[parameters('al').context.item.type]",
31          "alItemId": "[parameters('al').context.item.id]",
32          "alItemName": "[parameters('al').context.item.name]",
33          "alItemVersion": "[parameters('al').context.item.version.
34              name]",
35          "alConfigId": "[parameters('al').context.config.id]",
36          "alConfigName": "[parameters('al').context.config.name]"
37      },
38      "source": "[parameters('al').input.disk.image]",
39      "isFromImage": "[not(contains(variables('source'), 'diskId'))]"
40      ,
41      "generation": "[if(contains(variables('custom'), 'generation'),
42          variables('custom').generation, 'V2')]",
43      "vmSize": "[if(contains(variables('custom'), 'vmSize'),
44          variables('custom').vmSize, 'Standard_D2s_v3')]",
45      "storageSku": "[if(contains(variables('custom'), 'storageSku'),
46          variables('custom').storageSku, 'StandardSSD_LRS')]",
47      "trustedLaunch": "[if(contains(variables('custom'),
48          trustedLaunch), variables('custom').trustedLaunch, false())]",
49      "secureBoot": "[if(contains(variables('custom'),
50          secureBoot), variables('custom').secureBoot, variables('trustedLaunch'))]",
51      "vTpm": "[if(contains(variables('custom'), 'vTpm'), variables('
52          custom').vTpm, variables('trustedLaunch'))]",
53      "acceleratedNetworking": "[if(contains(variables('custom'),
54          acceleratedNetworking), variables('custom').
55          acceleratedNetworking, false())]",
56      "NVMe": "[if(contains(variables('custom'), 'NVMe'), variables('
57          custom').NVMe, false())]",
58      "securityProfile": {
59          "securityType": "TrustedLaunch",
60          "uefiSettings": {
61              "secureBootEnabled": "[variables('secureBoot')]",
62              "vTpmEnabled": "[variables('vTpm')]"
63          }
64      }
65  }

```

```
57  ,
58      "dataDiskTemplate": {
59
60          "createOption": "Attach",
61          "deleteOption": "Delete",
62          "managedDisk": {
63
64              "id": "ToBeReplaced"
65          }
66
67      }
68
69      "dataDiskName": "[concat(variables('name'), '-disk{
70  0 }
71  ')]", "additionalDisks": "[if(contains(parameters('al').input, 'additionalDisks'), parameters('al').input.additionalDisks, createArray()))]", "additionalDiskIds": "[map(variables('additionalDisks'), lambda('x', lambdaVariables('x').diskId))]", "dataDiskIds": "[map(variables('additionalDiskIds'), lambda('x', 'i', resourceId('Microsoft.Compute/disks', format(variables('dataDiskName'), lambdaVariables('i')))))]"}
72
73
74
75
76      }
77
78      "resources": [
79          {
80
81              "type": "Microsoft.Network/networkInterfaces",
82              "apiVersion": "2020-11-01",
83              "name": "[variables('nicName')]",
84              "location": "[variables('location')]",
85              "tags": "[variables('tags')]",
86              "properties": {
87
88                  "ipConfigurations": [
89                      {
90
91                          "name": "ipconfig1",
92                          "properties": {
93
94                              "privateIPAllocationMethod": "Dynamic",
95                              "subnet": {
96
97                                  "id": "[variables('custom').subnetId]"
98                              }
99
100
101                          "primary": true,
102                          "privateIPAddressVersion": "IPv4"
103                      }
104                  }
105              }
106          }
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
259
260
261
262
263
264
265
266
267
268
269
269
270
271
272
273
274
275
276
277
278
279
279
280
281
282
283
284
285
286
287
287
288
289
289
290
291
292
293
294
295
296
297
298
299
299
300
301
302
303
304
305
306
307
308
309
309
310
311
312
313
314
315
316
317
318
319
319
320
321
322
323
324
325
326
327
328
329
329
330
331
332
333
334
335
336
337
338
339
339
340
341
342
343
344
345
346
347
348
349
349
350
351
352
353
354
355
356
357
358
359
359
360
361
362
363
364
365
366
367
368
369
369
370
371
372
373
374
375
376
377
378
379
379
380
381
382
383
384
385
386
387
387
388
389
389
390
391
392
393
394
395
396
397
398
399
399
400
401
402
403
404
405
406
407
408
409
409
410
411
412
413
414
415
416
417
418
419
419
420
421
422
423
424
425
426
427
428
429
429
430
431
432
433
434
435
436
437
438
439
439
440
441
442
443
444
445
446
447
448
449
449
450
451
452
453
454
455
456
457
458
459
459
460
461
462
463
464
465
466
467
468
469
469
470
471
472
473
474
475
476
477
478
479
479
480
481
482
483
484
485
486
487
488
489
489
490
491
492
493
494
495
496
497
498
499
499
500
501
502
503
504
505
506
507
508
509
509
510
511
512
513
514
515
516
517
518
519
519
520
521
522
523
524
525
526
527
528
529
529
530
531
532
533
534
535
536
537
538
539
539
540
541
542
543
544
545
546
547
548
549
549
550
551
552
553
554
555
556
557
558
559
559
560
561
562
563
564
565
566
567
568
569
569
570
571
572
573
574
575
576
577
578
579
579
580
581
582
583
584
585
586
587
588
589
589
590
591
592
593
594
595
596
597
598
599
599
600
601
602
603
604
605
606
607
608
609
609
610
611
612
613
614
615
616
617
618
619
619
620
621
622
623
624
625
626
627
628
629
629
630
631
632
633
634
635
636
637
638
639
639
640
641
642
643
644
645
646
647
648
649
649
650
651
652
653
654
655
656
657
658
659
659
660
661
662
663
664
665
666
667
668
669
669
670
671
672
673
674
675
676
677
678
679
679
680
681
682
683
684
685
686
687
687
688
689
689
690
691
692
693
694
695
696
697
697
698
699
699
700
701
702
703
704
705
706
707
708
709
709
710
711
712
713
714
715
716
717
717
718
719
719
720
721
722
723
724
725
726
727
728
729
729
730
731
732
733
734
735
736
737
738
739
739
740
741
742
743
744
745
746
747
748
749
749
750
751
752
753
754
755
756
757
758
759
759
760
761
762
763
764
765
766
767
768
769
769
770
771
772
773
774
775
776
777
778
778
779
779
780
781
782
783
784
785
786
787
787
788
789
789
790
791
792
793
794
795
796
797
797
798
799
799
800
801
802
803
804
805
806
807
808
809
809
810
811
812
813
814
815
816
817
818
818
819
819
820
821
822
823
824
825
826
827
828
829
829
830
831
832
833
834
835
836
837
838
839
839
840
841
842
843
844
845
846
847
848
849
849
850
851
852
853
854
855
856
857
858
859
859
860
861
862
863
864
865
866
867
868
869
869
870
871
872
873
874
875
876
877
878
879
879
880
881
882
883
884
885
886
887
888
889
889
890
891
892
893
894
895
896
897
898
899
899
900
901
902
903
904
905
906
907
908
909
909
910
911
912
913
914
915
916
917
917
918
919
919
920
921
922
923
924
925
926
927
928
929
929
930
931
932
933
934
935
936
937
938
939
939
940
941
942
943
944
945
946
947
947
948
949
949
950
951
952
953
954
955
956
957
958
959
959
960
961
962
963
964
965
966
967
968
969
969
970
971
972
973
974
975
976
977
978
978
979
979
980
981
982
983
984
985
986
987
987
988
989
989
990
991
992
993
994
995
996
997
997
998
999
999
1000
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1018
1019
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1039
1040
1041
1042
1043
1044
1045
1046
1047
1047
1048
1049
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1078
1079
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1088
1089
1089
1090
1091
1092
1093
1094
1095
1096
1096
1097
1097
1098
1099
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1118
1119
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1128
1129
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1138
1139
1139
1140
1141
1142
1143
1144
1145
1146
1147
1147
1148
1148
1149
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1158
1159
1159
1160
1161
1162
1163
1164
1165
1166
1167
1167
1168
1168
1169
1169
1170
1171
1172
1173
1174
1175
1176
1176
1177
1177
1178
1178
1179
1179
1180
1181
1182
1183
1184
1185
1186
1186
1187
1187
1188
1188
1189
1189
1190
1191
1192
1193
1194
1195
1195
1196
1196
1197
1197
1198
1198
1199
1199
1200
1201
1202
1203
1204
1205
1206
1206
1207
1207
1208
1208
1209
1209
1210
1211
1212
1213
1214
1215
1215
1216
1216
1217
1217
1218
1218
1219
1219
1220
1221
1222
1223
1224
1225
1226
1226
1227
1227
1228
1228
1229
1229
1230
1231
1232
1233
1234
1235
1236
1236
1237
1237
1238
1238
1239
1239
1240
1241
1242
1243
1244
1245
1246
1246
1247
1247
1248
1248
1249
1249
1250
1251
1252
1253
1254
1255
1256
1256
1257
1257
1258
1258
1259
1259
1260
1261
1262
1263
1264
1265
1266
1266
1267
1267
1268
1268
1269
1269
1270
1271
1272
1273
1274
1275
1275
1276
1276
1277
1277
1278
1278
1279
1279
1280
1281
1282
1283
1284
1285
1285
1286
1286
1287
1287
1288
1288
1289
1289
1290
1291
1292
1293
1294
1295
1295
1296
1296
1297
1297
1298
1298
1299
1299
1300
1301
1302
1303
1304
1305
1305
1306
1306
1307
1307
1308
1308
1309
1309
1310
1311
1312
1313
1314
1315
1315
1316
1316
1317
1317
1318
1318
1319
1319
1320
1321
1322
1323
1324
1325
1325
1326
1326
1327
1327
1328
1328
1329
1329
1330
1331
1332
1333
1334
1335
1335
1336
1336
1337
1337
1338
1338
1339
1339
1340
1341
1342
1343
1344
1345
1345
1346
1346
1347
1347
1348
1348
1349
1349
1350
1351
1352
1353
1354
1355
1355
1356
1356
1357
1357
1358
1358
1359
1359
1360
1361
1362
1363
1364
1365
1365
1366
1366
1367
1367
1368
1368
1369
1369
1370
1371
1372
1373
1374
1375
1375
1376
1376
1377
1377
1378
1378
1379
1379
1380
1381
1382
1383
1384
1385
1385
1386
1386
1387
1387
1388
1388
1389
1389
1390
1391
1392
1393
1394
1395
1395
1396
1396
1397
1397
1398
1398
1399
1399
1400
1401
1402
1403
1404
1404
1405
1405
1406
1406
1407
1407
1408
1408
1409
1409
1410
1411
1412
1413
1414
1415
1415
1416
1416
1417
1417
1418
1418
1419
1419
1420
1421
1422
1423
1424
1425
1425
1426
1426
1427
1427
1428
1428
1429
1429
1430
1431
1432
1433
1434
1435
1435
1436
1436
1437
1437
1438
1438
1439
1439
1440
1441
1442
1443
1444
1445
1445
1446
1446
1447
1447
1448
1448
1449
1449
1450
1451
1452
1453
1454
1455
1455
1456
1456
1457
1457
1458
1458
1459
1459
1460
1461
1462
1463
1464
1465
1465
1466
1466
1467
1467
1468
1468
1469
1469
1470
1471
1472
1473
1474
1475
1475
1476
1476
1477
1477
1478
1478
1479
1479
1480
1481
1482
1483
1484
1485
1485
1486
1486
1487
1487
1488
1488
1489
1489
1490
1491
1492
1493
1494
1495
1495
1496
1496
1497
1497
1498
1498
1499
1499
1500
1501
1502
1503
1504
1504
1505
1505
1506
1506
1507
1507
1508
1508
1509
1509
1510
1511
1512
1513
1514
1515
1515
1516
1516
1517
1517
1518
1518
1519
1519
1520
1521
1522
1523
1524
1525
1525
1526
1526
1527
1527
1528
1528
1529
1529
1530
1531
1532
1533
1534
1535
1535
1536
1536
1537
1537
1538
1538
1539
1539
1540
1541
1542
1543
1544
1545
1545
1546
1546
1547
1547
1548
1548
1549
1549
1550
1551
1552
1553
1554
1555
1555
1556
1556
1557
1557
1558
1558
1559
1559
1560
1561
1562
1563
1564
1565
1565
1566
1566
1567
1567
1568
1568
1569
1569
1570
1571
1572
1573
1574
1575
1575
1576
1576
1577
1577
1578
1578
1579
1579
1580
1581
1582
1583
1584
1585
1585
1586
1586
1587
1587
1588
1588
1589
1589
1590
1591
1592
1593
1594
1595
1595
1596
1596
1597
1597
1598
1598
1599
1599
1600
1601
1602
1603
1604
1604
1605
1605
1606
1606
1607
1607
1608
1608
1609
1609
1610
1611
1612
1613
1614
1615
1615
1616
1616
1617
1617
1618
1618
1619
1619
1620
1621
1622
1623
1624
1625
1625
1626
1626
1627
1627
1628
1628
1629
1629
1630
1631
1632
1633
1634
1635
1635
1636
1636
1637
1637
1638
1638
1639
1639
1640
1641
1642
1643
1644
1645
1645
1646
1646
1647
1647
1648
1648
1649
1649
1650
1651
1652
1653
1654
1655
1655
1656
1656
1657
1657
1658
1658
1659
1659
1660
1661
1662
1663
1664
1665
1665
1666
1666
1667
1667
1668
1668
1669
1669
1670
1671
1672
1673
1674
1675
1675
1676
1676
1677
1677
1678
1678
1679
1679
1680
1681
1682
1683
1684
1685
1685
1686
1686
1687
1687
1688
1688
1689
1689
1690
1691
1692
1693
1694
1695
1695
1696
1696
1697
1697
1698
1698
1699
1699
1700
1701
1702
1703
1704
1704
1705
1705
1706
1706
1707
1707
1708
1708
1709
1709
1710
1711
1712
1713
1714
1715
1715
1716
1716
1717
1717
1718
1718
1719
1719
1720
1721
1722
1723
1724
1725
1725
1726
1726
1727
1727
1728
1728
1729
1729
1730
1731
1732
1733
1734
1735
1735
1736
1736
1737
1737
1738
1738
1739
1739
1740
1741
1742
1743
1744
1745
1745
1746
1746
1747
1747
1748
1748
1749
1749
1750
1751
1752
1753
1754
1755
1755
1756
1756
1757
1757
1758
1758
1759
1759
1760
1761
1762
1763
1764
1765
1765
1766
1766
1767
1767
1768
1768
1769
1769
1770
1771
1772
1773
1774
1775
1775
1776
1776
1777
1777
1778
1778
1779
1779
1780
1781
1782
1783
1784
1785
1785
1786
1786
1787
1787
1788
1788
1789
1789
1790
1791
1792
1793
1794
1795
1795
1796
1796
1797
1797
1798
1798
1799
1799
1800
1801
1802
1803
1804
1804
1805
1805
1806
1806
1807
1807
1808
1808
1809
1809
1810
1811
1812
1813
1814
1815
1815
1816
1816
1817
1817
1818
1818
1819
1819
1820
1821
1822
1823
1824
1825
1825
1826
1826
1827
1827
1828
1828
1829
1829
1830
1831
1832
1833
1834
1835
1835
1836
1836
1837
1837
1838
1838
1839
1839
1840
1841
1842
1843
1844
1845
1845
1846
1846
1847
1847
1848
1848
1849
1849
1850
1851
1852
1853
1854
1855
1855
1856
1856
1857
1857
1858
1858
1859
1859
1860
1861
1862
1863
1864
1865
1865
1866
1866
1867
1867
1868
1868
1869
1869
1870
1871
1872
1873
1874
1875
1875
1876
1876
1877
1877
1878
1878
1879
1879
1880
1881
1882
1883
1884
1885
1885
1886
1886
1887
1887
1888
1888
1889
1889
1890
1891
1892
1893
1894
1895
1895
1896
1896
1897
1897
1898
1898
1899
1899
1900
1901
1902
1903
1904
1904
1905
1905
1906
1906
1907
1907
1908
1908
1909
1909
1910
1911
1912
1913
1914
1915
1915
1916
1916
1917
1917
1918
1918
1919
1919
1920
1921
1922
1923
1924
1925
1925
1926
1926
1927
1927
1928
1928
1929
1929
1930
1931
1932
1933
1934
1935
1935
1936
1936
1937

```

```
105     ],
106     "dnsSettings": {
107       "dnsServers": []
108     }
109   },
110   "enableAcceleratedNetworking": "[variables('
111     acceleratedNetworking')]",
112   "enableIPForwarding": false
113 }
114
115
116 }
117
118 {
119   "condition": "[not(variables('isFromImage'))]",
120   "type": "Microsoft.Compute/disks",
121   "apiVersion": "2021-08-01",
122   "name": "[variables('diskName')]",
123   "location": "[variables('location')]",
124   "tags": "[variables('tags')]",
125   "sku": {
126     "name": "[variables('storageSku')]"
127   }
128
129   "properties": {
130     "osType": "Windows",
131     "hyperVGeneration": "[variables('generation')]",
132     "creationData": {
133       "createOption": "Copy",
134       "sourceResourceId": "[variables('source').diskId]"
135     }
136   }
137   "diskSizeGB": "[parameters('al').input.disk.size]",
138   "networkAccessPolicy": "DenyAll",
139   "publicNetworkAccess": "Disabled",
140   "supportedCapabilities": {
141     "acceleratedNetwork": "[variables('
142       acceleratedNetworking')]"
143   }
144
145   "copy": {
```

```
156         "name": "data-disk-loop",
157         "count": "[length(variables('additionalDiskIds'))]",
158         "mode": "Parallel"
159     }
160 ,
161 ,
162     "type": "Microsoft.Compute/disks",
163     "apiVersion": "2021-08-01",
164     "name": "[format(variables('dataDiskName'), copyIndex())]",
165     "location": "[variables('location')]",
166     "tags": "[variables('tags')]",
167     "sku": {
168
169         "name": "[variables('storageSku')]"
170     }
171 ,
172 ,
173     "properties": {
174
175         "creationData": {
176
177             "createOption": "Copy",
178             "sourceResourceId": "[variables('additionalDiskIds'
179             ')[copyIndex()]]"
180         }
181 ,
182         "networkAccessPolicy": "DenyAll",
183         "publicNetworkAccess": "Disabled"
184     }
185 ,
186     {
187
188         "type": "Microsoft.Compute/virtualMachines",
189         "apiVersion": "2022-08-01",
190         "name": "[variables('vmName')]",
191         "location": "[variables('location')]",
192         "dependsOn": [
193             "[resourceId('Microsoft.Network/networkInterfaces',
194                 variables('nicName'))]",
195             "[variables('diskId')]",
196             "data-disk-loop"
197         ],
198         "tags": "[variables('tags')]",
199         "properties": {
200
201             "hardwareProfile": {
202
203                 "vmSize": "[variables('vmSize')]"
204             }
205 ,
206             "securityProfile": "[if(variables('trustedLaunch'),
207                 variables('securityProfile'), null())]"
208         }
209     }
210 }
```

```
206     "storageProfile": {
207
208         "imageReference": "[if(variables('isFromImage'),
209             createObject('id', variables('source').id), null
210             ())]",
211         "diskControllerType": "[if(variables('NVMe'), 'NVMe',
212             'null())]",
213         "osDisk": {
214
215             "osType": "Windows",
216             "createOption": "[if(variables('isFromImage'),
217                 'FromImage', 'Attach')]",
218             "caching": "ReadWrite",
219             "deleteOption": "Delete",
220             "diskSizeGB": "[parameters('a1').input.disk.
221                 size]",
222             "managedDisk": "[if(variables('isFromImage'),
223                 createObject('storageAccountType', variables
224                     ('storageSku')), createObject('id',
225                     variables('diskId')))]"
226         }
227     ,
228
229     ,
230         "dataDisks": "[map(variables('dataDiskIds'), lambda
231             ('x', 'i', union(variables('dataDiskTemplate'),
232                 createObject('lun', lambdaVariables('i'),
233                     'managedDisk', createObject('id',
234                     lambdaVariables
235                         ('x'))))))]"
236     }
237
238     ,
239         "networkProfile": {
240
241             "networkInterfaces": [
242
243                 {
244
245                     "id": "[resourceId('Microsoft.Network/
246                         networkInterfaces', variables('nicName'))
247                         )]"
248                 }
249             ]
250         }
251
252     ,
253         "diagnosticsProfile": {
254
255             "bootDiagnostics": {
256
257                 "enabled": true
258             }
259         }
260
261     ,
262         "licenseType": "[if(contains(variables('custom'),
263             'licenseType'), variables('custom').licenseType, null
```

```
244         ()])",
245         "userData": "[parameters('al').input.vm.userData]"
246     }
247 }
248 ]
249 ,
250 "outputs": {
251     "diskId": {
252         "type": "string",
253         "value": "[reference(variables('vmName')).storageProfile.
254             osDisk.managedDisk.id]"
255     }
256 ,
257     "message": {
258         "type": "string",
259         "value": "[format('See [link=\\\"{"
260             0 }
261             /#{@{
262             1 }
263             /resource/{"
264             2 }
265             \\\"]{"
266             2 }
267             [/link]. IP Address: {
268             3 }
269             ', environment().portal, tenant().tenantId, resourceId('Microsoft.
270                 Compute/virtualMachines', variables('vmName')), reference(variables
271                 ('nicName')).ipConfigurations[0].properties.privateIPAddress)]"
272             }
273     }
274 }
275
276 }
```

Layered Image

The Layered Image deployment can create two types of resources, a gallery image version and/or a managed disk. Both resource types are named using the name of the App Layering image template being published and its constructed version number. The version number is constructed using major and minor numbers from App Layering image template disk name. If the disk name is not formatted as a numeric version ([number.number](#)), then 1.0 is applied by default. The patch number is the version number of the App Layering image template (the number of times that it's been published). The gallery image is assigned the name and the gallery image version is assigned the version number. The managed disk is assigned the name appended with the version number or the current timestamp depending on [diskNameFormat](#).

The gallery image and version are created in the gallery specified by the custom data. When an image is published multiple times, a new version is added to the compute gallery image, and the old version stays.

Layered Image custom data

R = Required, O = Optional, - = Not used

Name	Type	Description	Default	Gallery Image	Managed Disk
publishAs	string or array	The type of resource(s) to create. Specified as an array or string consisting of 'galleryImage' and/or 'managedDisk'	[“galleryImage”]	O	R
location	string	The region where resources are created.	Same region as the target resource group	O	O
generation	string	The generation of the virtual machine that the image will support.	“V2”	O	O

App Layering

Name	Type	Description	Default	Gallery Image	Managed Disk
NVMe	boolean	true to enable NVMe disk controllers, false otherwise. This must be the same value for all deployment types.	false	0	0
acceleratedNetworking	boolean	true to enable accelerated networking, false otherwise. This must be the same value for all deployment types.	false	0	0
trustedLaunch	boolean	true to enable Trusted Launch, false otherwise. This must be the same value for all deployment types.	false	0	0

App Layering

Name	Type	Description	Default	Gallery Image	Managed Disk
gallery	string	The name of the compute gallery to create the image in. The gallery must be in the target resource group.	-	R	-
replicaCount	number	The default number of replicas per region of the gallery image version	1	O	-
imageStorageAccountType	string	Specifies the storage account type to be used to store the image.	“Standard_LRS”	O	-

Name	Type	Description	Default	Gallery Image	Managed Disk
targetRegions	array	The target regions of the gallery image version. This is specified as an array of region name strings and/or TargetRegion objects. The array must contain the region of the source disk (as specified by the location custom data).	The region specified by the location custom data	0	-
storageSku	string	The name of the SKU to use.	“Standard-SSD_LRS”	-	0
diskAccessId	string	The resource ID of the disk access to use when uploading the disk contents. If specified, then the disk is created with public network access disabled.	null	-	0

Name	Type	Description	Default	Gallery Image	Managed Disk
diskNameFormatstring		Controls the name of images published as managed disks. “version” to append the version to the image name, “timestamp” to append the time of publish.	“version”	-	0

Layered Image template

- 1.2.0.0
 - Added support for accelerated networking
 - Added support for NVMe disk controllers
 - Added ability to specify the storage account type of gallery image versions
 - Added choice to name managed disks using the publish time rather than image version (similar to legacy naming)
 - Fixed a bug that made `galleryName` required even when not publishing as a gallery image version
- 1.1.0.0
 - Added support for Trusted Launch
 - Added replica count and target region support for gallery image versions
 - Added support for publishing as a managed disk
- 1.0.0.0 - Initial version

```
1  {
2
3  "$schema": "https://schema.management.azure.com/schemas/2019-04-01/
   deploymentTemplate.json#",
4  "contentVersion": "1.2.0.0",
5  "parameters": {
```

```
6         "al": {
7             "type": "object"
8         }
9     ,
10    "timestamp": {
11        "type": "string",
12        "defaultValue": "[utcNow('yyyy-MM-dd_HH-mm-ss')]"
13    }
14 }
15 ,
16 "variables": {
17     "invalidChars": [ " ", "(", ")", "[[", "]", "{",
18     "!", "@", "#", "$", "%", "^", "&", "*",
19     "+", "/", "\\", "", "\\", "\\", "\\", "\",
20     "|", "`", "~", "<", ">", ",",
21     "?", "*"], [
22     "numericChars": [ "0", "1", "2", "3", "4", "5", "6", "7", "8",
23     "9" ],
24     "custom": "[parameters('al').context.config.custom]",
25     "location": "[if(contains(variables('custom'), 'location'),
26         variables('custom').location, resourceGroup().location)]",
27     "name": "[join(split(parameters('al').context.item.name,
28         variables('invalidChars')), '_')]",
29     "tags": {
30         "alTaskId": "[parameters('al').context.taskId]",
31         "alUser": "[parameters('al').context.user]",
32         "alComment": "[parameters('al').context.comment]",
33         "alItemType": "[parameters('al').context.item.type]",
34         "alItemId": "[parameters('al').context.item.id]",
35         "alItemName": "[parameters('al').context.item.name]",
36         "alItemVersion": "[parameters('al').context.item.version.
37             name]",
38         "alConfigId": "[parameters('al').context.config.id]",
39         "alConfigName": "[parameters('al').context.config.name]"
40     }
41 },
42     "splitVer": "[split(parameters('al').input.diskName, '.')]",
43     "major": "[if(equals(0, length(join(split(variables('splitVer')
44         [0], variables('numericChars')), ''))), variables('splitVer
45         ')[0], '1'))",
46     "minor": "[if(greater(length(variables('splitVer')), 1), if(
47         equals(0, length(join(split(variables('splitVer')[1],
48         variables('numericChars')), ''))), variables('splitVer')[1],
49         '0'), '0')]",
50     "version": "[format('{"
51     0 }
52     .{
53     1 }
```

```
49  .{
50  2 }
51  ', variables('major'), variables('minor'), parameters('al').context.
52    item.version.number)]",
53    "galleryName": "[if(contains(variables('custom'), 'gallery'),
54      variables('custom').gallery, 'galleryName.required.in.custom
55      .data')]",
56    "generation": "[if(contains(variables('custom'), 'generation'),
57      variables('custom').generation, 'V2')]",
58    "trustedLaunch": "[if(contains(variables('custom'),
59      trustedLaunch'), variables('custom').trustedLaunch, false())
60    ]",
61    "acceleratedNetworking": "[if(contains(variables('custom'),
62      acceleratedNetworking), variables('custom').
63      acceleratedNetworking, false())]",
64    "NVMe": "[if(contains(variables('custom'), 'NVMe'), variables('
65      custom').NVMe, false())]",
66    "replicaCount": "[if(contains(variables('custom'),
67      replicaCount'), variables('custom').replicaCount, 1)]",
68    "targetRegions": "[if(contains(variables('custom'),
69      targetRegions'), variables('custom').targetRegions,
70      createArray(variables('location')))]",
71  "diskNames": {
72    "version": "[format('{"
73      0 }
74      _{
75      1 }
76      -{
77      2 }
78      -{
79      3 }
80    ', variables('name'), variables('major'), variables('minor'),
81      parameters('al').context.item.version.number)]",
82    "timestamp": "[format('{"
83      0 }
84      _{
85      1 }
86    ', variables('name'), parameters('timestamp'))]"
87  ,
88    "diskName": "[variables('diskNames')[coalesce(tryGet(variables
89      ('custom'), 'diskNameFormat'), 'version')]]",
90    "hasDiskAccess": "[contains(variables('custom'), 'diskAccessId
91      ')]",
92    "storageSku": "[if(contains(variables('custom'), 'storageSku'),
93      variables('custom').storageSku, 'StandardSSD_LRS')]",
94    "imageStorageAccountType": "[coalesce(tryGet(variables('custom
95      '), 'imageStorageAccountType'), 'Standard_LRS')]",
96    "publishAs": "[if(contains(variables('custom'), 'publishAs'),
97      variables('custom').publishAs, createArray('galleryImage'))]
98    ",
99    "galleryLink": "[format('[link=\\"{"
100      0 }
101      _{
102      1 }
103      -{
104      2 }
105      -{
106      3 }
107    ', variables('name'))]"
108  ]
```

```
83  0 }
84  #@{
85  1 }
86  /resource|{
87  2 }
88  \"]{
89  2 }
90  [/link].', environment().portal, tenant().tenantId, resourceId('
  Microsoft.Compute/galleries/images/versions', variables('
    galleryName'), variables('name'), variables('version')))]",
91  "diskLink": "[format('[link=\\"{
92  0 }]
93  #@{
94  1 }
95  /resource|{
96  2 }
97  \"]{
98  2 }
99  [/link].', environment().portal, tenant().tenantId, resourceId('
  Microsoft.Compute/disks', variables('diskName')))]",
100  "outputLinks": "[filter(createArray(if(contains(variables('
    publishAs'), 'galleryImage'), variables('galleryLink'), null
    ()), if(contains(variables('publishAs'), 'managedDisk'),
    variables('diskLink'), null()), lambda('link', not>equals(
      lambdaVariables('link'), null())))))]",
101  "possibleFeatures": [
102  {
103
104    "include": "[variables('trustedLaunch')]",
105    "feature": {
106
107      "name": "SecurityType",
108      "value": "TrustedLaunch"
109    }
110  }
111  ,
112  ,
113  {
114
115    "include": "[variables('acceleratedNetworking')]",
116    "feature": {
117
118      "name": "IsAcceleratedNetworkSupported",
119      "value": "True"
120    }
121  }
122  ,
123  ,
124  {
125
126    "include": "[variables('NVMe')]",
127    "feature": {
```

```
129             "name": "DiskControllerTypes",
130             "value": "NVMe, SCSI"
131         }
132
133     ]
134
135   }
136
137   ,
138   "resources": [
139     {
140
141       "condition": "[contains(variables('publishAs'), 'galleryImage')]",
142       "type": "Microsoft.Compute/galleries/images",
143       "name": "[format('{",
144       0 }]
145     /{
146     1 }
147     ', variables('galleryName'), variables('name'))]",
148     "apiVersion": "2021-07-01",
149     "location": "[variables('location')]",
150     "properties": {
151
152       "description": "[parameters('al').context.item.description]",
153       "features": "[map(filter(variables('possibleFeatures'), lambda('x', lambdaVariables('x').include)), lambda('x', lambdaVariables('x').feature))]",
154       "hyperVGeneration": "[variables('generation')]",
155       "osType": "Windows",
156       "osState": "Specialized",
157       "endOfLifeDate": "2030-01-01T00:00:00Z",
158       "identifier": {
159
160         "publisher": "AppLayering",
161         "offer": "[variables('name')]",
162         "sku": "[variables('generation')]"
163       }
164
165     }
166
167   ,
168   "tags": "[variables('tags')]",
169   "resources": [
170     {
171
172       "condition": "[contains(variables('publishAs'), 'galleryImage')]",
173       "type": "versions",
174       "apiVersion": "2022-03-03",
175       "name": "[variables('version')]",
176       "location": "[variables('location')]",
177       "dependsOn": [
```

```
177         "[resourceId('Microsoft.Compute/galleries/  
178             images', variables('galleryName'), variables  
179                 ('name'))]"  
180     ],  
181     "tags": "[variables('tags')]",  
182     "properties": {  
183         "publishingProfile": {  
184             "replicaCount": "[variables('replicaCount')  
185                 ]",  
186             "storageAccountType": "[variables('  
187                 imageStorageAccountType')]",  
188             "targetRegions": "[map(variables('  
189                 targetRegions'), lambda('item', if(  
190                 contains(lambdaVariables('item'), 'name  
191                 '), lambdaVariables('item'),  
192                 createObject('name', lambdaVariables('  
193                     item')))))]"  
194         }  
195     },  
196     "storageProfile": {  
197         "osDiskImage": {  
198             "source": {  
199                 "id": "[parameters('al').input.  
200                     source.diskId]"  
201             }  
202         }  
203     }  
204 }  
205 ]  
206 }  
207 },  
208 {  
209     "condition": "[contains(variables('publishAs'), '  
210         managedDisk')]",  
211     "type": "Microsoft.Compute/disks",  
212     "apiVersion": "2021-08-01",  
213     "name": "[variables('diskName')]",  
214     "location": "[variables('location')]",  
215     "tags": "[variables('tags')]",  
216     "sku": {  
217 }
```

```
219         "name": "[variables('storageSku')]"
220     }
221     ,
222     "properties": {
223
224         "osType": "Windows",
225         "hyperVGeneration": "[variables('generation')]",
226         "creationData": {
227
228             "createOption": "Copy",
229             "sourceResourceId": "[parameters('al').input.source
230                 .diskId]"
231         }
232         ,
233         "diskAccessId": "[if(variables('hasDiskAccess'),
234             variables('custom').diskAccessId, null())]",
235         "networkAccessPolicy": "[if(variables('hasDiskAccess'),
236             'AllowPrivate', 'AllowAll')]",
237         "publicNetworkAccess": "[if(variables('hasDiskAccess'),
238             'Disabled', 'Enabled')]"
239     }
240     ],
241     "outputs": {
242
243         "message": {
244
245             "type": "string",
246             "value": "[if(empty(variables('outputLinks')), null(),
247                 format('See {
248                     0
249                 .', join(variables('outputLinks'), ' and ')))]"
250         }
251     }
252 }
```

Template parameters

January 5, 2026

This article describes objects passed to all templates associated with an Azure Deployments connector. The object is passed as the `al` parameter and can be accessed via the ARM template function `parameters`, such as `[parameters('al').context.user]`.

```
1 {
```

```
2
3     "input": [
4         {
5             "type": "BootImage",
6             "source": "any"
7         }
8     ],
9     // OR
10    [
11        {
12            "type": "CacheDisk",
13            "size": "number",
14            "uploadSize": "number"
15        }
16    ],
17    // OR
18    [
19        {
20            "type": "LayeredImage",
21            "diskName": "string",
22            "source": "any"
23        }
24    ],
25    // OR
26    [
27        {
28            "type": "Machine",
29            "disk": [
30                {
31                    "image": "any",
32                    "size": "number",
33                    "name": "string"
34                }
35            ],
36            "vm": [
37                {
38                    "userData": "string"
39                }
40            ]
41        }
42    ],
43    "context": [
44        {
45            "taskId": "number",
46            "type": "string",
47            "user": "string",
48            "comment": "string",
49            "config": [
50                {
51                    "id": "string",
52                    "name": "string",
53                    "custom": "any"
54                }
55            ]
56        }
57    ]
58 ]
```

```
55     }
56     ,
57     "item": {
58
59         "type": "string",
60         "id": "number",
61         "name": "string",
62         "description": "string",
63         "created": "string",
64         "modified": "string",
65         "version": {
66
67             "number": "number",
68             "name": "string",
69             "description": "string",
70             "created": "string"
71         }
72
73     }
74
75 }
76
77 }
```

AlParam

The set of data passed to every deployment as the *al* parameter. This can be accessed from within the ARM template by passing *al* to the

[parameters](#)

ARM template function, such as `parameters('al')`.

Name	Description	Value
input	A set of data that applies only to a specific deployment type. The type of this property depends on the deployment type specified by <code>context.type</code>	BootImageInput, CacheDiskInput, LayeredImageInput, MachineInput
context	A set of data that applies to every deployment type.	DeploymentContext

BootImageInput

Input data specific to the BootImage deployment type.

Name	Description	Value
type	The type of deployment to which this input applies.	‘BootImage’
source	The source for the image. This is an object with a property for each output specified in the ARM template associated with the Cache Disk deployment.	any

CacheDiskInput

Input data specific to the CacheDisk deployment type.

Name	Description	Value
type	The type of deployment to which this input applies.	‘CacheDisk’
size	The size of the managed disk to be created, in GiB.	number (32-bit integer)
uploadSize	The size of source disk that will be uploaded to the created disk, in bytes.	number (64-bit integer)

LayeredImageInput

Input data specific to the LayeredImage deployment type.

Name	Description	Value
type	The type of deployment to which this input applies.	‘LayeredImage’
diskName	The name of the disk as specified in the App Layering image template of the image currently being published.	string

Name	Description	Value
source	The source for the image. This is an object with a property for each output specified in the ARM template associated with the Machine deployment.	any

MachineInput

Input data specific to the Machine deployment type.

Name	Description	Value
type	The type of deployment to which this input applies.	‘Machine’
disk	Input data that applies to the OS disk of the machine.	DiskInput
vm	Input data that applies to the VM itself.	VmInput

DiskInput

Input data that applies to the OS disk of a machine.

Name	Description	Value
image	The source for the disk. This is an object with a property for each output specified in the ARM template associated with the Boot Image deployment, if the Boot Image deployment is specified in the connector configuration. If the Boot Image deployment is not specified, then output of the ARM template associated with the Cache Disk deployment is used.	any
size	The size of the disk to be created, in GiB.	number (64-bit integer)
name	The name of the disk as specified in the App Layering image template of the image currently being published. If this is not an image publish operation, then this is undefined.	string

VmInput

Input data that applies to a VM.

Name	Description	Value
userData	The value that must be applied to the <code>userData</code> property of the <code>VirtualMachineProperties</code> of the <code>virtualMachine</code> specified in the ARM template associated with the deployment.	string

DeploymentContext

Provides the context of the current deployment operation.

Name	Description	Value
taskId	The ID of the ELM task that caused the deployment.	number (64-bit integer)
type	The type of the current deployment.	‘CacheDisk’, ‘BootImage’, ‘Machine’, ‘LayeredImage’
user	The username of the user who started the ELM task that caused this deployment.	string
comment	The comment the user entered when starting the ELM task that caused this deployment.	string
config	The platform connector configuration that is associated with this deployment.	Config
item	The App Layering entity associated with the ELM task that caused this deployment. This will refer to a layer, a layered image template, or a boot image.	Item

Config

Contains the properties of the platform connector configuration associated with the current operation.

Name	Description	Value
id	The ID of the platform connector configuration. This is a guid in the form of xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx.	string
name	The name of platform connector configuration.	string

App Layering

Name	Description	Value
custom	The custom data specified in the platform connector configuration. The type and properties of this object depend on the JSON specified for the current deployment type by the user.	any

Item

The App Layering entity associated with an ELM task that caused a deployment.

This can be a layer, a layered image template, or a boot image.

Name	Description	Value
type	The item type of App Layering entity.	‘Image’, ‘AppLayer’, ‘OsLayer’, ‘PlatformLayer’, ‘Connector-CachedCeBootImage’
id	The ID of the item.	number (64-bit integer)
name	The name of the item.	string
description	The description of the item.	string
created	The date and time the item was originally created.	string (ISO 8601 datetime)
modified	The date and time the item was last changed.	string (ISO 8601 datetime)
version	Information about this particular version of the item.	Version

Version

The version of an App Layering entity associated with an ELM task that caused a deployment. This can be a layer version, a layered image template version, or a boot image version.

Name	Description	Value
number	An integer that represents the version. For layers, this is the number of attempted packagings. For layered image templates, this is the number of successful publishes of the template. For boot images, this is the timestamp of the image file.	number (32-bit integer)
name	The name of the version. For layers, this is the version name specified by the user. For layered image templates, this is <i>Publish{Number}_. For boot images, this is the version (x.y.z) of App Layering creating the deployment.</i>	string
description	The description of the version. For layers, this is the version description specified by the user. For layered image templates, this is an empty string. For boot images, this is a generic description that includes the appliance version number.	string
created	The date and time the version was created. For layers, this is when the version was created. For layered image templates, this is when the template was last edited. For boot images, this is the timestamp of the image file.	string (ISO 8601 datetime)

Image Portability for Azure

January 5, 2026

The Image Portability for Azure connector is a specialized version of the [Azure Deployments connector](#) optimized for IPS tasks in Azure. This article describes the connector configuration settings. For more information about connector configurations and how to add new ones, see [Connector configurations](#).

Additional information about Image Portability is available on our [Image Portability](#) page.

Identity

Information on the **Identity** field can be found on our [Azure](#) page.

Location

Location is the Azure region where Azure resources created by the connector are hosted.

Virtual Machine Settings

Virtual Machine Settings are the fields used to specify where and how virtual machines are created within Azure by the connector. Refer to the following descriptions for information on each required and optional field on the connector configuration screen.

Resource Group

The **Azure Resource Group** in which to create the temporary virtual machines that will be used to process images. All resources related to the virtual machines will be created in this resource group.

The managed identity or registration credentials specified in the Identity section must have permission to manage the following types of resources within this resource group:

- Managed Disks
- Network Interface Cards
- Virtual Machines
- Network Security Groups (for the optional Network Isolation field)

Assigning the general [Contributor role](#) to the managed identity or registration credentials on the resource group grants the required permissions.

Virtual Network/Subnet

The **Azure Virtual Network/Subnet** to which to connect the virtual machine.

Note:

The identity of the connector is configured to use must have permission to connect a device to this virtual network or subnet. If the virtual network isn't located in any of the resource groups specified in the connector configuration, then roles assigned to the managed identity or registration credentials for those resource groups won't apply to the virtual network and a role must be assigned directly to the virtual network.

Machine Size

The **Machine Size** determines the compute resources allocated to the virtual machine. This field defaults to a value of Standard_D2s_v3 if a value is not provided. For more information on machine sizes, refer to [Azure's documentation](#).

Network Isolation (Optional)

The **Network Isolation** feature prevents created VMs from communicating with anything outside of the specified virtual network. This helps prevent updates from occurring for the OS and Applications within an image while the image is being prepared.

Encryption at Host (Optional)

The **Encryption at Host** feature ensures that all data stored on the physical host machine is encrypted using platform-managed keys before being written to disk. For more information on disk encryption, refer to [Azure's documentation](#).

Output Settings

Output Settings are fields to identify where and how the prepared image will be created within Azure. Refer to the following descriptions for information on each required field on the connector configuration screen.

Resource Group

This is the **Resource Group** that will contain the output image.

The managed identity or registration credentials specified in the Identity section must have permission to manage the following types of resources within this resource group:

- Managed Disks
- Snapshots

Assigning the general [Contributor role](#) to the managed identity or registration credentials on the resource group grants the required permissions.

Type

The **Type** of image to be created, either a Managed Disk or Snapshot.

Optional Settings

Optional settings that apply to all resources created with this connector configuration. Refer to the following descriptions for information on each optional field on the connector configuration screen.

Resource prefix

A prefix to prepend to the name of all resources created with this connector configuration.

Tags

The tags to apply to all resources created with this connector configuration.

XenServer®

January 5, 2026

XenServer is a highly optimized hypervisor platform for Citrix Virtual Apps and Desktops™, enabling Windows and Linux Virtual Apps and Desktops delivery to any device and hundreds of employees in just minutes.

With the XenServer architecture, you can interact with individual servers or a cluster of servers instead of a central management server. You can manage XenServer by using command-line access or management software, such as XenCenter. You can install XenCenter on your desktop and connect individually to each host or a cluster of hosts.

Before you start

You can use your XenServer environment to create layers and publish layered images. Each Connector configuration accesses a specific storage location in your XenServer environment. You might need more than one XenServer Connector configuration to access the correct location for each purpose. Further, you can publish each layered image to a location convenient to the system you provision with the published image. You can also export images from your XenServer environment via the [Image Portability](#) feature. For more information about Connectors and Connector configurations, see [Connector configurations](#).

Layering

To package layers using a XenServer virtual machine, specify a XenServer Connector configuration when [creating](#) or [versioning](#) a layer. Specify a XenServer or Machine Creation for XenServer Connector configuration in an [image template](#) to have that image created in XenServer when [published](#).

Image Portability

You can export existing images from your XenServer environment by using the XenServer Connector in conjunction with the [Windows File Share Connector](#) via the [Image Portability](#) export feature. Specify a XenServer Connector with Offload Compositing enabled in the **Connector Configuration** dropdown of a Windows File Share Connector and then specify that Windows File Share Connector to the IPS export command. This will allow Image Portability to access your XenServer environment, and export images to the configured file share location. Use the uuid of a VDI to specify to the export command which image to export.

Machine Creation for XenServer® Connector

The Machine Creation version of the XenServer Connector publishes layered images in a state that is ready for use by Citrix Machine Creation Service in your XenServer environment. Published layered images are booted for final configuration, shut down, and then a snapshot is taken.

Notes:

The Machine Creation for XenServer Connector configuration is for publishing layered images. You cannot package layers with it, nor use it with the Image Portability feature. Instead, use a XenServer Connector configuration.

The Machine Creation for XenServer Connector does not support caching

Required information for XenServer Connector configuration settings

Refer to the following required, recommended, and optional fields for configuration of your XenServer Connector.

- **XenServer Configuration Name:** A useful name to help identify and keep track of this connector configuration.
- **XenServer Address:** The name of the XenServer host with which the appliance integrates.
- **User Name and Password:** The credentials for the account that the appliance uses to connect to the XenServer.
- **Use Secured Communications:** SSL encryption for the API connection traffic between the App Layering Connector and XenServer. This field is selected by default.
- **Virtual Machine Template:** The virtual machine template for cloning. The list of choices contains custom virtual machine templates only, rather than actual virtual machines or any of the built-in templates. The selected template must not have any disks attached and must have at least one network card attached. If it does not, you see an error when trying to validate or save the configuration.
- **Storage Repository:** The storage repository in which to store created disks (VDIs). The list is filtered to show only repositories that can contain virtual hard disks (ISO repositories are filtered out).
- **Packaging Cache Size in GB (optional):** Specifies the maximum amount of space to use within the storage repository for caching. By default, this is 250 GB.
- **Use HTTPS for File Transfers:** Encrypts the image file transfers. HTTPS is selected by default for more secure uploads and downloads but can be cleared for increased performance. This does not apply when you enable Offload Compositing.
- **Offload Compositing (recommended):** Offloads compositing and other operations to a VM created within XenServer. This feature increases performance and enables features such as UEFI, Secure Boot, and Image Portability support. Disabling this is deprecated.
- **ISO Storage Repository:** Repository in which to store ISOs that are uploaded when Offload Compositing is enabled. The list is filtered to show only ISO repositories. SMB and NFS are supported.
- **ISO Share Path:** Automatically populates for selected ISO storage repository by the ISO share path configured. For display only.
- **ISO Share Username:** User name for the selected ISO Share. Only valid for SMB ISO Share. NFS ISO Share does not support a user name or password.
- **ISO Share Password:** Password for the selected ISO share. Only valid for SMB ISO Share. NFS ISO Share does not support a user name or password.

Virtual Machine folder

Virtual machines created by the XenServer connector, whether packaging machines or layered images, can use either folders or tags to organize the virtual machines. XenServer allows you to organize virtual machines by folder or by tag. These organizational tools are optional when creating and managing virtual machines through XenCenter or other tools. Although XenServer connector configurations do not allow you to specify folders or tags, the virtual machines created by the XenServer connector, both packaging machines and published layered images, can use both organizational tools.

Caching tags

If the template specified in the XenServer connector configuration has any tags, then those tags are carried over to any virtual machine cloned from that template. All packaging virtual machines or published layered images are tagged with the same tags that the template has. Also, the XenServer connector adds three tags.

- **App Layering** - All virtual machines created by the XenServer connector can be found by this tag regardless of their purpose or image.
- **Purpose Tag** - All packaging machines are tagged with **App Layering Packaging Machine** while all published layered image virtual machines are tagged with **App Layering Published Images**.
- **Image or Layer Name** - All packaging machines are tagged with the layer name for the layer for which they are generated, while all published layered images are tagged with the template image name.

If you are using XenCenter, you can view your virtual machines by tag by selecting **Organization Views** and then selecting **By Tag**.

Folder

By default, virtual machines created by the XenServer Connector are not placed in a folder. However, if the template specified in the XenServer Connector configuration resides in a folder, then any virtual machine that the Connector creates from that template also resides in the same folder. All packaging VMs and published layered images are placed in that same folder. There are no separate subfolders for packaging VMs or published layered images.

Machine network connectivity

The virtual network settings of the source template specified in the XenServer Connector configuration are carried over when creating any VMs through the XenServer Connector. There is no option in the Connector configuration UI to override the network settings.

XenServer pools

The XenServer Connector does not work correctly with XenServer pools if Offload Compositing is disabled. If the host specified in the configuration is part of a pool, then it must be the pool coordinator for the connector to work. However, this means that anytime the pool coordinator host goes down and a new pool coordinator is elected, the XenServer configuration must be updated.

To add a Connector Configuration

1. Click the **Connectors** page.
2. Click **Add Connector Configuration**. A dialog box opens.
3. Select the **Connector Type** for the platform and location where you are creating the layer or publishing the image.
4. Click **New** to open the Connector configuration page.
5. Type the configuration **Name**, XenServer address, user name, password, and setting for the **Use Secured Communications** check box. For more information, see the above field definitions.
6. Click **CONNECT** below the XenServer configuration fields. The **Virtual Machine Clone Settings** fields are then enabled if the user name and password are correct.

Note:

If there is a certificate error, the following error message is displayed:

One or more problems with the service certificate were found

. You can enable them to be ignored, or you must update the certificate on the server.

You can click **Ignore Certificate Errors and Continue**.

7. Select the required **Virtual Machine Template**.
8. Select the **Storage Repository**.
9. Ensure **Packaging Cache Size in GB** is set to the desired value (or use default). Setting the value to 0 results in no caching.
10. Select the setting for **Use HTTPS for File Transfers**.
11. Select the setting for **Use Offload Compositing**.
12. If **Use Offload Compositing** is selected, select **ISO Storage Repository**. The **ISO Share Path** is auto-populated.
13. If an **SMB ISO** share is selected, enter the **SMB ISO** share user name and password.
14. Click **CONFIRM AND COMPLETE**. A configuration summary is displayed.

15. Click **Save**. If no errors are displayed, the new connector configuration is saved and displayed on the **Connector** page.

Citrix Provisioning™ (XenServer, VMware, Hyper-V, Nutanix, Azure)

January 5, 2026

The Citrix Provisioning connector enables you to publish layered images to Citrix Provisioning for use with targets running on MS Hyper-V, VMware vSphere, Nutanix AHV Acropolis, XenServer, or Azure.

Before you start

The Citrix Provisioning connector configuration requires an account that the App Layering appliance can use to access the virtual machine where you are creating a layer or publishing layered images.

We recommend using the same VM hardware for creating the platform layer as you do for creating the Target Devices in Citrix Provisioning. In most cases, this means selecting the VM template used to create Target Devices in the connector configuration used to create the layer. This practice ensures that the published image and the target devices have the same baseline VM specs. For more information about Connectors and Connector configurations, see [Connector configurations](#).

Note:

The Citrix Provisioning Connector configuration does not currently support Image Portability features.

To publish an image for PVS, run the publish task using the [Windows File Share](#) connector configured with a share that corresponds to the PVS server's store location. This will create a VHD(x) containing the image within the PVS store. You can then register the disk with PVS manually, using the PVS console or PowerShell cmdlets.

For more information on Image Portability, see our [Image Portability](#) documentation.

Layering

Publishing Prerequisites

Refer to the following for prerequisites when publishing a layered image.

- **Install the App Layering agent on Citrix Provisioning servers**

Install the agent on Citrix Provisioning servers, and wherever your connector is configured to run scripts. Register each of the agents with the App Layering appliance.

- **Install the Citrix Provisioning console where the agent is installed**

The Citrix Provisioning console must be installed on all Citrix Provisioning servers where the agent is installed.

- **Install PowerShell Snap-in**

Install the appropriate PowerShell Snap-in.

- **Unique CMID for each target device (if using KMS)**

When using KMS licensing, Citrix Provisioning requires that each target device has a unique CMID. Rearming KMS is covered in the steps to [Create a Platform layer](#).

- **More Citrix Provisioning settings to use in your environment**

Configure Citrix Provisioning on your platform layer. Install the version of the Target Device software that matches the version of the server where you are publishing images. Make sure that the settings match the environment where the layered image is to be used.

- **Citrix Provisioning for Hyper-V**

Requires a Legacy Network Adapter to Pre-Boot Execution Environment boot.

Citrix Provisioning requirements

- Domain accounts have permission to access the Citrix Provisioning store and the local system account does not. If your Citrix Provisioning server is configured to use the local system account, which is the default setting, you can change the account by running the Citrix Provisioning configuration wizard. The wizard gives you the option to run as the **local system** or use a **domain account**. Choose a **domain account**.
- The domain user account in the connector configuration must be in the local Administrators group on the Citrix Provisioning server.
- Citrix Provisioning server and account information - For App Layering to access the location in your Citrix Provisioning environment where you want to publish a layered image, you supply the credentials and location in a Citrix Provisioning connector configuration.
- The App Layering agent must be installed on each of your Citrix Provisioning servers. For details, see the agent installation instructions.

Citrix Provisioning connector configuration

The information you need for the Citrix Provisioning connector configuration includes.

- **Config Name:** A useful name to identify and keep track of this connector configuration.

Citrix Provisioning Server Configuration

- **Console:** The *name* of the Citrix Provisioning server on which the App Layering agent is deployed. This is the server to which the Personal vDisk is published.

Note:

The host name is required, rather than the FQDN so that the Citrix Provisioning server can access the App Layering appliance if it is on a different domain.

- **Domain User:** User name of a domain account that has permission to manage Citrix Provisioning. This account is used by the agent to run Provisioning Services PowerShell commands. This account must have **Read/Write** access to the Citrix Provisioning store for writing the published Personal vDisk.
- **Password:** The password for the domain user account.

vDisk Settings

- **Site Name:** Name of the Site this Personal vDisk is to be a member of.
- **Store Name:** Name of the Store that this Personal vDisk is a member of.
- **Write Cache:** When a new Disk is being created, this value sets the **Write Cache** type of the new Disk. Possible values include:
 - Cache on Server
 - Cache on Server, Persistent
 - Cache in Device RAM
 - Cache in Device RAM with Overflow on Hard Disk
 - Cache on Device Hard Drive

Important:

When choosing a **Write Cache** option, see [Selecting the write cache destination for standard Personal vDisk images](#) to ensure that the Citrix Provisioning servers and target devices that use this Personal vDisk are properly configured for the type you select.

- **License Mode:** Sets the Windows License Mode to:
 - KMS - Key Management Service
 - MAK - Multiple Activation Keys
 - None
- **Enable Active Directory machine account password management:** Enables Active Directory password management. The default value is **Enabled**.

- **Enable Load Balancing:** Enables load balancing for the streaming of the Personal vDisk.
- **Enable Printer Management:** When enabled, invalid printers are deleted from the Device.

Compositing Settings Offload Connector Configuration: A hypervisor connector configuration with **Offload compositing** enabled. This connector configuration composites layer on behalf of the Citrix Provisioning connector. The virtual machine settings used by the offload compositing engine are from this connector configuration. For example, if the Offload Connector Configuration is set up to create UEFI machines, the resulting vDisk is in UEFI format.

Disk Format: The Disk Format of the Citrix Provisioning vDisk on the Citrix Provisioning Server. The format specified here overrides the format in the associated Offload Connector Configuration.

File Share Path: The UNC path corresponds to the Citrix Provisioning Store selected in the vDisk Settings. Requirements include:

- If the Citrix Provisioning Store does not point to a UNC File share, configure the local path as an SMB share.
- The File Share Path is accessible to the compositing engine and selected Citrix Provisioning Store.

If you change the Store selection when Offload Compositing is selected, the connector attempts to resolve the File Share Path. If the File Share Path cannot be resolved automatically, it remains blank.

Script configuration (Optional, advanced feature)

When creating a connector configuration, you can configure an optional PowerShell script on any Windows machine running an App Layering agent, the same agent used on the Citrix Provisioning server. Store these scripts on the machine where the App Layering agent is installed. Only run the scripts after a successful deployment of a layered image. Some preset variables are available to enable scripts to be reusable with different template images and different connector configurations. These variables also contain information needed to identify the virtual machine created as part of the published layered image in Citrix Provisioning.

Running the scripts do not affect the outcome of the publish job, and the progress of commands run in the script aren't visible. The Citrix Provisioning connector logs contain the output of the script that ran.

Configure a script

Remember that this procedure is optional. If you want a script to run each time a layered image is published, complete these steps using the values described in the sections that follow.

1. Complete and save the connector configuration.

Note:

Before selecting the Script configuration page, you must save (or discard) any edits to the connector configuration settings,

2. If the Navigation menu on the left is not open, select it and click **Script Configuration** to open the Script Path page.
3. Complete the required fields, and click **Save**. Field descriptions follow.

Script Configuration fields

- **Enable script:** Select this check box to enable the remaining fields. This allows you to enter a script that runs each time a Layered Image is published.
- **Script Agent:** The agent machine where the scripts are located and run from.
- **Username (optional):** The user name to *impersonate* when running the script. This name can be used to ensure the script runs in the context of a user that has the needed rights/permissions to perform the operations in the script.
- **Password (optional):** The password for the specified user name.
- **Script Path:** A full path and file name on the agent machine where the script file resides.

Other Script Configuration values

PowerShell variables Use any of these Variables in the PowerShell script:

Value	Applies to connector types	Value determined by which code	Description
connectorCfgName	All	Common code	The name of the connector configuration with which the script configuration is associated.
imageName	All	Common code	The name of the layered image template that is used to build/publish the layered image.

Value	Applies to connector types	Value determined by which code	Description
osType	All	Common code	The OS type of the published layered image. It can be one of the following values: Windows7; Windows764; Windows200864; Windows201264; Windows10; Windows1064
diskLocatorId	All	Provisioning Services	The internal ID for the Personal vDisk.

User Impersonation The App Layering Agent, which runs as a service on a Windows machine, runs under either the local system account or the network account. Either of these accounts can have some special privileges, but they are often restricted when running specific commands or seeing files in the file system. Therefore, App Layering gives you the option of adding a domain user and password that can be used to “impersonate” a user. This means that the script can be run as if that user had logged on to the system so that any commands or data are accessible subject to those user rights and permissions. If a user name or password is not entered, the script runs using the account under which the service is configured to run.

Script Execution Policy Script execution policy requirements are up to you. If you intend to run unsigned scripts, you must configure the execution policy to one of the more lenient policies. However, if you sign your own scripts, you can choose to use a more restrictive execution policy.

Google Cloud

January 5, 2026

A connector configuration contains the credentials that the appliance uses to access a specific project on Google Cloud. You need a connector configuration for each Google Cloud project that you want to access with the appliance.

Creating layers on Google Cloud requires a Google Cloud connector configuration. This article describes the values required for the connector. See [Connector configurations](#) for more about configurations and how to create them.

Note:

The **Google Cloud** connector is currently not compatible with the Image Portability Service feature.

Before you start

Google Cloud requirements

This section explains:

- The Google Cloud account information required to create this connector configuration.
- The Google Cloud storage you need for App Layering.

Required Google Cloud Service Account and Service Account Key The Google Cloud connector configuration requires the following information.

- **Project** - The Project Id of a Google Cloud project.
- **Service Account Key File** - For making API calls as the service account on behalf of the connector configuration.
- **Storage Bucket:** A storage location in Google Cloud for storing virtual disks uploaded by the connector.
- **Instance Template:** A Google Cloud VM template with the desired settings for creating a virtual machine.
- **Disk Type:** The type of [Google Cloud storage](#).
- **Zone:** The Google Cloud Zone where you plan to create layers or publish images using the connector configuration.

Required Google Cloud storage bucket Any account you use for App Layering must meet the following requirements:

- Must be separate from the storage bucket used for the appliance.
- Must be in the Google Cloud location where you plan to deploy virtual machines.

App Layering considerations

OS layer If you plan to use an OS image created on another platform, be sure to prepare to:

- Export the OS layer from the App Layering appliance running on another platform.
- Import the OS layer using the App Layering appliance running on the Google Cloud.
- Add a version to the imported OS layer, to prepare it to run on the Google Cloud.

Platform layer To create layers or publish layered images on Google Cloud, you only need a platform layer when you are publishing images to a provisioning service, such as Citrix® machine creation.

When you create a platform layer, there is no need to install Google Cloud tools on it. When Google Cloud tools aren't present, Google Cloud installs them onto the packaging machine when the machine is started.

When you do create a platform layer, any software installers you need (for example, Provisioning Service software) must be accessible from the packaging machine.

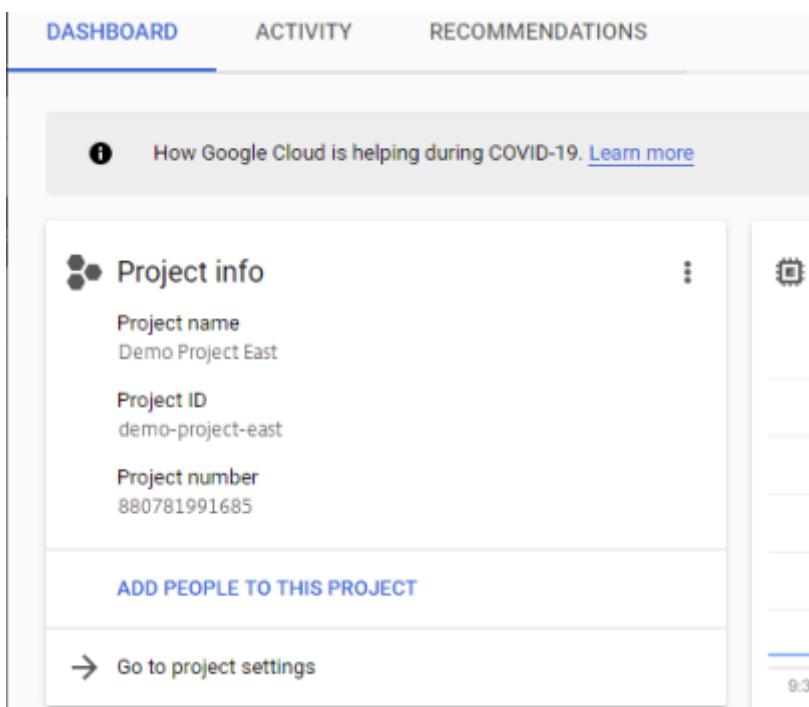
Notes:

- If you plan to publish images to Google Cloud, use the **Machine Creation for Google Cloud Connector**, which can also be set up with the steps below.
- Both the **Google Cloud Connector** and **Machine Creation for Google Cloud Connector** are not compatible with the Image Portability Service feature.

Create a Google Service Account and Service Account Key File

Use the following procedures for each Google Cloud project that you want to connect with the App Layering appliance.

1. Log into **console.cloud.google.com**.
2. Select the project, and click **Go To Project Settings**.



3. On the left pane, **click Service Accounts**.
4. At the top of the page, click **+ CREATE SERVICE ACCOUNT**.
5. Add the details for your service account
 - Name for this service account: Descriptive name. For example, TestEast1
 - Service Account Id: filled in with what you put in the name field, use that value.
6. Click the **Create** button.

Add the roles for your service account

In the **Select Role** box, select the roles required:

- Storage Admin
- Compute Admin
- Service Account User

1. Click the **Continue** button.
2. Click the **Done** button. User access to the role is not necessary.

Create a Service Account Key File

1. From the project, select the **Service Account** tab on the left.

2. Click the three vertical dots to the right of the service account you want to create the key file for.
3. Select **Create key**.
4. On the popup, select the JSON radio button and click the **CREATE** button.
5. When you are prompted to save the key to a file, enter the name of your choice or keep the default name. You are returned to the Service Accounts screen.
6. Click your service account to verify that the key is there.

Edit Service Account Roles

To edit the roles:

1. Navigate to the project UI screen as you did when creating the service account.
2. From the hamburger menu on the top left, expand the **IAM & ADMIN** and select **Manage Resources**.
3. Select your project, and on the right side type the **Service Account** you want to edit into the Filter Tree. The service account name is displayed.
4. Edit each of the roles to include or to remove from the roles and save the results.

Note:

If you remove all the privilege from the service account, it will no longer show on the IAM page (because IAM page only lists those account with at least one roles attached). But it will show up on Service Account screen. You would need to go to IAM page to add privileges back.

5. Click the **ADD** button.
6. On the **Add members, roles to “current” project** panel, in the **New members** field, type the name of the member to whom you want to give privileges.
7. Select the roles you want to give the member, and click the **SAVE** button.

MS Hyper-V

January 5, 2026

An MS-Hyper-V connector configuration includes the credentials and storage location the appliance needs to connect to Hyper-V, and it identifies the properties to be associated with the vDisk.

You can select a Hyper-V connector for importing the OS, creating other layers, or publishing layered images. Each connector configuration is set to access a storage location using a specific account.

The Hyper-V connector uses Microsoft's Background Intelligent Transfer Service (BITS) to copy disks to and from the appliance. With BITS the appliance reports progress as a percentage complete, and the connector no longer requires a CIFS share.

The Hyper-V connector includes an **Offload Compositing** checkbox that enables layer packaging and image publishing to be done on the Hyper-V server, rather than on the App Layering appliance. Offload Compositing greatly increases the speed of layer packaging and image publishing. It also automates layer finalization and lets you create layers and publish images as Hyper-V Generation 2 machines on VHD or VHDX disks.

For more about connectors and connector configurations in general, see [Connector configurations](#).

Note:

The **Microsoft Hyper-V** connector is currently not compatible with the Image Portability Service feature.

Before you start

The first time you create an image template for publishing layered images to your Microsoft Hyper-V environment, you need to add a new Hyper-V connector configuration for that location.

Hyper-V requirements

You can use a Hyper-V connector configuration makes it easy to connect to a location in your Microsoft Hyper-V environment.

The information you need for the Hyper-V connector configuration includes:

- **Hyper-V Configuration**
 - **Agent** - App Layering agent from the list of agents registered with the appliance. The agent runs under the machine account for the machine it is running on.
- **Virtual Machine Settings**
 - **Template VM (Optional)** - A template that can be used to clone a Hyper-V virtual machine with the desired hardware settings (memory, CPUs and video settings). You can specify the host, data store and network for configuring the resulting virtual machines. The template must have at least one network card attached, and it must not have any disks attached. Otherwise, you receive an error when attempting to validate or save the configuration.

- **Number of CPUs** - Number of CPUs to use for creating a Packaging Machine or publishing a layered image. The default value is 4 CPUs.
- **Memory (Mbs)** - Amount of memory allocated for creating the Packaging Machine or the layered image Machine. The default value is 8192 Mb.
- **Network** - Network switch. You can select from a list of network switches known by the agent.
- **Generation** - Generation 1 machines are supported in all cases. Generation 2 machines are supported only when the Offload Compositing option is selected.
- **Disk Format** - VHD or VHDX are supported on Generation 1 machines. Only VHDX is supported on Generation 2 machines.

Note:

VHDX disk format requires **Offload Compositing**, even on Generation 1 machines.

- **Offload Compositing** - Enables the layer packaging or image publishing process to run on the specified Hyper-V server. This feature increases performance, and it allows you to use VHDX disk format and Generation 2 VMs. With UEFI, you can also use **Secure Boot** if it is enabled on the VM.

When Offload Compositing is selected:

- If you do not provide a virtual machine template, the virtual machine defaults to BIOS.
- If you provide a template configured for BIOS or UEFI, the resulting virtual machine is the type you chose.
- If you provide a template with UEFI-Secure Boot enabled and selected, the resulting VM is UEFI-Secure Boot.

When Offload Compositing is not selected:

- If you do not provide a template, the virtual machine defaults to BIOS.
- If you provide a template configured for BIOS, the resulting virtual machine is BIOS.
- If you provide a template configured for UEFI, the machine fails to boot, and results in a blue screen. (Offload Compositing is required for UEFI.)

- **Storage Settings**

- **Path to Storage, eg: C:\Virtual Machines** - Path for the local or remote location where the App Layering software creates layer disks, packaging machines, and layered image disks. This value:
 - Must be the same as the location specified in the UNC path.
 - Cannot be a mapped drive.

- **Layer Disk Cache Size in GB (optional)** - Amount of File Share space (in gigabytes) to use for caching layer disks. A value of:
 - 0 disables layer caching.
 - 1+ enables layer caching and specifies the amount of space to allow for caching layer disks on the File Share.Copies of boot disks and packaging disks are stored and then reused wherever possible to reduce the time it takes to package an app layer.

App Layering requirements

The App Layering agent is required to use a Hyper-V connector. The App Layering agent must be:

- Installed on the Microsoft Hyper-V server where you want to create layers or publish layered images. For details, see [Install the App Layering agent](#) in the App Layering installation topic.
- Registered with the App Layering appliance. For details, see [Register with the App Layering appliance manually](#) in the App Layering installation topic.

Additional considerations

Generation 2 machines To start using Generation 2 machines in Hyper-V when your existing machines are Generation 1:

- Create a Hyper-V connector configuration with the **Offload compositing** feature enabled and Generation 2 selected.
- Create a Generation 2 platform layer.
- Publish Generation 2 images using this platform layer and your existing Generation 1 app layers.

If you are starting with a fresh deployment in Hyper-V, you can create a Generation 2 OS layer using either of the following approaches:

- Create the OS layer from a Generation 1 OS image. Then:
 - Create a Hyper-V connector configuration with the **Offload compositing** feature enabled and **Generation 2** selected.
 - Add a Generation 2 version to the OS layer.
 - Create Generation 2 app layers.
 - Create an image template with the new connector selected, and publish the images.
- Create the OS layer from a Generation 2 OS image, by using the [OS import script](#):
 - Locate the ImportOSLayer.ps1 script included in the OS Machine Tools download.
 - Import the OS using the script. The script supports the import of UEFI machines and completes the import faster than the management console.

Machine Creation for Hyper-V

Although there is no Machine Creation version of the Hyper-V connector, you can use a Hyper-V connector configuration to publish layered images to a specific machine creation location within your MS Hyper-V environment.

Create a Hyper-V connector configuration

To use the Hyper-V connector for layering or publishing, you:

1. Click the **Connectors** page.
2. Click **Add Connector Configuration**. This opens a small dialog box.
3. Select the **Microsoft Hyper-V** Connector Type. Then click **New** to open the Connector Configuration page.
4. Complete the fields on the Connector Configuration page. For guidance, see the field definitions in the Hyper-V requirements section above.
5. Click the **TEST** button to verify that the appliance can access the location specified using the credentials supplied.
6. Click **SAVE**. The new connector configuration should now be listed on the Connector tab.

HTTPS and certificate errors

By default, HTTPS is turned off and certificate errors are ignored, because the self-signed certificate that comes with the appliance would fail over HTTPS. When you upload your own certificate, you can configure the connector to use HTTPS and to stop ignoring certificate errors.

If you want to enable HTTPS for Hyper-V

If you want to enable HTTPS for Hyper-V, edit the settings for turning on HTTPS and for ignoring certificate errors in the config.json file.

1. Open the Hyper-V connector config.json file:

```
/usr/local/lib/node_modules/unidesk-hyperv-connector/config.json
```

2. Set **useHttpsFileTransfer** to true:

```
“useHttpsFileTransfer”: true
```

3. Set **ignoreCertificateErrors** to false:

```
“ignoreCertificateErrors”: false
```

4. Restart the Hyper-V connector.

Certificate errors

Once certificate errors are no longer ignored, you will receive the following error if your certificate expires:

- 1 Failed copying file to D:\path\file.vhdx. The certificate authority is invalid or incorrect.
- 2 The error occurred **while** the remote file was being processed.

Nutanix AHV (Acropolis)

January 5, 2026

A Nutanix AHV connector configuration contains the credentials and storage location that the appliance needs to connect to Nutanix Acropolis. Use the Nutanix connector to package layers and publish images to Nutanix Acropolis.

Before you start

You can use your Nutanix Acropolis environment to create layers and publish layered images. Each Connector configuration accesses a specific storage location in your Nutanix Acropolis environment. You might need more than one Nutanix Connector configuration to access the correct location for each purpose. Further, you can publish each layered image to a location convenient to the system you provision with the published image. You can also export images from your Nutanix Acropolis environment via the [Image Portability](#) feature. For more information about Connectors and Connector configurations, see [Connector configurations](#).

Layering

To package layers using a Nutanix virtual machine, specify a Nutanix Connector configuration when [creating](#) or [versioning](#) a layer. Specify a Nutanix or Machine Creation for Nutanix Connector configuration in an [image template](#) to have that image created in Nutanix when [published](#).

Image Portability

You can export existing images from your Nutanix Acropolis environment by using the Nutanix Connector in conjunction with the [Windows File Share Connector](#) via the [Image Portability](#) export feature. Specify a Nutanix Connector with Offload Compositing enabled in the **Connector Configuration** dropdown of a Windows File Share Connector and then specify that Windows File Share Connector to the

IPS export command. This will allow Image Portability to access your Nutanix Acropolis environment, and export images to the configured file share location.

Use one of the following to specify to the export command which image to export:

- The name or uuid of a disk image
- The name or uuid of a VM with exactly one disk attached
- The uuid of a virtual disk

Machine Creation for Nutanix Connector

The Machine Creation version of the Nutanix Connector publishes layered images in a state that is ready for use by Citrix Machine Creation Service in your Nutanix Acropolis environment. Published layered images are booted for final configuration, shut down, and then a snapshot is taken.

Notes:

The Machine Creation for Nutanix Connector configuration is for publishing layered images. You cannot package layers with it, nor use it with the Image Portability feature. Instead, use a Nutanix Connector configuration.

The Machine Creation for Nutanix Connector does not support caching

Prerequisites and permissions

Specify the Nutanix Prism Elements console

App Layering uses the Prism Elements web console and does *not* support the Prism Central console.

Important:

When using Nutanix connectors, App Layering requires direct NFS access to the hosts to work correctly. In older versions of Nutanix AHV (5.6 and 5.7), this direct NFS access to hosts was not allowed if a Prism Element host or cluster was registered with Prism Central. Make sure that your Nutanix setup allows this access. For details about this issue on various Nutanix versions, see [Adding layer versions with Nutanix fails with error: Failed to execute the script](#)

When configuring the Nutanix connector be sure to enter the **URL** for the Prism Elements console.

Note:

If Prism Central is used in the connector configuration, you will receive the error “internal error 500”.

Add the Citrix App Layering™ appliance to the Nutanix allow list

Ensure that the appliance is added to your Nutanix allow list so that it can access the appropriate storage containers, as needed. This can be accomplished by configuring the file system and container-level allow list settings. For details about adding an allow list with Nutanix, see the Nutanix documentation.

Permissions

Refer to the following list for required permissions when using the Nutanix Connector configuration.

- A Nutanix Prism account (new or existing) must have privileges to perform the following operations:
 - VM operations:
 - clone
 - delete
 - power on/off
 - attach virtual disks
 - Image operations:
 - create
 - update (aka upload)
 - delete
 - Virtual disks:
 - create
 - attach to VMs

Required information for Nutanix connector configuration settings

Refer to the following required, recommended, and optional fields for configuration of your Nutanix Connector.

- **Connector Configuration Name:** A useful name to help identify this connector configuration.
- **Web Console (Prism) Address:** The host name (resolvable via DNS) or IP address of the Prism Web Console. This address is the same one that you use to access the Nutanix Prism Web Console.
- **User Name/Password:** The credentials for the account that the appliance uses to connect to the Nutanix system.
- **Virtual Machine Template (recommended):** Virtual Machine Template that can be used to clone a VM with the hardware settings for Nutanix, including memory, CPUs, and video settings.

You can specify the host, datastore, and network for configuring the resulting VMs. Since there is no concept of a “template” at Nutanix, these “templates” are actual VMs. The OS version used by the selected “template” must match the OS version that you are using for building layers or publishing layered images. The template must not have any disks attached and must have at least one network card attached. If it does not, you see an error when trying to validate or save the configuration.

- **Storage Container:** Allows you to select the storage container for the images (virtual disks, VHDs) that are uploaded, and the resulting virtual disks that are created from those images. When creating app layers and OS layer versions, mount the storage container as an NFS mount point.

Configure the **allow list** using the Nutanix web console or Nutanix CLI tools. Set the allow list to the cluster and every storage container on the cluster, even the ones you are not using.

Note: If the appliance is not allow-listed for the selected storage container, the validation phase fails, and the error is indicated with the storage container selection.

- **Layer Disk Cache Size in GB (optional):** Specifies the size of the cache allowed for each layer.
- **Offload Compositing:** Enables the layer packaging or image publishing process to run on the specified Nutanix server. This feature increases performance and it allows you to use a native disk format and either BIOS or UEFI virtual machines. This is enabled by default.
- **Packaging Cache Size in GB (recommended):** Amount of cache size space (in GB) to use for packaging. Accept the recommended value or modify it.

How Virtual Machines are Organized

Nutanix does not provide a mechanism for organizing virtual machines. Because of this, it could be difficult to find the virtual machines created by your appliance when the total number of virtual machines is large. To help you find these VMs, the following naming conventions are used:

- **Packaging Machines** (virtual machines created during the process of creating an App Layer or OS Version)
 - The virtual machine name starts with the layer name that is being created/modified
 - The virtual machine names end with the following text: (Packaging Machine)
- **Layered Image Virtual Machines** (virtual machines created as a result of publishing a layered image)
 - The virtual machine name starts with the image name that was published
 - The virtual machine name ends with the following text: (Published Image)

When viewing virtual machines through the Nutanix web console, you can search for virtual machines by filtering on:

- “Citrix App Layering” to find all virtual machines created by the App Layering service.

- “Citrix App Layering Packaging Machine” to find all virtual machines created for layer management jobs.
- “Citrix App Layering Published Image” to find all virtual machines created to publish a layered image.
- Image name or layer name to find virtual machines related to a specific layered image publishing job or App or OS creation.

Virtual Machine Network Connectivity

The virtual network settings of the source template specified in the Nutanix AHV connector configuration will be carried over when creating any VMs through the Nutanix Acropolis Hypervisor (AHV) Connector. There is no option in the Connector Configuration UI to override the network settings.

To add a connector configuration

1. Log in to the management console as an administrator.
2. Select the **Connectors > Add connector configuration**.
3. Select **Nutanix AHV** from the connector **Type** drop-down menu and click **New**. This opens the connector configuration.
4. Enter the configuration **Name**, and the Acropolis Address, User Name, and Password. For guidance, see the above field definitions.
5. Click the **Connect** button below the Acropolis Configuration field. The **Virtual Machine Clone Settings** field is then enabled if the connection is successful. Any connection problems are reported on the connector configuration blade. If there were server certificate errors found, you see an **Ignore Certificate Errors and Continue** button.
6. Select the Virtual Machine Template.
7. Select the Storage Repository.
8. Click **Confirm and Complete**. If there are no errors, a summary page is displayed.
9. Click **Save**. Verify that the new connector configuration is listed on the **Connectors** page.

VMware vSphere

January 5, 2026

A VMware vSphere connector configuration contains the credentials and storage location that the appliance needs to connect to vSphere. Use the vSphere connector to package layers and publish images to VMware vSphere or VMware Cloud on AWS.

Before you start

You can use your VMware vSphere environment to create layers and publish layered images. Each connector configuration accesses a specific storage location in your VMware vSphere environment. You might need more than one VMware vSphere connector configuration to access the correct location for each purpose. Further, you can publish each layered image to a location convenient to the system you provision with the published image. You can also export images from your VMware vSphere environment via the [Image Portability](#) feature. For more information about Connectors and connector configurations, see [Connector configurations](#).

Layering

To package layers using a VMware vSphere virtual machine, specify a VMware vSphere connector configuration when [creating](#) or [versioning](#) a layer. Specify a VMware vSphere or Machine Creation for VMware vSphere connector configuration in an [image template](#) to have that image created in VMware vSphere when [published](#).

Image Portability

You can export existing images from your VMware vSphere environment by using the VMware vSphere connector in conjunction with the [Windows File Share](#) connector via the [Image Portability](#) export feature. Specify a VMware vSphere connector with Offload Compositing enabled in the **Connector Configuration** dropdown of a Windows File Share connector and then specify that Windows File Share connector to the IPS export command. This will allow Image Portability to access your VMware vSphere environment, and export images to the configured file share location. Use the datastore path of a vmdk (e.g. `[vsanDatastore] d6650865-56fd-065b-20f6-e4434b003954/myimage.vmdk`) to specify to the export command which image to export. The vmdk must be on the same datastore as the one specified in the connector configuration.

Machine Creation for vSphere Connector

The Machine Creation version of the VMware vSphere connector publishes layered images in a state that is ready for use by Citrix Machine Creation Service in your vSphere environment. Published layered images are booted for final configuration, shut down, and then a snapshot is taken.

Notes:

The Machine Creation for vSphere connector configuration is for publishing layered images. You cannot package layers with it, nor use it with the Image Portability feature. Instead, use a VMware vSphere connector configuration.

The Machine Creation for vSphere connector does not support caching

Additional considerations

Virtual controllers

You can use either the default LSI Logic SAS controller, or a VMware paravirtual SCSI controller.

To use the default LSI Logic SAS controller, simply select it for the layer's virtual machine, and make sure that all of your layers use the same controller.

To use a VMware paravirtual SCSI controller, you need a pre-existing Template VM with a VMware Paravirtual SCSI controller and without any disks.

To use an existing LSI OS Layer with a VMware Paravirtual SCSI controller If you have an OS layer with an LSI Logic SAS controller, and you want to use it with a VMware Paravirtual SCSI controller you can use either of the following approaches:

- Add a version to the OS layer, using a VMware vSphere connector with an LSI Logic SAS VM template. When the packaging machine is created, follow the steps below to make the OS layer Paravirtual enabled.
- Add a new platform layer with an LSI OS layer, and a platform connector with an LSI Logic SAS VM template. When the packaging machine is created, follow the steps below to make the platform layer Paravirtual enabled.

Note:

The following changes must be performed on the OS layer and the platform layer.

When the packaging machine from your chosen approach is ready:

1. Log in to the virtual machine and shut it down.
2. In the vSphere Web Client open the **Edit Settings** page for the packaging machine.
3. Add a new SCSI controller, by selecting **SCSI Controller** from the **New Device** menu, and click **Add**.
4. Expand the **New SCSI controller** section that was added, and set **Change Type** to **VMware Paravirtual**.
5. Add a new hard disk, by selecting **New Hard Disk** from the **New device** menu, and clicking **Add**.
6. Expand the New Hard disk section and set the following parameters:
 - Size: 1 GB
 - Disk Provisioning: Thin provision
 - Virtual Device Node: New SCSI controller default bus

7. Click **OK**.
8. Install the Paravirtual drivers by powering on the packaging machine, logging in, and then shutting down.
9. In the vSphere Web Client, open the **Edit Settings** page for the packaging machine.
10. Remove both the hard disk and the Paravirtual controller that you added earlier in this procedure.
11. Power on the packaging machine, log in, and click **Shut down For Finalize**.

Once you finish creating the layer, you can use it to create an image with a Paravirtual controller.

Required information for vSphere connector configuration settings

Refer to the following required, recommended, and optional fields for configuration of your VMware vSphere connector.

- **Connector Configuration Name**- A useful name to help identify and track this connector configuration.
- **vCenter Server**- The address, hostname, or FQDN of the vCenter server with which the appliance integrates.
- **vCenter User Name**- The user name of the account that the appliance uses to connect to vSphere.
- **vCenter Password**- The password of the account that the appliance uses to connect to vSphere.
- **DataCenter Name**- The name of the vSphere data center in which the App Layering appliance creates and retrieves virtual machines.
- **Packaging Cache Size in GB (Recommended)**- The size of the Disk Cache that App Layering uses when creating layers. If you leave the size blank or set it to 0, App Layering does not use a Disk Cache. If you specify a size, App Layering uses a Disk Cache of up to this size to keep copies of boot disks and packaging disks, and reuses these disks to create packaging machines. The reuse of these boot disks and packaging disks reduces the time that it takes to package an App layer.
- **Virtual Machine Template** - (Optional) Virtual Machine Template that clones a virtual machine with the hardware settings for VMware, including memory, CPUs, and video settings. This setting lets you specify the host, datastore, and network for configuring the resulting virtual machines.

Important:

When publishing to VMware Cloud, a **VMware Virtual Machine template** (not a regular VM template) is required for the virtual machine's network to work correctly.

When selecting a template virtual machine:

- Answer **Yes** to the prompt asking to update settings, but do **not** change the network.
- Make sure that the OS version that's used by the selected template matches the OS version that you're using for building layers or publishing layered images.
- The template must not have any disks attached, and must have at least one network card attached. Otherwise, you receive an error when trying to validate or save the configuration.
- **ESXHost Name**- The name of the vSphereESX Host on which the appliance creates and retrieves virtual machines.
- **DataStore Name**- The name of the vSphere DataStore in which the appliance creates and retrieves virtual machines.
- **Network Name**- The name of the vSphere Network in which the appliance creates and retrieves virtual machines.
- **Virtual Machine Folder Name**- The name of the vSphere Folder in which the appliance creates and retrieves virtual machines.
- **Offload Compositing** - Enables the layer packaging or image publishing process to run on the specified vSphere server. This feature increases performance, and it allows you to use VMDK disk format and either BIOS or UEFI virtual machines. With UEFI, you can also use **Secure Boot** if it's enabled on the VM.

Notes:

- Disabling **Offload Compositing** is a deprecated feature. The ability to disable offload compositing will be removed in a future release.
- When using a vSphere connector configuration with VMware Cloud and a vSAN 7.0 Update 2 (or later) datastore, **Offload Compositing** must be selected.

When Offload Compositing is selected:

- If you do *not* provide a virtual machine template, the virtual machine defaults to BIOS.
- If you provide a template configured for BIOS or UEFI, the resulting virtual machine is the type you chose.
- If you provide a template with UEFI-Secure Boot enabled and selected, the resulting VM is UEFI-Secure Boot.

When Offload Compositing isn't selected:

- If you do not provide a template, the virtual machine defaults to BIOS.
- If you provide a template configured for BIOS, the resulting virtual machine is BIOS.

- If you provide a template configured for UEFI, the machine fails to boot, and results in a blue screen. (Offload Compositing is required for UEFI.)

Required Privileges for the connector's vSphere Client Administrator Role

Set the VMware privileges required by the vSphere Client Administrator to match the permissions for the [App Layering appliance](#).

Once you have set the permissions, verify them by clicking **Save** in the Create Connector configuration summary blade.

Create a connector configuration

To enter values:

- The first three vCenter fields must be entered manually. Once the credentials in those fields are validated, you can select values for the remaining fields from the drop-down menus.
- To enter values manually, click to put the cursor in the field and type the value, making sure that the case matches the value in vCenter.
- To select a value from a drop-down list, click once to put the cursor in the field to display the list of possible values.

To add a connector configuration

1. Click the **Connectors** page.
2. Click **Add Connector Configuration**. A dialog box opens.
3. Select the Connector Type for the platform and location where you're creating the layer or publishing the image. Then click **New** to open the Connector Configuration page.
4. Enter the configuration *Name*, and the *vCenter Server*, *vCenter User Name*, and *vCenter Password*. For guidance, see the previous field definitions.
5. Click the **Connect** button below the vCenter fields. The data center field is then enabled with a list of data centers available.
6. Select the data center, enabling the remaining drop-down lists.
7. Complete the remaining fields and click the **TEST** button to verify that App Layering can access the location specified using the credentials supplied.
8. Click **Save**. Verify that the new connector configuration is listed on the **Connectors** page.

Error messages

If you receive ENOTFOUND errors when you deploy a packaging machine or publish an image, use the IP Address in place of the FQDN for the vCenter server.

Windows File Share

January 5, 2026

A Windows File Share configuration contains the credentials and storage location that the appliance needs to connect to a file share. Use **Windows File Share Connector Configuration** to publish Layered Images.

For more about Connectors and Connector Configurations, see [Connect](#).

Before you start

Ensure that you have the following configurations:

- A connector configuration with Offload Compositing enabled that supports packaging layers.
- A network file share that can be accessed over SMB from packaging machines created by the packaging connector.

If this is your first time using App Layering

If this is your first time using App Layering and you want to publish layered image disks to a network file share, you can create a connector configuration for each of your publishing locations.

When publishing a layered image, you can select a connector configuration. If you don't yet have the right connector configuration for the task, you can create one by clicking **Add Connector Configuration** on the **Connectors** page.

Required information for Windows File Share Connector Configuration settings

In the **Management Console**, select **Connector > Add Connector Configuration > Windows File Share** and click **New**.

To configure the Windows File Share, specify the following values:

- **Config name:** A useful name to help identify and track this connector configuration.

- **SMB File Share Path:** Network share path where layered image disks must be published.
- **User Name:** Name of a user with read and write permissions to the given share path.
- **Password:** Password of the given user.

Compositing Settings

- **Offload Connector Configuration:** A connector configuration with Offload Compositing enabled. This connector configuration composites the layer on behalf of the **Windows File Share** connector. The virtual machine settings used by the Offload Compositing engine are from this connector configuration. For example, if the Offload Connector Configuration is set up to create UEFI machines, the resulting image is in UEFI format.
- **Disk Format:** The virtual disk format to publish images as in the network file share. This can be either VHD or VHDX.

Image Portability

The Windows File Share connector is a required component for using Image Portability. It works in conjunction with other IPS-compatible connectors to support **export** and **publish** operations.

During an **export** or **publish** task, the Windows File Share connector facilitates transferring a platform native image to shared storage as a VHD(x). When configuring this connector for IPS tasks, ensure that the **Connector Configuration** field references an IPS-compatible connector associated with the source environment (the platform from which the image is being exported/published).

In the case of a **publish** task, the shared storage is the target location for creating a PVS vDisk using an image which was previously prepared for use with PVS. This may be the PVS vDisk store Windows folder which has been shared on the network. You then register the disk with PVS manually, using the PVS console or PowerShell cmdlets.

Layering

January 5, 2026

This document provides an overview of layering and the different layer types involved.

Layers

The App Layering architecture lets you manage just one copy of your Windows OS and apps, regardless of your hypervisor. You can maintain one set of apps for two environments. For example, you can

deploy an OS and its app layers in an on-premises hypervisor and in a cloud-based hypervisor. The same layers run on each.

You can create layers for your operating system, platform tools, and applications. To preserve users' settings and data, enable user layers on your image templates.

OS layer

The OS layer includes your operating system and hypervisor software and settings. It is an essential building block for all other layers that you create. You only need one OS layer for a specific Windows OS. For example, if you support both a Windows desktop OS and a Windows Server OS, create one OS layer for each. Layered images can only contain app and platform layers that were built using the OS layer selected for the image.

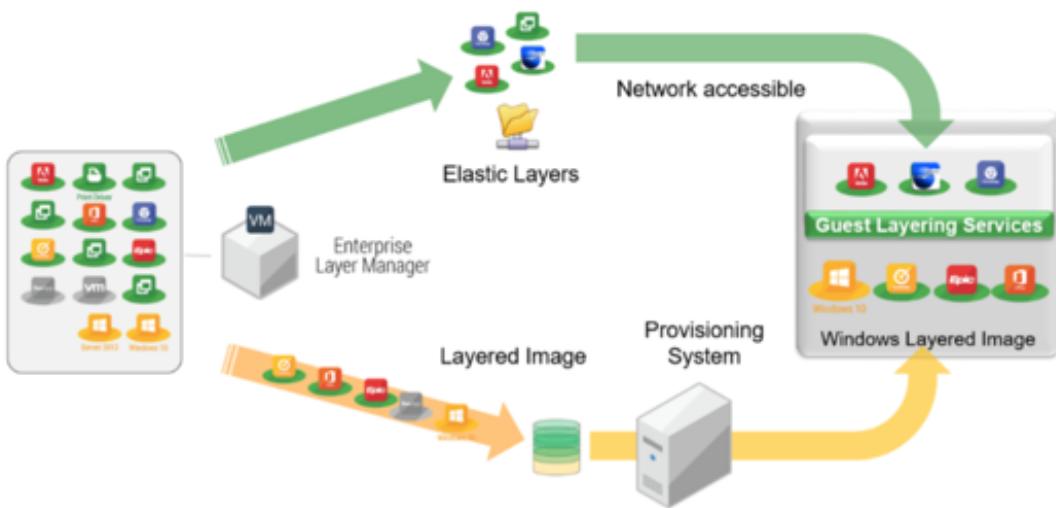
When you add an update to one of the OS layers, the platform and app layers built with that OS continue to run on it.

Platform layer

The platform layer includes the provisioning software and connection broker tools. The platform layer ensures that your OS and app layers run flawlessly in a specific on-premises or cloud environment. You can reuse your OS and App layers, and select a different platform layer for each hypervisor or Provisioning Service.

App layers

App layers include the software for each of your applications. Many App layers can be combined and deployed in different layered images. They also can be assigned to users, groups, or machines and deployed as [Elastic layers](#). Layered images are used to provision users' systems, while Elastic layers are delivered when the user logs in.



Elastic layers

Elastic layers are app layers that are assigned to users, groups, or machines for delivery on demand rather than being built into a layered image. Assignment of elastic layers can be done through the App Layering management console or the Studio App Packages node. See [Elastic layers](#) for more details.

User layers

User layers enable persistence on non-persistent machines by storing a user's settings, files, and installed apps. The layer is created at first login to a machine running a layered image with user layers enabled, and is reapplied whenever the user logs into any machine with user layers enabled. This delivers a consistent, personalized desktop experience while supporting image updates, machine pools, and on-demand provisioning.

User layers are only compatible with single session, non-persistent machines.

Create and manage layers

January 5, 2026

A layer is a virtual disk that contains the software for your operating system, platform tools, apps, or the user's data and settings.

When you create a layer, the appliance saves the new layer as a virtual disk in your hypervisor environment, and attaches the disk to a packaging machine.

Once created, each layer is stored in a repository as a virtual disk.

Types of layers

You can use the following types of layers:

- Layers you create in your hypervisor and include in the image templates you use to publish layered images.
- Layers you enable on image templates, and therefore on the layered images you publish.

Layers to include in image templates and layered images

You can create layers for your OS, your platform tools, and the applications that you want to deliver to users.

- **OS layer:** The layer where you install the Windows OS from ISO. You can reuse the same OS layer with all compatible platform and app layers. We recommend creating just one OS layer for each major Windows version, for example, one for Windows 10 and one for Windows Server 2016. You can add new versions of a layer for each follow-on release. For example, if you have a Windows 10, version 1709 layer, you add a version to it for version 1809, and one for 1903. When you update the OS layer, it is not necessary to update the app layers, but *do* update the platform layer. For more about creating an OS layer, see [Create the OS layer](#).
- **Platform layers:** A layer where you install and configure the software for a specific on-premises or cloud environment. When you isolate your infrastructure software in a platform layer, you can reuse the same OS layer and app layers on multiple hypervisors. You can create a platform layer for each part of your infrastructure if for example you use more than one hypervisor.

If you use any of the following software in your infrastructure, you normally install them on a platform layer:

- Connection broker software
- Provisioning software
- System Center Configuration Manager (SCCM)

Though it might not seem intuitive, it is crucial to install the software for the first *hypervisor* you support on the *OS layer*.

If you support more than one hypervisor, you can then create another platform layer for the additional hypervisor. The platform layer for an additional hypervisor must include the hypervisor software, along with the provisioning, connection broker, and SCCM software. When you create this additional platform layer, be sure to select the check box for the new hypervisor you

are supporting. This ensures that the App Layering appliance removes the original hypervisor files and settings from the OS layer so that they do not interfere with performance.

For more about creating an OS layer, see [Create platform layer](#).

- **App layers:** The layers where you install applications. Typically, we recommend installing one app on each layer, though you can include more. For easy maintenance, include apps that are on the same update schedule. If an application requires other apps, create the layer for the required application first. For more about creating an app layer, see [Create or clone an app layer](#). For tips about layering a specific application, see [App Layering Recipes](#).

Delivering layers elastically

Besides the layers that you include in layered images, you can enable Elastic and User layers on them through settings in the image template:

- **Elastic layers:** App layers that are assigned to specific users and delivered when the users log in. An elastic app layer is not included in the base image, but is delivered on it. Elastic apps appear on the user's desktop.

Important:

An app layer can be delivered to a user either as part of the layered image or as an elastic layer.

There are a few applications that cannot be used as elastic layers, for example, Microsoft Office. To find out whether an application has this limitation, check the App Layering recipes [here](#). If no limitations are specified for your app, you can assign it as an elastic layer. For more about enabling and assigning elastic layers on a layered image, see [Deploy app layers as elastic layers](#).

- **User layers:** Enabling user layers on a layered image allows you to persist a user's data and settings, and any applications that they install themselves. When enabled, a user layer is created for each user the first time they log on to an image. To enable this feature, select the **User layers** setting in the image template that you use to publish the layered image. For more about enabling elastic layers on a layered image, see [Deploy user layers](#), and [Create or clone an image template](#).

Don't assign application layers that have browsers like Chrome or Firefox to users when those users are logging into an image that is using full user layers. Browsers update frequently, which can cause revision conflicts between the user's writeable layer and the revision on the bootable image. You can, however, assign those layers to users if they are logging in to an image that is a session host.

Applications that require a local user or administrator

The OS layer preserves any local users or groups that you add, but app layers, platform layers, and user layers do not. For example, users and groups that you add or change while installing an application on an app layer, platform layer, or user layer do not persist. You can either:

- Add the local user or administrator to the OS layer before installing the application.
- Install the application on the OS layer.

Layer integrity overview

When creating an OS, app, or a platform layer, you begin layer creation in the App Layering management console, then install the software in the specified VM in your hypervisor. When the layer is in the state that you want it to be in when users start their desktops, you shut down the machine and finalize the layer.

When you shut down a layer to finalize it, Windows `Ngen.exe` operations display messages about pending tasks that must be completed before shutting down. You must let these jobs complete, but you can expedite the `Ngen.exe` operations, if necessary. Details about `Ngen.exe` messages and how to expedite operations are included in each of the related layering articles:

- [Prepare your OS image for layering in XenServer, Hyper-V, or vSphere](#)
- [Prepare your OS image for layering in Azure](#)
- [Prepare your OS image for layering in Nutanix](#)
- [Create platform layer](#)
- [Create or clone an app layer](#)
- [Update layer](#)
- [Troubleshoot layer integrity issues](#)

Layer priority

Layer priority defines layer order when creating the Windows file system and registry. Layer priority is important when:

- Compositing layers as part of publishing layered images.
- Searching layers for file and registry settings.
- Delivering elastic layers and user layers to users' desktops.

The App Layering software assigns a priority to each layer, and applies the layers in order, from the lowest priority to the highest.

In Windows, the highest priority layer takes precedence. If a file or registry entry exists in two layers, Windows uses the file or registry entry from the layer with highest priority.

How layer priority is determined

A layer's priority is based on the layer type and, for App layers, the order in which the layers were created.

Layers within the base image Layers that are part of the layered image are applied in order, with the Platform Layer always applied last, as the highest priority layer.

As the following table shows, the priority assigned to app layers is based on the order in which the layers are created. The newest app layers are given a higher priority than older layers.

Priority	Layer type
High	Platform layer
	App layer created last
Medium	App layers in order by creation date
	App layer created first
Low	OS layer

If the layers have a file or registry entry in common, the file or registry entry from the higher priority layers are used.

Layers enabled on the base image When a published image boots, more layers can be applied, if the layers are enabled in the image template for your layered image:

- Elastic layers (app layers assigned to users as elastic layers)
- User Layers

When merging layers onto an image, User layers are always the highest priority. Elastic layers are next, and the layers in the base image last.

As shown in the following table, the priority of elastic layers is the same as the priority of the original app layers, but applied to the base image. Elastic layer priority does *not* depend on the order in which the layers are attached to the published image.

Priority	Layer type
High	User layer
	Elastic layer - App layer created last

Priority	Layer type
Medium	Elastic layers - App layers in creation order
	Elastic App - App layer created first
Low	Layered image - All layers within base image

Layer priority conflicts

Most app layers work, but in some situations, the order in which you install applications can cause conflicts on the desktop.

If one app must be installed before another, create the layers in the order required. The App Layering software applies the layers in the same order.

If two layers conflict and you suspect that it is due to the order in which they are incorporated into the image, you have two choices:

- Recreate the layer that you want to install last, so that it is incorporated in the correct order.
- Request assistance from Technical Support.

Prepare the OS for layering

January 5, 2026

You can prepare your operating system for layering at any time, even if the App Layering software has not yet been installed. It is important that you meet all requirements so that the OS layer works correctly in your environment.

Once you have met the requirements and have familiarized yourself with the guidelines for what to include in the OS layer, use the instructions for preparing the OS in your hypervisor environment. If you later expand support to another hypervisor, you can reuse this OS layer by installing the tools for the second hypervisor on the Platform layer that you create for that second environment.

Requirements and recommendations

When preparing an OS image, meet the following requirements and consider the related recommendations.

- **One OS layer for each Windows version you are managing (recommended):** Citrix® recommends that you prepare a single OS image for each Windows version that you are managing, along with a set of Platform and App layers for each.
- **Fresh OS image:** Start with a fresh image of a [supported Windows OS](#) from your hypervisor. This ensures that the image is optimized for your environment.
- **IP address from DHCP:** Make sure that the OS image is *not* in a domain. Ensure that the image gets its IP address from DHCP. Otherwise, you cannot install the App Layering OS Machine Tools. Domain join can be done in the platform layer.
- **App Layering OS Machine Tools:** Locate the OS Machine Tools in the App Layering installation package.
- **Optimization script for MS Office:** If you are going to run MS Office, you must use the optimization script included in the installation package.

XenServer®, MS Hyper-V, or VMware vSphere

In the rare case that you need to run Windows Mini Setup, you can edit the unattend.hta file we supply for your needs.

- **Answer file for unattended installation (optional):** The answer file is included in the App Layering download.

Notes:

Avoid using third-party scripts, because they can change services and features that the App Layering service uses, for example, Universal Plug and Play and the 8.3 file names setting.

What to include in the OS layer

Include the following software and settings in the OS layer:

- **Hypervisor tools:** You must include your hypervisor tools in the OS layer. You can upgrade the tools by adding a new version to the layer.

Notes:

- When you upgrade the hypervisor tools on the OS layer, test the existing platform layer to see if it needs updating. Depending on the platform and what else is installed on it, you may need to recreate the platform layer.
- If you are using the same OS layer with multiple hypervisors, it makes sense to install the hypervisor tools in purpose-built platform layers for those given hypervisors.

- **.NET Framework v4.0 or later:** Include .NET Framework v4.0 or later so that Windows updates are only required on the OS layer. For example, .NET 4.8 is required for Citrix Virtual Apps and Desktops (CVAD) 2303 to add a VDA.
- **.NET Framework 3.5 (when creating an MS Office layer):** For ease in updating, install all versions of the .NET Framework on the OS layer *before* creating the Office layer. If .NET Framework v3.5 is not present when you install Office, Office installs it for you, and it is not recommended to have .NET Framework versions or updates installed in app layers.
- **Disable Windows updates using *local GPO*:** Disable Windows updates on the OS layer, and do so using Local GPO rather than the Windows Update Service.
- **Windows Store app removal:** If you remove Windows Store apps, remove them from the OS layer, not on an App layer.
- **Windows activation:** Use KMS for Windows Activation. When creating your OS layer, run **SetKMSVersion.exe** to configure the startup scripts that activate the correct version of Windows.
- **User accounts and groups:** Any extra user accounts or groups must be created in the OS layer. Any domain group membership changes must be done through Group Policy.
- **Applications that create local users:** Include apps that create local users to ensure that changes to local groups and local users are captured, something that is not done on Platform and App layers.

What *not* to include in the OS layer

Do not include the following software on the OS layer.

- **Provisioning software:** Software associated with your Provisioning Service must be installed on your Platform layer, not on the OS layer.
- **Connection broker software:** Your connection broker software must also be installed on your Platform layer, not on the OS layer.
- **MS Office and other apps:** Do *not* include MS Office or other applications on the OS layer, except for the few apps that create local users. Generally, applications should be installed on App layers.
- **Domain join:** Do *not* join the OS layer to an Active Directory domain. Instead, join the domain in the Platform layer. This allows you to use the same OS in different domains.
- **Debug flag:** The Debug flag cannot be enabled in any BCD boot entry in your OS layer if you are using Secure Boot. Whether the flag is true or false does not matter; the flag itself cannot be present, as it is known to cause issues.

For detailed steps to prepare the OS, select your hypervisor:

- [XenServer](#), [Hyper-V](#), [vSphere](#)
- [Azure](#)
- [Nutanix](#)

Prepare your OS image for layering in XenServer®, Hyper-V, or vSphere

January 6, 2026

Before you start, ensure that you meet the [requirements](#). While preparing the image, you can [Expedite a Microsoft Ngen.exe operation, if necessary](#), if you think it is taking too long.

If using Windows 10, you can speed up desktop start times as long as you are *not* running Citrix Provisioning, machine creation, or VMware View. In this situation, you can [remove Windows 10 built-in applications](#). We recommend removing the apps on a *new version* of the OS layer, rather than in the OS image itself.

Note:

XenServer supports UEFI-based machines after a new XenServer connector was added.

Install the OS on a virtual machine

It is crucial to start with an OS freshly installed from ISO, preferably from your hypervisor.

In this procedure, be sure to follow steps and notes specific to the Windows version you are installing.

1. Log in to your hypervisor client.
2. Create a virtual machine with the correct CPU, RAM, hard drive, and network settings for your operating system type. Guidance:
 - **XenServer virtual machine:** Ensure that only one network is selected.
 - **vSphere virtual machine:**
 - **Network:** (Required) Select the **VMXNET 3 network adapter**.

Important:

You can have one, and only one, network device, and the E1000 NIC must *never have been used*. The default E1000 adapter (or even a ghost NIC leftover from an E1000

adapter) can cause customization timeout errors on the virtual machines.

- **Thin Provision:** Select **Thin Provision**.
- **All hypervisors:**
 - **Hard drive:** Ensure that the appliance can access the hard drive that you create.

3. Attach the ISO and install the operating system. **This machine must not be joined to the domain.** Domain join must be done in the Platform layer, and any domain group membership changes must be done through Group Policy.
4. Install the hypervisor tools for the platform where you plan to package layers. If you support multiple hypervisors, put the tools for the hypervisor you plan to use for publishing images in the Platform layer.
 - **For Hyper-V:** Use the Microsoft Windows Integration Services Setup Disk to install Hyper-V Integration Services.

If using a Server OS, install the Remote Desktop Session Host feature

When using a Windows Server, you need to install the **Remote Desktop Session Host** feature. When the **Remote Desktop Session Host** role is installed in the OS layer, it is updated as part of Windows. You can install the role on the platform layer with the VDA instead if you prefer.

If you install RDS in the OS layer, you need to use local GPOs to define the RDS license servers. Otherwise, over time, you will lose the ability to log in to packaging machines.

To install the Session Host feature:

1. In the **Server Manager**, select **Add roles and features**.
2. For the **Installation Type**, select **Role-based** or **Feature-based** installation.
3. For the **Server** role, select **Remote Desktop Services > Remote Desktop Session Host (Installed)**. This installs the C++ library and the RDS role.
4. Complete the process of adding the Server Roles.

Ensure the correct versions of .NET Framework are installed (Windows 10 and Windows Server 2016)

The .NET Framework is a software framework provided by Microsoft, and it is required for many third-party applications to run. Any installation of the .NET Framework must be included in the OS layer. This includes .NET 3.5 and .NET 4.0 or later.

Note:

Citrix Virtual Apps and Desktops (CVAD) 2303 requires .NET version 4.8 to add VDAs.

Be sure to install the .NET Framework and any updates on your OS layer.

Install Windows updates

Be sure to install all Windows updates.

1. Install all important updates.
2. Check for updates again after the virtual machine is rebooted. Some updates became available only after others are installed.
3. Install all required service packs:
 - If using Windows 2008 with Citrix Provisioning™, install Windows Server 2008 R2 Service Pack 1 (SP1).

Note:

If KB3125574 is installed, uninstall it before installing this service pack.

4. Clear **Windows Automatic Updates** and disable **Windows System Restore** using the local group policy editor, [gpedit.msc](#). The system handles restore points for you. Layer versions allow you to specify when updates occur.
5. **Windows 10:** Clear Hibernation by entering this command:

```
1 powercfg.exe /hibernate off
```

6. Enable the built-in administrator and select **Password never expires**.
7. If using Key Management Service (KMS) licensing, run a command window as Administrator, and enter these commands:

```
1 slmgr /skms <kmsserverhost>
2 slmgr /rearm
3 reboot
4 slmgr /ipk XXXX-YOUR-KMS-KEY-XXXX
5 slmgr /ato
```

8. If using a server OS, run the following commands in PowerShell:

```
1 Set-ExecutionPolicy Unrestricted
2 Enable-PSRemoting
```

Expedite a Microsoft Ngen.exe operation, if necessary

Once all software updates have been installed, you must allow [Ngen.exe](#) to essentially recompile [.NET](#) byte code into native images and construct the registry entries to manage them.

[Ngen.exe](#) is the Microsoft Native Image Generator, which is part of the [.NET](#) system. Windows determines when to run [Ngen.exe](#) based on what software is being installed and what Windows detects in the configuration.

Important:

When [Ngen.exe](#) is running, you must let it complete. An interrupted [Ngen.exe](#) operation can leave you with non-functioning [.NET](#) assemblies or other problems in the [.NET](#) system.

Normally, [Ngen.exe](#) is a background operation that pauses when there is a foreground activity. If you want to expedite an [Ngen.exe](#) operation, you can bring the task into the foreground to complete it as quickly as possible.

To bring the task into the foreground:

1. Open a command prompt as Administrator.
2. Go to the [Microsoft .NET\Framework](#) directory for the version currently in use:

```
1 cd C:\Windows\Microsoft.NET\FrameworkNN\vX.X.XXXXX
```

3. Enter the following [Ngen.exe](#) command to run all queued items. This command processes queued component installs before building assemblies.

```
1 ngen eqi 3
```

Note:

This variation of the [ngen](#) command has been tested and is the variation that works in this situation in App Layering.

The [Ngen.exe](#) task moves to the foreground in the command prompt and lists the assemblies being compiled. It is OK if you see compilation messages.

You can use the **Task Manager** to see if an instance of [MSCORSVW.EXE](#) is running. If it is, allow it to complete, or run [ngen eqi 3](#).

Warning:

Do not reboot to stop the task. Allow the task to complete!

4. Ensure that all [Ngen.exe](#) processes have run to completion.

Run the App Layering OS Machine Tools on the image

To prepare the OS image to run in a layer, you run the OS Machine Tools file on the image. This executable runs a GPO setup script (gposetup.cmd), and a script to set the Key Management Service (KMS) version. The script is called [SetKMSVersion.hta](#).

1. Download the following zip file onto the OS image:

App_Layering_Citrix_App_Layering_OS_Machine_Tools_20.x.zip

2. Extract the files to:

```
1 c:\windows\setup\scripts
```

Note:

The file must be extracted to the directory. Do not change the directory.

If using KMS, configure license activation

Once the Key Management Service (KMS) scripts are extracted, the [SetKMSVersion](#) utility asks you to choose whether to use KMS licensing.

Note:

Publishing images into environments where both KMS and Active Directory-based activation (ADBA) are being used at the same time causes problems with activation.

1. In the dialog box that appears, select whether to use Key Management Service (KMS) licensing.



To configure scripts for KMS, do the following.

1. Go to:

`c:\windows\setup\scripts`

2. Run **SetKMSVersion.hta** as Administrator to create a script in the `c:\windows\setup\scripts\kmsdir` folder.

When the operating system starts, the appropriate KMS activation script is run.

Install the App Layering services

1. In the `c:\windows\setup\scripts` folder, run the **setup_x86.exe (32-bit)** or **setup_x64.exe (64-bit)**.

You are ready to [import the image](#) into a new OS layer.

Note:

Ensure that the image preparation tools installer is run once before the OS has been imported. Don't run the image preparation tools installer after the OS is imported as this might cause unknown issues.

Prepare your OS image for layering on Google Cloud

January 5, 2026

This topic explains how to prepare a clean OS image for import into a new OS layer. Before you start, make sure that you meet the [requirements](#). While preparing the image, you can [Expedite Microsoft Ngen.exe operations](#), if you think it is taking too long.

If using Windows 10 and *not* running Citrix Provisioning, Citrix machine creation, or View, you can speed up desktop start times by [removing Windows 10 built-in applications](#). However, we recommend removing the apps on a *new version* of the OS layer, *not* in the OS image itself.

Install the OS on a virtual machine

1. Get familiar with the guidelines for [preparing an OS layer](#), including requirements and recommendations. Be sure to read the sections on what to include, and what not to include in an OS layer layer.
2. Navigate to the [Google Cloud portal](#).
3. Select **Marketplace** in the left column, and deploy a new virtual machine.

Note:

When configuring the new instance network, make sure the VM is on a network that is accessible to the appliance.

4. If you are using a Windows Server operating system, scroll to **Operating systems**, and select a **Windows Server 2019** or **Windows Server 2016** operating system.
5. If you are bringing Windows 10 from another platform (it is not available in the Marketplace), follow the steps in the [Bringing your own licenses tutorial](#).
6. Configure the new instance:
 - When selecting a network for the new instance, make sure that the VM is on a network that is accessible to the appliance.
 - When selecting storage, any type of storage is fine.

Run the App Layering OS Machine Tools on the image

1. On the new machine, open a web browser, navigate to the Download Center and download the **OS Machine Tools**.
2. Download the following zip file onto the OS image:

```
1 Citrix_App_Layering_OS_Machine_Tools_20.x.x.exe
```

3. Run the file, and it copies files to:

```
1 c:\windows\setup\scripts
```

Note:

The file must be extracted to the above directory. Do not change the directory.

If using Key Management Service, configure license activation

Once the scripts are extracted, the [SetKMSVersion](#) utility asks you to choose whether to use Key Management Service (KMS) licensing.

Note:

Publishing images into environments where both KMS and Active Directory-based activation (ADBA) are being used at the same time causes problems with activation.

1. In the following dialog box, select whether to use Key Management Service (KMS) licensing.



To configure scripts for KMS, do the following.

1. Navigate to:

```
1 c:\windows\setup\scripts
```

2. Run **SetKMSVersion.exe** as Administrator to create a script file in the `c:\windows\setup\scripts\kmsdir` folder.

When the operating system starts, the appropriate KMS activation script is run.

Install the App Layering services

1. On the new machine, navigate to `C:\Windows\Setup\scripts` and run `setup_x64.exe` to install the App Layering drivers on the OS machine.
2. The installation prompts you for the location of the `Unattend.xml` file (the default location is '`C:\windows\panther`).
3. Ensure that this machine is not joined to a domain.
4. Perform pending reboots on the OS machine so that you can import this image into a layer.
5. Make sure that the new OS machine is in one of the following states before proceeding.
 - Running
 - Stopped
 - Stopped (deallocated)

Expedite a Microsoft Ngen.exe operation, if necessary

Once all software updates have been installed, you must allow [Ngen.exe](#) to essentially recompile .NET byte code into native images and construct the registry entries to manage them.

[Ngen.exe](#) is the Microsoft Native Image Generator, which is part of the .NET system. Windows determines when to run [Ngen.exe](#) based on what software is being installed and what Windows detects in the configuration.

Important:

When [Ngen.exe](#) is running, you must let it complete. An interrupted [Ngen.exe](#) operation can leave you with non-functioning .NET assemblies or other problems in the .NET system.

Normally, [Ngen.exe](#) is a background operation that pauses when there is foreground activity. If you want to expedite an [Ngen.exe](#) operation, you can bring the task into the foreground to complete it as quickly as possible.

To bring the task into the foreground:

1. Open a command prompt as Administrator.
2. Go to the [Microsoft.NET\Framework](#) directory for the version currently in use:

```
1 cd C:\Windows\Microsoft.NET\FrameworkNN\vX.X.XXXX
```

3. Enter the following [Ngen.exe](#) command to run all queued items. This command processes queued component installs before building assemblies.

```
1 ngen eqi 3
```

The [Ngen.exe](#) task moves to the foreground in the command prompt, and lists the assemblies being compiled. It is OK if you see compilation messages.

You can use the **Task Manager** to see if an instance of [MSCORSVW.EXE](#) is running. If it is, allow it to complete, or run [ngen update eqi 3](#).

Caution:

Do not reboot to stop the task. Allow the task to complete!

4. Ensure that all [Ngen.exe](#) processes have run to completion.

Prepare your OS image for layering in Azure

January 5, 2026

This topic explains how to prepare a clean OS image for import into a new OS layer. Before you start, make sure that you meet the [requirements](#). While preparing the image, you can [Expedite Microsoft Ngen.exe operations](#), if you think it is taking too long.

If using Windows 10 and *not* running Citrix Provisioning, machine creation, or View, you can speed up desktop start times by [removing Windows 10 built-in applications](#). However, we recommend removing the apps on a *new version* of the OS layer, *not* in the OS image itself.

Note:

Do *not* use an unattend file with a Machine Creation Services™ (MCS) Azure connector. The App Layering software removes the unattend file if it is present, because it is not necessary or recommended for an MCS Azure connector.

Install the OS on a virtual machine

1. In the [Microsoft Azure portal](#), create a new virtual machine from the Windows Server Remote Desktop image by selecting:

New > Compute > Virtual Machine

2. Complete the Create virtual machine wizard:

Basics:

- **Name:** The name you specify for the new machine must comply with Azure naming conventions.
- **Username and password:** The user name and password of the new server machine you specify are used for any packaging machines that are created containing this OS layer.
- **Resource group location:** Be sure that the value for the Resource group location matches the Storage account location that you configured in the connector configuration.

3. Select required network settings.
4. Review the summary and create the virtual machine.
5. Log into the new virtual machine, and reboot the machine.
6. Install all important updates. Be sure to reboot the system and check for more updates. Some updates become available only after others are installed.
7. Run Windows Ngen.exe.
8. Remove or rename the Unattend file in [C:\Windows\OEM](#).
9. Clear Windows Automatic Updates by selecting:
Control Panel > System and Security > Windows Update > Change Settings
10. Ensure that this machine is not joined to a domain.

11. Enable the built-in administrator and check **Password never expires**.
12. If this is a server OS, run the following commands in PowerShell:

```
1 Set-ExecutionPolicy Unrestricted  
2 Enable-PSRemoting
```

Run the App Layering OS Machine Tools on the image

1. On the new machine, open a web browser, navigate to the Download Center and download the OS Machine Tools.
2. Download the following zip file onto the OS image:

```
1 Citrix_App_Layering_OS_Machine_Tools_20.x.x.exe
```

3. Run the file, and it copies files to:

```
c:\windows\setup\scripts
```

Note:

The file must be extracted to the above directory. Do not change the directory.

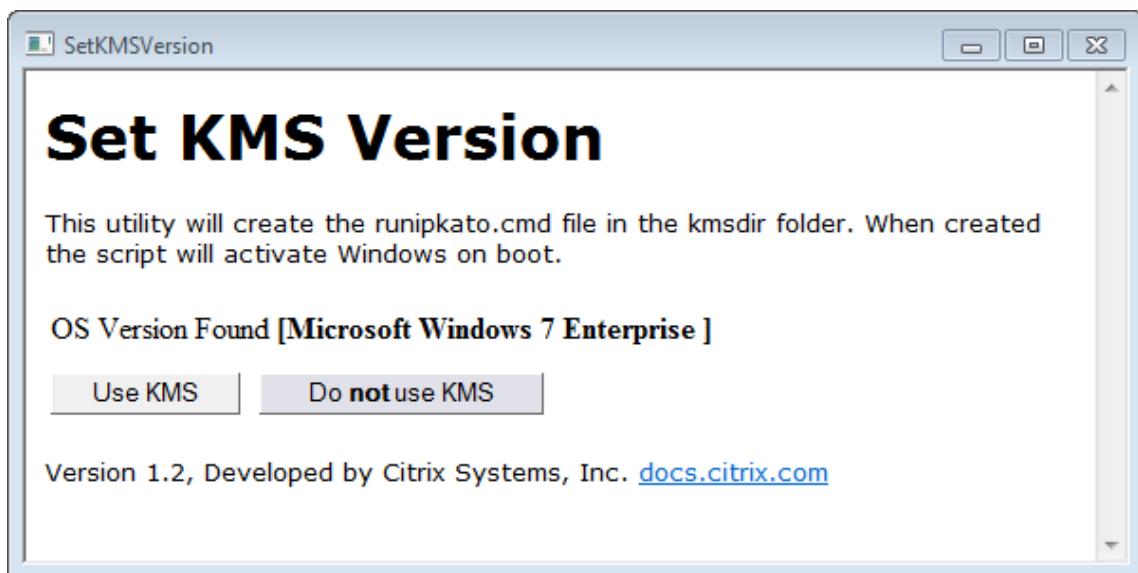
If using Key Management Service, configure license activation

Once the scripts are extracted, the `SetKMSVersion` utility asks you to choose whether to use Key Management Service (KMS) licensing.

Note:

Publishing images into environments where both KMS and Active Directory-based activation (ADBA) are being used at the same time causes problems with activation.

1. In the following dialog box, select whether to use Key Management Service (KMS) licensing.



To configure scripts for KMS, do the following.

1. Navigate to:

`c:\windows\setup\scripts`

2. Run **SetKMSVersion.exe** as Administrator to create a script file in the `c:\windows\setup\scripts\kmsdir` folder.

When the operating system starts, the appropriate KMS activation script is run.

Install the App Layering services

1. On the new machine, navigate to `C:\Windows\Setup\scripts` and run **setup_x64.exe** to install the App Layering drivers on the OS machine.
2. The installation prompts you for the location of the `Unattend.xml` file (the default location is '`C:\windows\panther`).
3. Ensure that this machine is not joined to a domain.
4. Perform pending reboots on the OS machine so that you can import this image into a layer.
5. Make sure that the new OS machine is in one of the following states before proceeding.
 - Running
 - Stopped
 - Stopped (deallocated)

Expedite a Microsoft Ngen.exe operation, if necessary

Once all software updates have been installed, you must allow **Ngen.exe** to essentially recompile `.NET` byte code into native images and construct the registry entries to manage them.

`Ngen.exe` is the Microsoft Native Image Generator, which is part of the `.NET` system. Windows determines when to run `Ngen.exe` based on what software is being installed and what Windows detects in the configuration.

Important:

When `Ngen.exe` is running, you must let it complete. An interrupted `Ngen.exe` operation can leave you with non-functioning `.NET` assemblies or other problems in the `.NET` system.

Normally, `Ngen.exe` is a background operation that pauses when there is foreground activity. If you want to expedite an `Ngen.exe` operation, you can bring the task into the foreground to complete it as quickly as possible.

To bring the task into the foreground:

1. Open a command prompt as Administrator.
2. Go to the `Microsoft.NET\Framework` directory for the version currently in use:

```
cd C:\Windows\Microsoft.NET\FrameworkNN\vX.X.XXXX
```

3. Enter the following `Ngen.exe` command to run all queued items. This command processes queued component installs before building assemblies.

```
ngen eqi 3
```

The `Ngen.exe` task moves to the foreground in the command prompt, and lists the assemblies being compiled. It is OK if you see compilation messages.

You can use the **Task Manager** to see if an instance of `MSCORSVW.EXE` is running. If it is, allow it to complete, or run `ngen update eqi 3`.

Caution:

Do not reboot to stop the task. Allow the task to complete!

4. Ensure that all `Ngen.exe` processes have run to completion.

Prepare your OS image for layering in Nutanix

January 5, 2026

This topic explains how to prepare a clean OS image for import into a new OS layer. Before you start, make sure that you meet the [requirements](#). While preparing the image, you can [Expedite a Microsoft Ngen.exe operation, if necessary](#), if you think it is taking too long.

If using Windows 10 and *not* running PVS, machine creation, or View, you can speed up desktop start times by [removing Windows 10 built-in applications](#). However, we recommend removing the apps on a *new version* of the OS layer, *not* in the OS image itself.

Note:

Do *not* use an unattend file in Nutanix. The App Layering software removes the unattend file if it is present, because it is not necessary or recommended in Nutanix.

Install the OS on a virtual machine

As part of this procedure, you can set up Key Management Service (KMS) activation.

Note:

Publishing images into environments where both KMS and Active Directory-based activation(ADBA) are being used at the same time causes problems with activation.

1. Log into the Prism Console.
2. Select **Task > VM**, and switch to **Table View** to see the existing virtual machines.
3. Click **+Create VM** in the upper right corner, and enter the specifics about the new virtual machine:
 - a) Enter a **Name** and add a **Description**.
 - b) Select the number of **VCPU**s.
 - c) Set the **Cores per CPU**.
 - d) Set **Memory**.
 - e) Select **Disks**, and create a virtual machine with three disks. The first CD-ROM is the ISO for the OS. The second CD-ROM is for the Nutanix VIRTIO drivers that allow the Nutanix virtual machine to access the disk where you install the OS. One CD-ROM is assigned in the beginning.
 - i. Edit the values for the assigned **CD-ROM**:
 - ii. For Operation, select Clone from **ADSF** file.
 - iii. For Bus Type, select **IDE**.
 - iv. Enter the path to your Windows ISO. The path is the combination of the Storage Container and the ISO Name. For example:
`/ISOSTore/en_windows_10_enterprise_version_1511_x64_dvd_7224901
.iso`
 - v. Click **Update**.
 - f) Add another disk by clicking the **+Add New Disk** button:
 - i. Set the Type to **CDROM**.

- ii. Set the Operation to **Clone from ADSF file**.
- iii. Set the Bus Type to **IDE**
- iv. Enter the path to the Windows VIRTIO Drivers. For example:
`/ISOStore/virtio-win-0.1.102.iso`
- v. Click **Add**.
- g) Click the **+Add New Disk** button.
 - i. Set the **Type** to **Disk**.
 - ii. Set the **Operation** to **Allocate on Container**.
 - iii. Set the **Bus Type** to **SCSI**.
 - iv. Select the **Container** you want to use.
 - v. Enter the **Size**.
 - vi. Click **Add**.
- h) Click **+Add new Nic**, and enter the **VLAN Name**.
- i) Click **Save**.

4. Power on the **virtual machine**.
 - a) Select **Tasks > VM**.
 - b) Switch to the **Table View** to see existing virtual machines.
 - c) Select the virtual machine in the **Table**, and click **Power On**.
5. Launch the console by selecting the VM and clicking Launch Console. When the VM boots it begins to install the Windows OS from the ISO disk. When the VM boots, it begins to install the Windows OS from the ISO disk.
 - a) When asked, “Where do you want to install Windows?” notice that even though you added a disk in the VM creation wizard, there is no disk.
 - b) Select the **Load Driver** option, and select **Browse**.
 - c) Select the CD with the **virtio-win-0.1.1** drivers.
 - d) Select the **vioscsi** folder, and choose the folder for your Windows OS.
6. After the OS is installed manually install the VirtIO drivers:
 - a) Launch **Device Manager**.
 - b) Select **Other Devices**, right-click **Ethernet Controller** and choose **Update Driver Software**.
 - c) Browse **My Computer**, and choose the **VirtIO CD**. The Ethernet drivers are stored in the **NetKVM** folder.
7. **Server OS:** If you need a session host feature:
 - a) Select **Add roles and features**.
 - b) For the Installation Type select **Feature-based installation**.

- c) For the Server Role, select **Remote Desktop Services > Remote Desktop Session Host**.
- d) Complete the process of adding server roles.
8. Install all important updates. Restart the system and check for more updates. Some updates only become available after others are installed.
9. Install all required service packs.
10. Disable **Windows System Restore** and **Windows Automatic Updates**.
11. Enable built-in administrator and check **Password never expires**.
12. If using Key Management Service (KMS) licensing, run a command window as Administrator, and enter these commands:

```
1 slmgr /skms <kmsserverhost>
2 slmgr /rearm
3 reboot
4 slmgr /ipk XXXX-YOUR-KMS-KEY-XXXX
5 slmgr /ato
```

13. **Server OS:** Add Domain Users to Remote setting for server OS.
14. Check for dead (ghost) NICs and delete if any exist. Enter the commands:

```
1 set devmgr_show_nonpresent_devices=1
2 devmgmt.msc
```

15. Uninstall any dead (ghost) NICs.
16. If this is a server OS, run the following commands in PowerShell:

```
1 Set-ExecutionPolicy Unrestricted
2 Enable-PSRemoting
```

Run the OS Machine Tools on the OS image

To prepare the OS image to run in a layer, you run the OS Machine Tools file on the image. This executable runs a GPO setup script (gposetup.cmd), and a Set KMS Version script (SetKMSVersion.hta).

1. Download the following executable file onto the OS image:

[Citrix_App_Layering_OS_Machine_Tools_20.x.x.exe](#)

2. Run the executable. Files are saved to:

`c:\windows\setup\scripts`

Note:

The file must be extracted to the `c:\windows\setup\scripts` directory. Do not change the directory.

If using Key Management Service (KMS), configure license activation

Once the scripts are extracted, the SetKMSVersion utility asks you to choose whether to use KMS licensing.

Note:

Publishing images into environments where both KMS and Active Directory-based activation (ADBA) are being used at the same time causes problems with activation.

1. In the following dialog box, select whether to use Key Management Service (KMS) licensing.



To configure scripts for KMS, do the following.

1. Navigate to:

`c:\windows\setup\scripts`

2. Run **SetKMSVersion.exe** as Administrator. This creates a script file in the `c:\windows\setup\scripts\kmsdir` folder.

When the operating system starts, the appropriate KMS activation script is run.

Make sure the correct versions of .NET Framework are installed (Windows 10 and Windows Server 2016)

The .NET Framework is a software framework provided by Microsoft, and it is required for many third party applications to run. Any installations of the .NET Framework must be included in the OS layer. This includes .NET 3.5 and .NET 4.0 or later.

Note:

.NET 4.8 is required by Citrix Virtual Apps and Desktops (CVAD) 2303 to add VDAs.

Be sure to install the .NET Framework and any updates on your OS layer.

Install the App Layering services

1. In the `c:\windows\setup\scripts` folder, run the **setup_x86.exe (32-bit)** or **setup_x64.exe (64-bit)**.
2. The installation prompts for the location of the `unattend` file. Do NOT use the `unattend` file in Nutanix.

Run the Optimization script, if using MS Office

The Optimization script included in the App Layering installation package is required to layer Microsoft Office. This script allows you to save memory and CPU by disabling services you don't need, enabling services you do need, and removing installation-specific drivers and settings.

You can run the Optimization script on the OS layer, and if needed, supersede it with a new version of the script in an App layer included in your image template. Since App layers are applied to the image after the OS layer, the script in the App layer overrides the original version in the OS layer.

1. In the `c:\windows\setup\scripts` folder, run the **optimizations.cmd** file to create a file to run when the image is created.
2. Follow the instructions to run `optimizations.cmd` on the OS image.

Expedite a Microsoft Ngen.exe operation, if necessary

Once all software updates have been installed, you must allow `Ngen.exe` to essentially recompile .NET byte code into native images and construct the registry entries to manage them.

`Ngen.exe` is the Microsoft Native Image Generator, which is part of the .NET system. Windows determines when to run `Ngen.exe` based on what software is being installed and what Windows detects in the configuration.

Important:

When `Ngen.exe` is running, you must let it complete. An interrupted `Ngen.exe` operation can leave you with non-functioning .NET assemblies or other problems in the .NET system.

Normally, `Ngen.exe` is a background operation that pauses when there is foreground activity. If you want to expedite an `Ngen.exe` operation, you can bring the task into the foreground to complete it as quickly as possible.

To bring the task into the foreground:

1. Open a command prompt as Administrator.
2. Go to the `Microsoft.NET\Framework` directory for the version currently in use:

```
cd C:\Windows\Microsoft.NET\FrameworkNN\vX.X.XXXX
```

3. Enter the following `Ngen.exe` command to run all queued items. This command processes queued component installs before building assemblies.

```
ngen eqi 3
```

The `Ngen.exe` task moves to the foreground in the command prompt, and lists the assemblies being compiled. It is OK if you see compilation messages.

You can use the **Task Manager** to see if an instance of `MSCORSVW.EXE` is running. If it is, allow it to complete, or run `ngen update eqi 3`.

Caution:

Do not reboot to stop the task. Allow the task to complete!

4. Ensure that all `Ngen.exe` processes have run to completion.

Create the OS layer

January 5, 2026

An OS layer includes the software and settings for the operating system that you deploy in layered images. The OS layer is necessary for creating:

- Platform layers
- App layers
- Layered images

Requirements

Before creating the OS layer, be sure to:

- [Prepare the OS image for importing into the layer.](#)
- [Install the appliance.](#)
- [Configure App Layering.](#)

Considerations for your OS layer

- To deploy Windows patches and updates, you can simply add a version to the layer. You can easily revert to the previous version of the layer, if necessary.
- You can select any version of the layer to use in an image template, and therefore in the published images.
- You can update the OS using Windows Update, Windows Server Update Services (WSUS), or offline standalone update packages. Do *not* use tools like SCCM.
- Platform and app layers are tied to the specific OS layer that you use to create them, though not to a specific version of the layer. When you add versions to the OS layer, the dependent app and platform layers continue to work.
- Windows updates must be applied to the OS layer before you update any other layers.

Import the OS using the ImportOSLayer.ps1 script

This procedure explains how to import the OS for your new OS layer using the [ImportOsLayer.ps1](#) script.

If you have downloaded and expanded the App Layering OS Machine Tools onto your OS image, the [ImportOsLayer.ps1](#) has been copied to c:\windows\setup\scripts.

Run the script

To import the OS:

1. Run the [ImportOsLayer.ps1](#) PoSH script as administrator:

```
1 C:\Windows\Setup\scripts\ImportOsLayer.ps1 -ElmAddress <Ip Address
     > [-IgnoreCertErrors]
2 C:\Windows\Setup\scripts\ImportOsLayer.ps1 -ElmAddress <FQDN> [-
     IgnoreCertErrors]
```

where

- `ElmAddress` is the IP Address or FQDN of the App Layering appliance. It specifies where the new OS layer is created.
- `IgnoreCertErrors` ignores certification errors when the script communicates with the App Layering appliance.

2. The `ImportOsLayer.ps1` script prompts you for the credentials to connect to the App Layering appliance (referred to as the ELM in the script). The script uses your credentials to create a session on the appliance.
3. The script then prompts you for details about the new OS layer:
 - LayerName (required)
 - VersionName (required)
 - LayerSizeGib (required, but defaults to 60 GB)
 - LayerDescription (optional)
 - VersionDescription (optional)
 - Comment (optional)

Once you've entered the required information, the script reboots the system into the compositing engine, imports the OS, and builds the layer. Monitor the progress of the job in the management console.

When the compositing engine is finished (success or failure), it reboots back into the Windows OS image.

Create Platform layer

January 5, 2026

A platform layer includes the platform software and settings required for your layers and layered images to run flawlessly in your environment.

You can create platform layers for two purposes:

- **For creating and packaging layers:** When you've imported the OS from a different hypervisor than the one where you create your layers, use this type of platform layer to create app layers.
- **For publishing layered images:** Use this type of Platform layer in your image template so that the published layered images run flawlessly in your environment.

Platform layers for packaging layers or publishing layered images

Use the following table to determine whether you need a platform layer. This table also shows what software to install on the platform layer, if you need one.

	Packaging layers	Publishing layered images
Platform layer required?	Required if the OS image originated on a different hypervisor. When an app requires the agent or SSO software, you can create a platform layer specifically for creating and updating that layer.	Required when publishing to a provisioning server and using a connection broker.
What to install	Hypervisor tools, when the OS originated on a different hypervisor. The SSO or agent software, if necessary to create an app layer.	Provisioning and connection broker software and settings. If publishing to a different hypervisor than the one where the OS originated, include the hypervisor tools.
Values to select	Select your hypervisor.	Select your hypervisor, provisioning software, and connection broker.
What you need	Installer for hypervisor	Installers for provisioning software and connection broker.

Other software and settings to include in the platform layer

Besides the platform software specified above, you must include the following settings and software on the platform layer:

- Domain join
- NVIDIA Drivers, if applicable
- Citrix Receiver™, for the single sign-on component
- Citrix Workspace™ Environment Management(WEM) agent

Note:

The RSA key generated by Citrix WEM causes issues when using WEM on the deployed image. If the RSA key is present when finalizing the layer, a message appears stating you must delete the RSA key file, which begins with the following path: `C:\ProgramData\Microsoft\Crypto\RSA\S-1-5-18\fb8cc9e38d3e60ab60c17cdfd6dd6d99_`.

- Citrix® uberAgent®
- Any software that impacts the logon stack, for example, Imprivata
- Citrix Provisioning™ on Hyper-V: Requires a Legacy Network Adapter to PXE boot
- Microsoft System Center Configuration Manager (SCCM) software, if you are using it

Process for creating a platform layer

The steps for creating a platform layer are to:

1. Create a platform layer in the management console.
2. Connect to and log in to the packaging machine.
3. Install your provisioning and connection broker software.
4. Is the appliance running on a different hypervisor than the one where you create layers and publish images? If yes, we recommend installing the hypervisor tools too.
5. Verify the layer and shut down the packaging machine.
 - If the connector configuration selected is set to use **Offload Compositing**, the layer is automatically finalized.
 - If the connector configuration is *not* set to **Offload Compositing**, [finalize the layer manually](#), as described in the detailed steps in this article.

When to update a platform layer

The platform layer is the highest priority layer. It is critical for deploying images, especially for network devices. Whenever you update the infrastructure software, add a version to the platform layer.

When you update the OS layer, the image sometimes has startup issues. To fix the problem, add a version to the platform layer using the new OS layer. Once the packaging machine starts, shut down the machine for finalization. The platform layer gathers the critical components from the new OS layer version and updates them in the platform so that they match the OS version.

Requirements

When creating a platform layer, the software installers must be available in a location the packaging machine can access. For example, your provisioning server and connection broker software must be accessible. If the appliance is running on a different hypervisor, also include the hypervisor tools.

For detailed requirements, select the environment where you create layers or publish images:

- [Citrix Provisioning](#)
- [XenServer](#)
- [MS Azure](#)
- [MS Hyper-V](#)
- [Nutanix AHV](#)
- [VMware vSphere](#)
- [Google Cloud](#)

A word on optimizations

The platform layer is the highest priority layer. You might think it would be the best place to include optimizations. However, on Windows 10, optimizations that remove Windows Apps only work on the OS layer. The Windows Apps are integrated with the Windows **Store**, which can only be modified in the OS layer.

Citrix offers an excellent optimization utility called the [Citrix Optimizer](#). We recommend using this utility, rather than the optimizer shipped with App Layering because the Citrix Optimizer can usually reverse the optimizations if needed.

To speed up user logins. Log in using a network user account, and reboot the desktop. Then, log in as an administrator, and delete the profile created. When the first network user logs on, some system files are updated, then login performance usually improves.

Start a new platform layer

To create a platform layer, follow these steps:

- Prepare the layer using **Create Platform Layer**.
- Deploy a packaging machine in your environment.
- Install the tools and configure the settings for your environment.
- Finalize the layer.

Follow these steps, starting in the Action bar:

1. Select **Layers > Platform Layers**. Then select **Create Platform Layer**.

2. In the Layer Details tab, enter a **Layer Name** and **Version**, both required values. Optionally, you can also enter other values.
3. In the Version Details tab:
 - a) (Required) Enter a New Version name. For example, enter the software version or other identifying information.
 - b) If you are adding a version to an existing layer, the **Base Version** field lets you choose which version to use as the starting point. The default choice is the latest version.
4. In the OS layer tab, select the OS layer you want to associate with this Platform layer.
5. In the Connector tab, choose a **Connector Configuration** for the platform where you are creating this layer.
6. In the **Platform Types** tab, select **This platform will be used for publishing layered images**, or **This platform will be used for packaging**. Then select the hypervisor, provisioning software, and connection broker where you are publishing the layered image.
Note: If you are *not* using provisioning or a connection broker, select **None** for each of those options.
7. In the Packaging Disk tab, enter a **file name** for the packaging disk. This disk is used for the packaging machine (the virtual machine) where you want to install the tools.
8. In the Icon Assignment tab, select an icon to assign to the layer. This icon represents the layer in the Layers module.
 - To use an existing image, select an image in the image box.
 - To import a new image, click **Browse** and select an image in PNG or JPG format.
 - If the layer uses one of the supplied icons and a connector with **Offload Compositing** selected, the packaging machine assigns an icon based on the layer's contents.
9. In the Confirm and Complete tab, review the details of the App layer, enter a comment if necessary, and click **Create Layer**. Any comments you enter appear in the **Information** view of the **Audit History**.
10. Select the **Tasks** page, and click the **Packaging Disk** task. Click the info icon to show the full task description.

Once the packaging disk has been created, the Tasks bar displays the location of the packaging disk in your environment.

Next, you can deploy the packaging machine for your layer.

Deploy a Packaging Machine

The App Layering system creates a packaging machine in the location defined in the connector configuration. The packaging machine is a temporary VM where you install the software for the layer. Once you finalize the layer, the packing machine is removed.

Citrix Provisioning for Hyper-V: Configuring two network cards

When using dual network cards and running Citrix Provisioning for Hyper-V, you must configure the cards as follows on *every new version of the Platform layer*.

Once your provisioning software is installed and the required reboots have been completed:

1. Open an administrative command prompt on the packaging machine.
2. Run the command: `ipconfig /all`
3. Match the IP address of the streaming NIC (Legacy Network Adapter in Hyper-V) with the correct adapter name.
4. Renew the DHCP lease on the streaming NIC.
5. Again in an administrative command prompt run `ipconfig /release *adapter-name*` followed by `ipconfig /renew *adapter-name*`. This command forces the App Layering drivers to select this adapter as the “primary NIC”.
6. Run **Shutdown for Finalize** and finalize the layer as you normally would.

Important:

If you select Shutdown for Finalize, but then need to turn the machine back on for any reason, you must rerun the **release** and **renew** commands.

Install the platform tools on the packaging machine

Next, install the software for the platform where you publish layered images. Platform tools include provisioning and connection broker software that layered images require in the target environment. Keep in mind, the state of the software when you finalize the layer is the state the image uses.

1. Log in remotely to the packaging machine. Be sure to log in using the user account you used to create the OS.
2. Install the tools that your layered images are configured to run. For example, include your provisioning, connection broker, and hypervisor tools. Don’t forget your drivers, boot-level applications, and any required files.
3. If the installation requires a system restart, restart it manually. The packaging machine does not restart automatically.

4. Make sure that the packaging machine is in the state you want it to be in when the image is booted:
 - If the tools you install require any post-installation setup or registration, complete those steps now.
 - Remove any settings, configurations, files, mapped drives, or applications that you do not want to include on the packaging machine.
5. (Optional) To customize the image deployed from the ELM before deployment to MCS, follow these steps:
 - a) Upgrade your master tools in OS revision 2308 and beyond.
 - b) Then, create the file:
`c:\windows\setup\scripts\kmsdir\Admin_Controlled_Shutdown.txt`

Note:
File contents are not important.
 - c) When the image is deployed from the ELM, the booted image remains running so that you can do your customizations. A reboot does not affect the machine's state.
 - d) After you've completed your customization, run the command:
`c:\windows\setup\scripts\kmsdir\CompleteDeployment.cmd`.
At this point, the machine shuts down and is finished with the deployment task. With this, you can deploy the machine to MCS.

Verify the layer and shut down the packaging machine

Once the tools are installed on the packaging machine, you can verify that the layer is ready to finalize. Any required post-installation processing needs to be completed. For example, a reboot or a Microsoft `ngen` process might need to complete.

To verify that outstanding processes are complete, run the *Shutdown For Finalize* tool. Look for the *Shutdown For Finalize* icon on the packaging machine's desktop.

Shut down the packaging machine so you can finalize the layer

1. If you are not logged into the packaging machine, remote login using the account set up during OS layer creation.
2. Double-click the *Shutdown For Finalize* icon. A command-line window displays messages detailing the layer verification process.
3. If there is an outstanding operation, you are prompted to complete the process. For example, if a Microsoft `ngen` operation must be completed, you can expedite the `ngen` operation, as detailed in the section, [Layer integrity messages during the finalization process](#).

4. Once pending operations are done, double-click the *Shutdown For Finalize* icon again.

The Layer is now ready to finalize.

- If the connector configuration selected is set to **Offload Compositing**, the layer is automatically finalized.
- If you are not using **Offload Compositing**, [finalize the layer manually](#).

Layer integrity messages during the finalization process

The following layer integrity messages tell you what queued operations must be completed before the layer is ready to finalize:

- A [RunOnce script is outstanding](#) – check and reboot the packaging machine.
- A [post-installation reboot is pending](#) – check and reboot the packaging machine.
- A [Microsoft ngen operation is in progress in the background](#). – An [MSI install operation is in progress](#) – check the packaging machine.
- A [reboot is pending to update drivers on the boot disk](#) – check and reboot the packaging machine.
- A [Microsoft ngen operation is needed](#).
- [Software Center Client is configured to run, but the SMSCFG.INI is still present](#). To learn more about deploying SCCM in a virtual environment, see the Microsoft TechNet article, [Implementing SCCM in a XenDesktop VDI environment](#).

For details about what the layer integrity messages mean and how to debug them, see [Debugging Layer Integrity Problems in Citrix App Layering 4.x and later](#).

You cannot bypass layer integrity messages by shutting down the machine. The App Layering software returns you to the packaging machine until the processes have been completed.

If a Microsoft [ngen](#) operation is in progress, you can try to expedite it, as described in the next section.

Expedite Microsoft Ngen.exe operations, if necessary

Once all software updates have been installed, you must allow [Ngen.exe](#) to essentially recompile [.NET](#) byte code into native images and construct the registry entries to manage them.

The [Ngen.exe](#) executable is the Microsoft Native Image Generator, which is part of the [.NET](#) system. Windows determines when to run [Ngen.exe](#) based on what software is being installed and what Windows detects in the configuration.

Important:

When `Ngen.exe` is running, you must let it complete. An interrupted `Ngen.exe` operation can leave you with non-functioning .NET assemblies or other problems in the .NET system.

Normally, `Ngen.exe` is a background operation that pauses when there is a foreground activity. To expedite an `Ngen.exe` operation, bring the task into the foreground to complete it.

To bring the task into the foreground:

1. Open a command prompt as Administrator.
2. Go to the `Microsoft.NET\Framework` directory for the version currently in use:

```
cd C:\Windows\Microsoft.NET\FrameworkNN\vX.X.XXXX
```

3. Enter the following `Ngen.exe` command to run all queued items. This command processes queued component installs before building assemblies.

```
ngen eqi 3
```

The `Ngen.exe` task moves to the foreground in the command prompt and lists the assemblies being compiled. It is OK if you see compilation messages.

You can use the **Task Manager** to see if an instance of `MSCORSVW.EXE` is running. If it is, allow it to complete, or run `ngen eqi 3`.

Caution:

Do not reboot to stop the task. Allow the task to complete!

4. Ensure that all `Ngen.exe` processes have run to completion.
5. When complete, shut down the virtual machine using the **Shutdown For Finalize** shortcut available on your desktop.

Finalize the layer manually

Now that the layer has been verified and shut down, it is ready to finalize.

NOTE:

If you are using a connector with **Offload Compositing** enabled, this finalization process is automated and you do not have to do these manual steps.

1. Return to the management console.
2. Select **Layers > Platform layers**, and the layer version on the **Version Information** tab you prepared.
3. Click **Finalize** to finish creating the layer.

4. Monitor the taskbar to verify that the action is completed successfully.

Once the layer is verified, the packaging machine is removed to minimize the storage space used.

Create or clone an app layer

January 5, 2026

An app layer is a virtual disk that includes one or more applications. Typically, an app layer includes one application. If you include more than one application in a layer, limit it to things that you normally update at the same time.

Create an app layer from scratch

This section walks you through app layer creation, including:

- Requirements and considerations
- Start a new app layer
- Deploy the packaging machine
- Install the application
- Layer integrity messages you might see
- Verify the layer and shut down the machine
- Expedite Microsoft `Ngen.exe` operations, if necessary
- Finalize the layer

Requirements and considerations

An app layer includes one or more applications and related settings. Always install MS Office in an app layer, and never in the OS layer.

- **Anti virus applications:** Always put your antivirus application in an App layer using the instructions provided [here](#). Be strategic with your virus definition file updates. Also, be aware of file marking features, for example, Symantec's Virtual Image Exception Tool. Consider host-based scanning engines, and keep in mind the delay at user logon. Be sure to scan the published layered image, not the layer. Scanning is only done on user access on Citrix Virtual Apps and Citrix Virtual Desktops.
- **MS Office:** Use this [recipe](#) to install Office. For VDI deployments of Office 2010 and later, consider KMS a requirement. For Office 2007 and earlier, consider Volume licensing a requirement. Using other licensing structures is not as convenient, because they require each license to

be activated on each desktop. To persist user settings and data, enable Office 365 User layer stores.OST and streaming files. The search indexes are not stored.

- **Recipes for layering certain applications:** Virtually any application can be layered, but some are easier to layer if you start with the tips we've assembled in our [App Layering Recipes](#) forum. Before you start, consult the forum for tips and procedures about the specific applications you are layering.
- **Applications that require you to add a local user or administrator.** A local user or administrator that you add or change while installing an application on an app layer does not persist. The OS layer preserves any local users or groups that you add, but your app layers do not. Either add the local user or administrator to the OS layer before installing the application or consider installing the application on the OS layer.

Start a new app layer

To create a packaging machine where you can install the application:

1. Log in to the management console and select **Layers > App Layers**.
2. Click **Create Layer** in the **Action** bar.
3. Enter a **Layer Name** and **Version**, both required values. You can also enter other values.
4. On the OS Layer tab, select the OS Layer you want to associate with this app layer.
5. (Optional) The Prerequisite layers tab gives you the option to specify other app layers that must be present while installing the apps on this layer. Only use this when the required apps cannot be included in the same layer. For more information about this advanced feature, see Prerequisite layers in the following sections.
Note: When adding a new version to an existing app layer, you must specify the Prerequisite layers you need. They are *not* carried over from version to version.
6. In the Connector tab, choose a connector configuration that includes the credentials for the platform where you plan to build the layer and the storage location. If the configuration you need isn't listed, click **New** to add one.
7. On the Packaging Disk tab, type a **file name** for the packaging disk, and choose the disk format. This disk is used for the packaging machine, the virtual machine where you install the application.
8. On the Icon Assignment tab, choose an icon to assign to the layer. This icon represents the layer in the Layers Module.
 - To use an existing image, select an image in the image dialog box.
 - To import a new image, click **Browse** and select an image in PNG or JPG format.
 - If you are using a connector with Offload Compositing selected, and you choose one of the icons that came with App Layering, the packaging machine attempts to assign an icon based on the layer's contents when the layer is finalized.

9. On the Confirm and Complete tab, review the details of the app layer and then click **Create Layer**. You can type an optional comment before creating the layer. Your comments appear in the Information view Audit History. After creating the packaging disk, the Tasks bar displays a link to the packaging disk in your hypervisor where you can deploy the packaging machine.
10. Select the **Tasks** page and click the **Packaging Disk** task. Click the info icon to show the full task description, including a link to the location where the packaging machine for this layer is published.

Next, you can deploy the packaging machine for your layer.

Deploy the packaging machine

Select your hypervisor:

- XenServer®, Hyper-V, Nutanix, or vSphere
- Azure
- Other hypervisor (Network File Share)
- Google Cloud Platform (GCP)

XenServer, Hyper-V, Nutanix, vSphere

1. Log in to your hypervisor client (XenServer, Hyper-V Manager, Nutanix Prism, or vSphere).
2. Log in to the App Layering management console, and select the **Tasks** page so you can see the current tasks.
3. Select the **Create App layer** task and click the info icon to see the full task description.
4. Use the URL provided in the task description to navigate to the packaging machine in your hypervisor client.
5. The packaging machine is powered on.

You can now install the applications for this layer on the packaging machine.

Azure The appliance opens the Azure *Custom deployment* template, where you can create the packaging machine.

1. Log in to the Azure portal (<https://portal.azure.com>). **Note:** You must log in before attempting the next step.
2. Go to the App Layering management console and select the **Tasks** page. Select the *Create App layer* task and click the info icon to see details.
3. Use the link in the task details to navigate to the packaging machine in Azure. The *Custom deployment* panel opens.
4. Set the Azure parameters.

- Packaging Machine Name - must conform to Azure virtual machine name requirements.
- Size –virtual machine size.
- Virtual Network and Subnet - for deploying the packaging machine.

IMPORTANT: Make sure the value of the **Resource group location** matches the **Storage account location** that you configured in the connector configuration. If these locations are not the same, the packaging machine fails to deploy. If your deployment does fail, you can paste the link into the browser again and start over.

5. Once your packaging machine is powered on, you can install the application you want to include in the layer.

Other hypervisor (by way of the appliance's Network File Share)

1. Locate the packaging disk in the following directory on the Network File Share:
 \Unidesk\Packaging Disks
2. Copy the packaging disk to a separate location on your hypervisor. This allows space for the files generated by your hypervisor when you use the disk to create a new virtual machine.

Important: Do *not* copy the disk to the Finalize folder until it is ready to finalize. A disk in the Finalize folder cannot be attached to the new virtual machine that you are going to create next.

3. Create a virtual machine using the packaging disk as the boot disk.
4. Power on the packaging machine.

Once your packaging machine is powered on, you can install the application you want to include in the layer.

Install the application

When installing your application on the packaging machine, leave the application as you want users to see it when they log in. The application's state is what users experience every time they access the app.

1. Log in remotely to the packaging machine with the *User account* used to create the operating system.
2. Install the application, along with any drivers, boot-level applications, or files required for the app.
3. If a system restart is required, restart it manually. The packaging machine does *not* restart automatically. If the application you install affects boot-level components, restart the packaging machine as part of finalizing the layer.
4. Make sure that the packaging machine is in the state you want it to be in for the user:

- If the application requires any post-installation setup or registration, complete those steps now.
- Remove any settings, configurations, files, mapped drives, or applications that you do *not* want to include on the packaging machine.

Verify the layer and shut down the machine

Once the application is installed on the packaging machine, verify that the layer is ready to be finalized. A layer is ready to be finalized when all post-installation processing is complete.

To verify that all outstanding processes are complete, you can run the **Shutdown For Finalize** tool on the packaging machine's desktop.

To use the Shutdown For Finalize tool:

1. If you are not logged into the packaging machine, remote login as the user who created the machine.
2. Double-click the **Shutdown For Finalize** icon. A command-line window displays messages detailing the layer verification process.
3. If there is an outstanding operation to be completed before the layer can be finalized, you are prompted to complete the process. If a [Microsoft Ngen.exe](#) operation must be completed, you might be able to expedite the [Ngen.exe](#) operation, as detailed later in this article.
4. Once any pending operations are complete, double-click the **Shutdown For Finalize** icon again. This shuts down the packaging machine, and the layer is ready to finalize.

Layer integrity messages you might see during the finalization process

The following layer integrity messages tell you what queued operations must be completed before the layer is ready to finalize:

- A RunOnce script is outstanding - Check and reboot the Packaging Machine.
- A post-installation reboot is pending - Check and reboot the packaging machine.
- A Microsoft Ngen.exe operation is in progress in the background.
- An MSI install operation is in progress - Check the packaging machine.
- A reboot is pending to update drivers on the boot disk - Check and reboot the packaging machine.
- A Microsoft Ngen.exe operation is needed.
- Software Center Client is configured to run, but the SMSCFG.INI is still present. To learn more about deploying SCCM in a layer, see the article, [App Layering Recipe: How to deploy Microsoft SCCM in a layer](#).

For details about what the layer integrity messages mean and how to debug them, see [Debugging Layer Integrity Problems in Citrix App Layering](#).

You cannot bypass layer integrity messages by shutting down the machine, because the App Layering software stops and returns you to the packaging machine until all of the processes have been completed.

If a Microsoft Ngen.exe operation is in progress, you might be able to expedite it, as described in the next section.

Expedite Microsoft Ngen.exe operations, if necessary

Once all software updates have been installed, you must allow `Ngen.exe` to essentially recompile .NET byte code into native images and construct the registry entries to manage them.

`Ngen.exe` is the Microsoft Native Image Generator, which is part of the .NET system. Windows determines when to run `Ngen.exe` based on what software is being installed and what Windows detects in the configuration.

Important:

When `Ngen.exe` is running, you must let it complete. An interrupted `Ngen.exe` operation can leave you with non-functioning .NET assemblies or other problems in the .NET system.

Normally, `Ngen.exe` is a background operation that pauses when there is a foreground activity. If you want to expedite an `Ngen.exe` operation, you can bring the task into the foreground to complete it as quickly as possible.

To bring the task into the foreground:

1. Open a command prompt as Administrator.
2. Go to the `Microsoft.NET\Framework` directory for the version currently in use:

```
cd C:\Windows\Microsoft.NET\FrameworkNN\vX.X.XXXXX
```

3. Enter the following `Ngen.exe` command to run all queued items. This command processes queued component installs before building assemblies.

```
ngen eqi 3
```

The `Ngen.exe` task moves to the foreground in the command prompt and lists the assemblies being compiled. It is OK if you see compilation messages.

You can use the **Task Manager** to see if an instance of `MSCORSVW.EXE` is running. If it is, allow it to complete, or run `ngen eqi 3`.

Caution:

Do not reboot to stop the task. Allow the task to complete!

4. Ensure that all `Ngen.exe` processes have run to completion.
5. When complete, shut down the virtual machine using the **Shutdown For Finalize** shortcut available on your desktop.

Finalize the Layer

Once the software has been installed and the packaging machine has been verified and shut down, you are ready to finalize the layer.

Hyper-V:

If **Offload Compositing** is selected in your connector configuration, finalization happens automatically as part of the compositing process.

XenServer, Azure, Hyper-V, Nutanix AHV, VMware vSphere Now that the Layer has been verified and shut down, it is ready to finalize.

Hyper-V:

If you are using a connector with **Offload Compositing** selected, this finalization process is automated and you do not have to do these manual steps.

1. Return to the management console.
2. Select **Layers > App layers**, and the layer you prepared.
3. Select your layer version on the **Version Information** tab and click **Finalize** in the Action bar.
4. Click **Finalize** to finish creating the layer.
5. Monitor the taskbar to verify that the action is completed successfully.

Once the layer is verified, the packaging machine is removed to minimize the storage space used.

Other hypervisor (Network File Share) Now that the Layer has been verified and shut down, it is ready to finalize.

1. Copy the Packaging Disk from the folder containing the Packaging Machine files to the Finalize folder on the Network File Share:
 \Unidesk\Finalize
2. Return to the Management Console.
3. Select **Layers > App Layers**.

4. Select your layer version on the **Version Information** tab and click **Finalize** in the Action bar.
5. Monitor the taskbar to verify that the action completes successfully and that the layer is deployable.

Clone an app layer

You can create an app layer that is the same as an existing layer by cloning a specific version of the layer. During the cloning process, you are prompted for information specific to the layer. You can update the app layer by adding versions to it. Since only one version of a layer is cloned, the new layer has just one version to start, even if the layer it was cloned from had many.

To clone a layer:

1. Select the app layer that you want to copy and click **Clone Layer** in the Action bar.
2. Select the Source Layer Version to clone. You can choose the version you want from the drop-down menu.
3. Enter a name for the layer, and a description, if the extra information is helpful. Descriptions are optional.
4. Enter the version, and a description of the version, if the extra information is helpful.
5. On the Icon Assignment tab, select the icon for the new layer.
6. On the Confirm and Clone tab, verify the settings and click the **Clone Layer** button.

A new layer is created with the same layer properties as the source, except for the icon. The layer [priority](#) is higher than that of the source layer because every new App layer has a higher priority than the last app layer created. The new layer size could be smaller than the original, but this just indicates that empty space was removed during cloning. The layer functions the same as the source.

You can use the new layer like any other layer, and it is *not* associated in any way with the original layer.

Advanced app layer options

When creating and updating app layers, keep in mind the following advanced features.

- Custom Layer script
- Layer caching
- Prerequisite layers

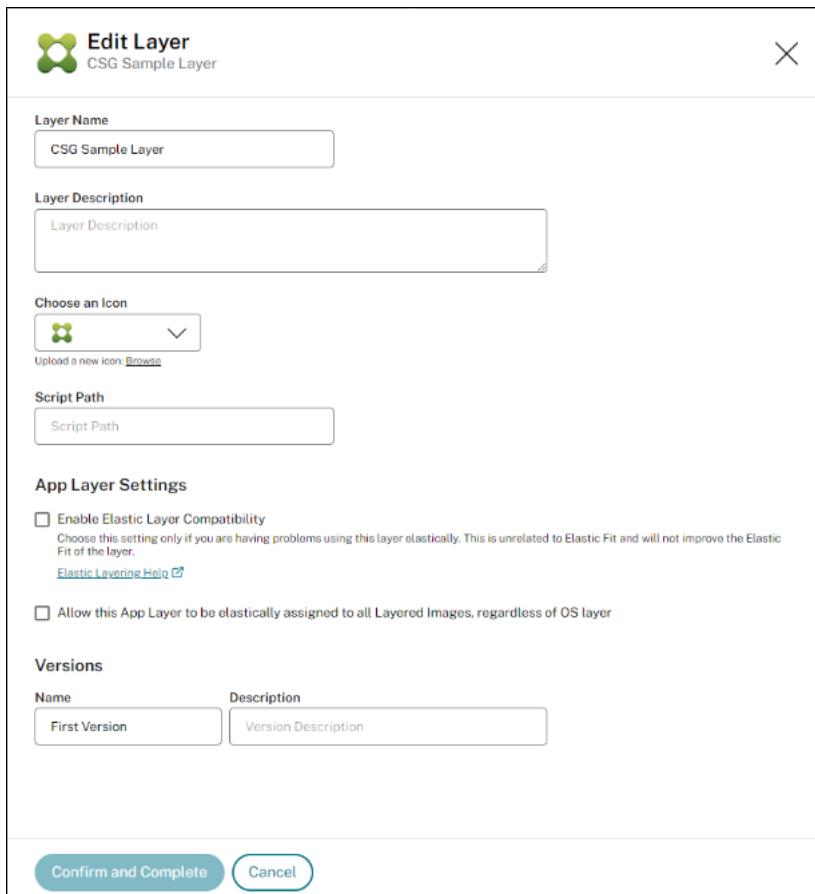
Custom Layer Script

You can include a script in an app layer that runs once, upon system startup. To configure the script, edit the properties of the application layer.

Note:

You can also edit the properties for the layer revision either while the revision is being created, or even after it has been finalized.

The script runs the first time any layered image that includes the app layer starts. If the app layer is elastically layered, the Custom Layer script runs when mounting the app layer disk. Custom Layer scripts are typically used for apps, such as MS Office, that require license activation the first time it starts.



Layer Caching for faster app layer creation

You can use layer caching to speed up layer creation times.

How caching works The first time you create an app layer, if the cache size is set to a large enough value, a template consisting of the boot disk and the empty packaging disk is saved in the cache. The boot disk includes the OS layer, Platform layer, and Prerequisite layer (if any) that are specified in the app layer settings.

Whenever you create an app layer that uses the same OS layer, Prerequisite layer, and Platform layer combination, the App Layering software reuses the template, significantly reducing creation time.

If you then create an app layer that uses a different OS layer, Prerequisite layer, and Platform layer combination, the App Layering software creates a template and adds it to the cache.

Recommended cache size The recommended cache size depends upon how many OS, Platform, and Prerequisite layer combinations you require for your app layers. The number of combinations determines the number of templates saved in the cache.

To estimate the space required for each template:

1. Select the  icon for each OS, Platform, and Prerequisite layer, and look up the **Maximum Layer Size**.
2. Add the Maximum Disk Sizes. The total is the Cache size that you need for that template.

To estimate the space required for the cache, add the size you determined for each of your templates.

Prerequisite layers

Rarely recommended, Prerequisite layers let you include one or more existing app layers on the packaging disk when creating a layer or adding a version to it.

Use prerequisite layers *only if they are required*, because they can add something into the layer that is not required to deploy the current application. This behavior can cause conflict in the future.

When to use Prerequisite layers Prerequisite layers can be required for several reasons:

- When installing the application on the current layer requires the presence of another application. For example, when you install an application that requires Java, and Java is located in a separate layer.
- When the installation of the software adds settings to an existing application. For example, when you install an Office add-in, you must install Microsoft Office first.
- When two applications change the same registry key, and the second application must add to a key and not replace it. For example, two applications that both change login keys in Windows, such as Citrix Agent and Imprivata.

Note

Some of these issues can also be handled by putting the two applications in the same layer rather than using prerequisite layers.

Prerequisite layer characteristics Prerequisite layers have the following characteristics:

- Prerequisite layers are *not included* in the app layer that they are used to create.
- The app layer that you create and each of its Prerequisite layers must use the same OS Layer.
- When adding a *version* to an app layer, Prerequisite layers are *not included* by default. Each time you add a version to a layer, you must select one or more Prerequisite layers.

Layer antivirus apps

January 5, 2026

This article provides the fundamental guidelines for deploying antivirus software in an App Layering or [User Personalization Layer \(UPL\)](#) environment.

For additional antivirus-specific details, see the vendor's documentation for VDI deployments.

Recommendations for all antivirus software

Create a new App Layer to install and maintain your chosen antivirus solution. Citrix® does not recommend installing antivirus software directly on an OS Layer as this makes maintenance more difficult and often leads to contamination of the antivirus state among packaged app and platform layers.

Note:

This does not apply to the UPL images where the antivirus software is required to be installed in the base image.

The following are key points common to most antivirus deployments in app layering (some apply to UPL too):

- If you have already installed antivirus software on your OS Layer, it must be uninstalled and reinstalled in a new app layer.
 - Windows Defender is an exception to this layer advice and is automatically prevented from contaminating other layers by filters built into the App Layering and UPL software.
- Avoid combining other applications with antivirus software on the same app layer.
- Follow the vendor's guidance for VDI deployment (including for UPL).
- Consider disabling automatic updates of the core antivirus software. These updates are better managed through app layer revisions or, with UPL, backups of the base image.
- Daily updates of virus definitions are fine and must not be affected by disabling major updates.

- Add a [UserExclusion](#) file to the antivirus layer to block files and directories from persisting in user layers (including for UPL). See the antivirus vendor's guidelines for non-persistent VDI deployments, for files and/or folders that must not be persisted.
- Add any vendor-recommended registry exclusions to the antivirus layer (including UPL). These are relatively rare, but if necessary, contact Citrix support.

In general, Citrix recommends making a new version of the app layer when the antivirus software has a major update. Once the layer has been updated, assign it to all the templates that use that antivirus app, and redeploy new images to take advantage of the changes in the antivirus software.

Published images, including the UPL master images, can be started outside of the desktop environment to allow antivirus pre-scanning of the assembled image, depending on the antivirus software used.

Elastic layer not enabled

If you are deploying images without elastic layering enabled, consider whether your images are non-persistent or persistent:

For persistent machines, you might want to enable auto-updates to keep the antivirus software up-to-date.

For non-persistent machines, you might not want to turn on auto updates, because the updates occur on the images after every reboot. (The non-persistent machine is reverted when it reboots.)

Layer Docker

January 5, 2026

App Layering supports the Docker platform. To deploy Docker in an App Layering environment, consider the following:

- Docker must be installed in the OS layer since it configures various Windows components.
- You can create an OS revision to install and enable Docker.
- Docker stays dormant until you issue Docker commands. Having Docker installed on the revision causes no issues on later revisions of the OS, packaging machines, or desktops deployed using the OS.
- Do not place any Docker images in the OS layer or any of the app layers. The [vhd \(x\)](#) files written to the disk when the image is downloaded must be moved to the user's writeable layer for the image to run. Docker must open the files for read and write access, which can only be done on the user's volume.

Note:

In a full user layer implementation, the user only needs to download the Docker image once. The image persists through logoff and logon sessions.

- Docker requires a hypervisor that can support nested hypervisors. The target hypervisor hosting the deployed VDI desktops must support nested hypervisors so when a user logs on the machine, they can download and run Docker images.

Important:

The storage location of the user's layer does not matter for this constraint. It's only the deployed VDI machine mounting the user's layer which is required to run on a hypervisor that allows nested hypervisors.

- Docker images can be large. The size of the user layers must be set to a larger size to accommodate the images when they're downloaded. We recommend user layers set to 100 GB following our testing.

Note:

User layers are dynamic disks, so the actual size of the disk won't be full size. Once disk space is used, it won't return to being sparse, so the storage location of the user's volumes needs sufficient disk space to accommodate the full size of the user's layers over time.

App Layering Recipes

January 6, 2026

You can layer most applications with no issues, but there are a few that require extra care.

- **Antivirus applications:** For detailed instructions, see [layering antivirus apps](#).
- **Application guidance for VDI deployments:** For the few applications that require special guidance in virtual environments, we offer more detailed steps in an online Support forum, called [Application Layer recipes](#).

Popular App Layering recipes

The following list includes a sampling of recipes for the few applications that require layering guidance. Unless otherwise noted, these recipes apply to all App Layering versions.

If your application is not listed in the [Application Layer recipes forum](#), you can most likely install it in a layer without any special guidance.

- [Adobe Reader](#)
- [AppSense](#)
- [Bit9](#)
- [Chrome](#)
- [Docker](#)
- [Dropbox](#)
- [Firefox](#)
- [IBM SPSS 21 Licensing Server](#)
- [MS Office, including Office 365](#)
- [MS Teams](#)
- [NVIDIA GRID](#)
- [Print Server](#)
- [SCCM 2012 Client](#)

Deploy elastic layers

January 7, 2026

With the Elastic layers feature, you can deliver narrowly targeted apps outside of the base image. In fact, you can assign layers to specific users on demand that are delivered during logon to Windows.

About elastic layers

An elastic layer is an app layer that you assign to individual users and groups for delivery on demand. Users receive the elastic layers assigned to them in addition to the apps included in the base image.

Elastic layers allow you to give each user a unique set of applications along with the common apps included in the base image. On session hosts, an elastic layer is used across sessions. On standalone desktops, elastic layers are used across floating pools and shared groups.

Studio App Packaging based assignment

Elastic Layers can be assigned through [Studio App Packages](#), enabling you to use a single workflow, from a single console, to deliver applications to all of your Citrix non-persistent workspaces. For more information, see [Studio App Packaging based assignments](#).

App Layering based assignment

App Layering allows you to assign elastic layers to AD users, groups, and machines. For more information, see [App Layering based assignments](#).

Elastic Layering Limitations

You cannot use elastic layer for the following:

- Microsoft Office, Office 365, Visual Studio.
- Applications with drivers that use the driver store. Example: a printer driver.
- Applications that modify the network stack or hardware. Example: a VPN client.
- Applications that have boot level drivers. Example: a virus scanner.

An app layer does *not* preserve a local user or administrator that you add for an app that requires it, but the OS layer does. Therefore, add the local user or administrator to the OS layer before installing the application. Once the app layer is working, you can assign it as an elastic layer.

The user account used to attach elastic layers

By default, when the first user assigned an elastic layer logs on to a machine, all elastic layers assigned to the user are mounted. Other users who log on to the same machine use the same connection as the first user. The connection lasts for 10 hours after the first login, and then all elastic layers are disconnected if the first user is no longer logged in. In a shift-based environment, users on the second shift would be impacted about two hours into the shift (or, 10 hours after the initial user logged on for the first shift).

If you are delivering elastic layers in a shift-based environment, you can change the account used to attach elastic layers. Instead of attaching as the first user who logs in, you can change the user used to attach all elastic layers to the `ulayer` service, which runs under the local `SYSTEM` account. The `SYSTEM` account corresponds to the domain machine account of the machine that the `ulayer` service is running on when accessing the share. This account must have `read-only` access to the file share containing your elastic layers.

- To change the account used to attach elastic layers to `SYSTEM`, create the registry `DWORD` value, and set it to **1**:

```
HKEY_LOCAL_MACHINE\Software\Unidesk\Ulayer:AsSelfAppAttach to  
**1**
```

- To revert to attaching elastic layers using the account of the first user to log in, set the registry `DWORD` value to **0**:

```
HKEY_LOCAL_MACHINE\Software\Unidesk\Ulayer:AsSelfAppAttach to  
**0**
```

- To revert to the default behavior, remove the **DWORD** value:

```
HKEY_LOCAL_MACHINE\Software\Unidesk\Ulayer:AsSelfAppAttach
```

Note:

For multi-session use cases, it is recommended to change the account used to attach elastic layers to the **SYSTEM** account.

Run the **Elastic Fit** analyzer on app layers

Before assigning an app layer elastically, use the **Elastic Fit Analyzer** to determine the likelihood that the layer assignment will be successful.

Elastic Fit analysis

In the Layer Details, the **Elastic Fit** rating indicates how likely it is that the layer works when elastically assigned.

Good Elastic Fit. This layer works when deployed elastically.



Poor Elastic Fit. Delivering the layer elastically is not likely to work when deployed elastically. The layer can behave differently than it does when it is deployed in a layered image.



Elastic Fit details

You can learn more about an app layer's Elastic Fit rating by expanding the Elastic Fit Analysis. If the Elastic Fit is less than ideal, the list of violated rules is displayed.

Low Severity Warning. Delivering the layer elastically is unlikely to cause any change in behavior or functionality for most applications.



Medium Severity Warning. Delivering the layer elastically can cause minor changes in behavior or functionality for some applications.



High Severity Warning. Delivering the layer elastically is likely to cause significant changes in behavior or functionality for many applications.



Note:

If you receive a warning that a master key file change has been detected, and you did not intentionally change that file, set the value of the `DeleteMasterKeys` flag in the registry location `HKLM\System\ControlSet001\Services\Uniservice` to 1 (true). Now when the app layer is finalized, master key files are deleted from the layer. This value is not persistent and only works per revision. It must be set each time a revision of the layer is created.

Analyze an app layer's Elastic Fit

All new versions of a layer version are analyzed for elastic layering compatibility when they are finalized. To analyze existing app layers for Elastic Fit:

1. Log in to the management console.
2. Select **Layers > App Layers**.
3. Select the layer to analyze, and click **Analyze Layer**.
4. On the Select Versions tab, choose the Layer Versions to analyze.
5. On the Confirm and Complete tab, click **Analyze Layer Versions**. The analysis takes seconds.
6. To see the **Elastic Fit Analysis**, select the app layers module, move the mouse pointer over the layer icon and click the **Info** icon.
7. Expand the **Version Information** for each layer version, and look for the Elastic Fit rating.
8. For a detailed report, expand the **Elastic Fit Details**. If the Elastic Fit is less than ideal, the list of violated rules will be displayed.
9. You can display the AD tree and hide the violated rules by clicking a button acknowledging that the layer is unlikely to work as expected.

Upgrading from earlier releases

After upgrading from an early App Layering release, the Elastic Fit Detail shows that existing layer versions have not been analyzed. The versions have a single *High* severity Elastic Fit Detail, and a *Poor* Elastic Fit. For an accurate reading, run the analysis on existing layer versions.

Troubleshooting elastic layer issues

You can diagnose the source of an elastic layering issue by finding out whether the layer is being delivered, and whether the layer is working correctly. If needed, collect data for support, as described

here.

Is the issue with layer delivery?

Are the things that you'd expect to see when this app is installed there?

- Do you see the files and registry entries for the layer?
- If the app is supposed to be in the Start menu, is it there?
- If you expect there to be a shortcut for the app on the user's desktop, is there one?

If you discover that app delivery is an issue, you can collect the following data, open a case, and send the data to support.

1. Collect the data from these logs:

- Windows App Event log –In the **Windows Event Viewer** under **Windows Logs**, export the application event log as an EVTX file.
- App Layering Service log (ulayersvc.log) –C:\ProgramData\Unidesk\Logs\ulayersvc.log

2. Collect the values of these Registry keys:

- HKEY_LOCAL_MACHINE\SOFTWARE\Unidesk\ULayer:AssignmentFile
- HKEY_LOCAL_MACHINE\SOFTWARE\Unidesk\ULayer:RepositoryPath

3. Collect the contents of the assignment (ElasticLayerAssignments.json) and Layers (Layers.json) files from the Repository Path.

4. Contact Support.

Is the issue an operational one?

Any of these behaviors can indicate an elastic layering issue:

- The app is being delivered but doesn't launch correctly.
- An operation within the app doesn't work correctly.
- A licensing problem or a security issue.
- The app launches, but then misbehaves, for example, it crashes on startup, or starts up but doesn't work right.

If the problem with the layer is operational, test the app layer in the base image to rule out general layering issues:

1. Add the app layer to an image template, and publish a layered image that includes the app layer.
2. Log in as a user who is *not* assigned the layer elastically, and make sure that the application is operational in the base image.

3. Contact Support with your findings.

App Layering based assignments

January 5, 2026

App Layering allows you to assign elastic layers to AD users, groups, and machines.

Prerequisites

- .NET Framework 4.6.2 or higher is required on any image where elastic layers are deployed.
- The app layers that you want to assign as elastic layers.

Elastic layer assignments

You can deliver a specific app layer version to members of a group each time they log in to their desktops. You assign the app layer version as an elastic layer. A copy of the layer is then stored in the appliance's Network File Share, and delivered on-demand to the assigned AD users and groups, in addition to the layers that they receive via the base image.

To use this feature, you add *Elastic Assignments* specifying which users and groups receive each revision of the app layer. You then publish your base image with the **Elastic Layering** setting enabled.

How users access elastic layers assigned to them

When a user logs in to their session or desktop, they receive an elastic layer in the following cases:

- The user (an AD user in the management console) is assigned the layer.
- An AD group that the user belongs to is assigned the layer.
- The machine that the user is logging into is a member of an AD Group that receives the elastic layer.
- The machine that the user is logging into is associated with an AD Group that is assigned to the layer via the management console.

When a user is assigned more than one version of a layer

When a layer is assigned directly to a user, and indirectly to one or more of the user's groups, they receive the most recent directly assigned version. For example:

- If the user is assigned **Version 2**, and a group that the user belongs to is assigned **Version 3**, the user gets **Version 2**.
- If two or more groups that the user belongs to are assigned different versions of the same layer, the user receives the most recent version of the layer assigned.

When a user receives an app layer both in the base image, and as an elastic layer

When an app layer is included in the base image, do not assign it to the same user as an elastic layer. If the user does end up with the same layer assigned both ways, they receive the elastic layer, no matter the version.

Enable elastic layers on your base images

You can enable elastic layers on your base (layered) images by configuring the image template that you use to publish them:

1. In the management console, select the image template to use for publishing your layered images.
2. Select the **Images** tab, and then the image template on which you want to enable elastic layering.
3. Select **Edit Template** from the Action bar.
4. Select the **Layered Image Disk** tab.
5. In the **Elastic layering** field, select **Application Layering**.
6. Select the Confirm and Complete tab, and click **Save Template and Publish**.
7. Use your provisioning system to distribute the virtual machines.

When the users log in, the desktop includes an icon for each of their elastic app layers.

Elastically assign an app layer to AD users and groups

The first time you assign an app layer elastically, we recommend starting with a simple app like **Notepad++** or **GIMP**.

1. Log in to the management console as an Admin user, and select **Layers > App Layers**.
2. Select an app layer that you do *not* plan to include in the base image, and select the app version you want to assign.
3. Click **Update Assignments**.
4. Select the version of the app layer that you want to assign users.

5. Skip **Image Template Assignment**. This is for assigning the layer to an image template.
6. Select the users and groups that you want to receive this app layer version.
7. Review your selections, and click **Assign Layers**.

When the users log in, there is an icon for each elastic layer they've been assigned.

Elastically assign an app layer to users via machine assignments and associations

You can assign layers to a machine by adding the machine to, or associating the machine with, the AD Group. Then elastically assign the app layers to the AD Group.

The layers assigned to the machine are available to every user who successfully logs into that machine. The App Layering Service scans for changes to the machine's AD group memberships and associations every 10 minutes. When the users log in, they see an icon for each elastic app layer they've been assigned.

Use Active Directory to add the machine to the AD Group

Assuming you have a published layered image booted in your environment, you can add the machine to an AD Group, and assign elastic layers to the AD Group.

1. Use Active Directory (AD) to add the machine to an AD Group.
2. Select an app layer that you do *not* plan to include in the base image, and elastically assign the layer to an AD Group.
3. You can wait for AD to propagate the changes and for the App Layering Service, or you can force the App Layering Service to update its list of machine groups by doing *one* of the following:
 - Wait for the App Layering Service to detect the changes (within 10 minutes by default).
 - Restart the App Layering Service.
 - Reboot the App Layering Service Machine.
 - Run the **refresh.groups** command:
`C:\Program Files\Unidesk\Layering Services\ulayer.exe refresh.groups`

Example

You start with an AD User, and AD Group, and a machine that you provisioned using a layered image.

- AD User: *Kenya*
- Kenya has no elastic assignments.

- AD Group: *Marketing*
 - The *Marketing* group includes the member Kenya.
- Machine: *ElasticTestMachine*
 - The *ElasticTestMachine* base image includes the *MS Office App Layer*.

In this example, you elastically assign the *Chrome App layer* to *ElasticTestMachine*:

1. In AD, you add the machine *ElasticTestMachine* to the *Marketing* AD Group.
2. In the management console you elastically assign the *Chrome App Layer* to the *Marketing* Group.
3. When Kenya, who is part of the *Marketing* group, logs into *ElasticTestMachine*, she receives both the *MS Office App layer*, which is in the base image, and the *Chrome App layer*.
4. When any user who is *not* in the *Marketing* group logs into *ElasticTestMachine*, they also receive both Layers: *MS Office* because it is in the base image, and *Chrome* because the *ElasticTestMachine* is a member of the *Marketing* AD Group.

Manage elastic assignments

You can:

- Add an elastic assignment.
- Update an app layer and elastically assign the new version of the layer.
- Remove elastic assignments.
- Debug an elastic assignment.

Update an app layer and its elastic assignments

You've added elastic assignments to an app layer, and users are accessing the app as expected. A new version of the application is released, so you update it with a new version to the layer. Now you need to assign the new version to the users who have the layer.

1. Log in to the management console and select **Layers > App Layers**.
2. Select the elastically assigned app layer that you updated.
3. Click **Version Information > Update Assignments**.
4. Select the new version.
5. Skip the **Image Template Assignment** tab.
6. In the **Elastic Assignment** tab, there's a list of Users and Groups who have been assigned a different version of the selected layer. Select the users and groups to whom you want to assign the new version of the layer.

Notes:

- If the list is long, use the **Search** field to filter the results.
- If the list is empty, click the check box called, **Show AD users and groups already at this version**. A list of grayed out names appears. These users have already been assigned the version.

7. On the Confirm and Complete tab, verify the Users and Groups that you want to receive the new version.
8. Click **Update Assignments**.

Remove a layer's elastic assignments

1. Log in to the management console and select **Layers > App Layers**.
2. Select the app layer for which you want to remove assignments, and select **Remove Assignments**.
3. Select the assigned templates from which you want to remove the layer. The assignments for the layer are listed.
If the list is long, use the Search field to filter the results.
4. On the Confirm and Complete tab, verify that the correct image templates are selected to receive the new version.
5. Click **Remove Assignments**.

OS layer switching for elastic layers

For best results, when assigning app layers as elastic layers, you can assign app layers that have the same OS layer as the one used in the layered image. However, with this traditional approach, you might need to create and manage additional copies of some app layers, one for each OS layer you deploy with.

For greater flexibility and to reduce the need to duplicate app layers, you can assign an elastic layer to users on layered images that use a different OS layer. You can enable this ability in the application layer properties by selecting the **Allow this App Layer to be elastically assigned to all Layered Images, regardless of OS layer** check box. All [elastic layering limitations](#) are valid when switching OS layers.

When it might work well

For simple applications that can be installed on any OS. Example: Notepad++, WinRAR, 7Zip

When it might not work well

For complex applications whose installation depends on the OS installed. Example:

- If you use a Windows 11 OS layer to create the app layer, and the image assigned as a Server 2022 OS layer, then the application might not work as expected.
- Applications that are dependent on a specific version of .Net might not run successfully if the new OS doesn't have the correct version of .Net installed.

Note:

- It is recommended that you use the same OS class and OS revisions that are close to each other. Example: You can use two Windows 11 25H2 revisions that are one week apart.

Elastic layer compatibility mode

By default, elastic layers assigned through App Layering are composited into the image after the user logs on. If an elastic layer doesn't load correctly, try enabling **Elastic Layer Compatibility Mode**. With Compatibility Mode enabled, the elastic layer starts loading before login is complete.

Notes:

- Compatibility mode is required when using published applications, as the layer must be mounted before launch
- We recommend disabling compatibility mode unless an elastic layer doesn't work as expected
- Enabling this setting on too many layers slows login times

Studio App Packaging based assignments

January 7, 2026

Elastic Layers can be assigned through [Studio App Packages](#), enabling you to use a single workflow, from a single console, to deliver applications to all of your Citrix non-persistent workspaces.

Prerequisites

- .NET Framework 4.6.2 or higher is required on any image where elastic layers are deployed.
- The app layers that you want to assign as elastic layers.
- Studio App Packaging assignment of elastic layers requires CVAD/DaaS 2511 or later.

Enabling app layers for Studio use

To enable an app layer for use with Studio App Packages:

- Go to **Layers** from the App Layering UI
- Edit the desired app layer
- Under **App Layer Settings** select the *Enable this App Layer for use with Studio App Packages* checkbox

Preparing an image for elastic layer assignment through Studio

For both single and multi-session Windows operating systems, when installing the VDA either within a platform layer to be included in a published image or within a non-layered image, the optional VDA component [App Packages Delivery Component](#) must be included to enable elastic layer discovery within Studio App Packages.

Layered Images

Enable *Application Layering* under the **Elastic Layering** dropdown in the image template. Ensure VDA 2511 or later is installed in a platform layer included in the published image.

Non-Layered Images

- **Multi-session VDA installation**

Install Virtual Delivery Agent (VDA) 2511 or later. If an older version of the VDA is already installed, uninstall the old version first. When installing the new version, be sure to select and install the optional component **Citrix Elastic App Layering**.

- **Single-session VDA installation**

Use [Citrix User Personalization Layer \(UPL\)](#) in a single-session, non-layered image to enable Studio-assigned elastic layers. Configure the image as stipulated for UPL.

OS compatibility for elastic layers

For the best results, elastic layers should be assigned to images running an OS that is as close as possible to the OS used to author the layers. It is recommended that you use the same OS class and OS revision, e.g. Windows 11 25H2. If you use a Windows 11 OS layer to create the app layer, and the image it is assigned to runs Server 2022, then the application might not work as expected.

Deploy user layers

January 5, 2026

User layers persist each user's:

- Profile settings
- Data
- Locally installed applications in non-persistent VDI environments

When you enable user layers on an image template, systems provisioned using the resulting layered images provide every user with a user layer.

When a user logs in to a desktop that is user layer-enabled, a new Search index database is created. The index incorporates search information from the user layer and any elastic layers. The Search feature is only available when the indexing is complete.

This topic explains how to enable user layers on an image template, and the resulting layered images. Systems that you provision using the images provide every user with a user layer.

Types of user layers

You can enable the following types of user layers:

- **Full:** All of a user's data, settings, and locally installed apps are stored on their user layer.
- **Office 365:** (Desktop systems) Only the user's Outlook data and settings are stored on their user layer.
- **Session Office 365:** (Session hosts) Only the user's Outlook data and settings are stored on their user layer.

You can enable a Full user layer, an Office 365 user layer, or a Session Office 365 user layer. The full user layer includes everything that the Office 365/Session Office 365 user layer saves, along with the settings and data for other applications.

Note:

Office 365 and Session Office 365 are deprecated.

Requirements

Before enabling user layers, be sure to meet the requirements that apply to the following types of user layers:

- All types of user layers
- Full user layers
- Office 365 and Session Office 365 user layers

All user layers

To enable user layers you need:

- Adequate network bandwidth. Bandwidth and latency have a significant effect on the user layer. Every write goes across the network.
- Enough storage space allocated for users' data, configuration settings, and their locally installed apps. (The appliance uses the main storage location for packaging layers, publishing layered images, and serving up Elastic layers.)

Full user layers

- When using Profile Management with a Full user layer you must clear the deletion of the user's information on logoff. Depending on how you are deploying the settings, you can clear deletion using either:
 - A Group Policy Object.
 - The policy on the Delivery Controller™ (DDC).

Office 365 and Session Office 365 user layers

- Use a profile manager, such as the Citrix Profile Manager. Otherwise, Outlook assumes that every user who logs in is new and creates OS files for them.
- The Office layer must be included in the image template and deployed in the layered image. However, you can use other Elastic layers with an Office 365 user layer.
- Microsoft Office is supported as an app layer in a published image only, not as an elastic layer.
- Any change to the default location of the search index files is *not* preserved in the Office 365 layer.
- This feature has been tested for one desktop per user at a time (Single sign-on).

Note:

Office 365 and Session Office 365 are deprecated.

Compatibility

Full user layers are supported on the following platforms:

- **Operating systems:**

All operating systems must be configured in single-user mode to work with user layers. Servers cannot be used in multi-user mode. User layers on session hosts are not supported.

- Windows 10, 64-bit
- Windows 11, 64-bit (only if deployed to a platform enabled for offload compositing)
- Windows Server 2016, single-user mode only
- Windows Server 2019, single-user mode only

- **Publishing platforms:**

User layers are supported on the following publishing platforms.

- Citrix Virtual Desktops™

Applications that are *not* supported on a user layer

The following applications are not supported on the user layer. Do *not* install these applications locally:

- Enterprise applications: Enterprise applications, such as MS Office and Visual Studio, must be installed in app layers. User layers are based on the same technology as elastic layers. As with elastic layers, never use user layers for these enterprise applications!
- Applications with drivers that use the driver store. Example: a printer driver.

Note:

You can make printers available using Group Policies. See GPO-installed printers in the following section.

- Applications that modify the network stack or hardware. Example: a VPN client.
- Applications that have boot-level drivers. Example: a virus scanner.
- Applications that require you to add a local user or group. Local users and groups that you add as part of installing an application only persist in the OS layer. Consider installing an application on a layer that will be included in the base image, with the required user or administrator added to the OS layer.

Windows updates

Windows updates must be disabled on the user layer.

Outlook store add-ins

Citrix Profile Management disables Store add-ins.

The first time Outlook starts, the **Store/Add-ins** icon on the ribbon displays a window with a long list of add-ins. During the initial login, if you install add-ins, they appear on the ribbon on subsequent logins. If you do not install the add-ins, the **Store/Add-ins** icon displays a blank white window.

GPO-installed printers

For users on non-persistent desktops running Windows 10, you can install printers using a Group Policy. With a policy in place, the printers are listed in users' Devices and Printers, application printer settings, and device manager.

To set up GPO-installed printers:

1. Enable user layers in the image template.
2. Ensure that the desktop is joined to the domain (on the Platform layer).
3. Create a group policy to deploy each network printer, and then assign it to the machine.
4. When logged in as a domain user, verify that the printer is listed in Devices and Printers, Notepad, and device manager.

User layer format

User layer virtual disks are created using the VHDX format. You can still use the existing user layer VHD files without converting or recreating them. However, when both the VHD and VHDX files exist in the same folder, the VHDX file takes precedence.

How to override VHDX format for user layer virtual disk creation

You can change the behavior to force the user layers that are created to use the VHD format. To override this behavior, use the following system registry parameter:

- Path: `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ulayer`
- Name: `DefaultUserLayerVHDXDisabled`
- Type: `DWORD`
- Data: `1`

User Layer/UPL space reclamation

You can use **User Layer/UPL space reclamation** to automatically optimizes the VHDX files every time the user logs off.

How to enable User Layer/UPL space reclamation Before enabling **User Layer/UPL space reclamation**:

- Optimize drives service `defragsvc` must be enabled and running
This service is disabled for existing OS layers.
- Citrix recommends creating a new OS layer version before enabling the feature.

You can enable **User Layer/UPL space reclamation** in either of the following ways:

- Using Citrix Studio
- Using Windows registry editor
 - Path: `HKLM\SOFTWARE\Policies\Citrix\UserPersonalizationLayerConfig`
 - Name: `UserLayerCompactionEnabled`
 - Type: `DWORD`
 - Data: 1 (Default value: 0)

This feature is disabled by default.

Enable user layers on a layered image

To deploy user layers, you enable the layers using the settings in an image template. For detailed steps, see [Create or clone an image template](#). The rest of this article supplies details about sizing, storing, configuring security on, moving, and repairing user layers. It also covers the steps for customizing notifications for end users.

User layer location

When an image template has user layers enabled, the images you publish persist users' data, settings, and locally installed apps.

When user layers are enabled, you must add storage locations for the layers.

Important:

Do *not* allow user layers to be saved on the appliance's main file share. Otherwise, space can be depleted for:

- Upgrading the software.
- Serving up elastic layers to users.
- Saving files that you are moving to a hypervisor for which there is no supported connector.

The first storage location added to the appliance becomes the default location for user layers that are not associated with any other storage location. When you add more storage locations, they are listed in priority order.

You can assign groups of users to each storage location that you add.

Where a user layer is stored when the user belongs to more than one group

If a user belongs to more than one group and those groups are assigned to different storage locations, the person's user layer is stored in the highest priority storage location.

If you change the priority order of the storage locations that the user is assigned to *after* the person's user layer was saved to the highest priority location, the data saved up until that point remains in the original location. To preserve the person's user layer, you *must* copy their user layer to the new highest priority location.

How to specify the user layer file share location on a specific image

You can support a user who needs to access two separate images at the same time, where both images:

- Need the persistence of user layers.
- Were created using the same OS layer.

To configure user layer file share assignments:

1. Add the following Registry key in one or more of your published images *before* any user logs in:

```
[HKLM\Software\Unidesk\ULayer]  
"UserLayerSharePath"
```

You can add the preceding key to the platform layer, to an app layer, or as a machine group policy.

If you add the **UserLayerSharePath** key to the image before a user logs in, the appliance ignores the user layer share assignments. Instead, all users on the machine use the specified share for user layer VHDX or VHDs. The `\Users` subtree is appended to this key to locate the actual layers.

How to specify a custom user layer path

You can set a custom path by creating a `REG_SZ` value called `CustomUserLayerPath` in the `HKLM\Software\Unidesk\Ulayer` key. The `HKLM\Software\Unidesk\Ulayer` key can include environment variables and Active Directory (AD) attributes.

In the `CustomUserLayerPath` value, all system variables can be expanded, but the only user variables that can be expanded are `%USERNAME%` and `%USERDOMAIN%`. The full path is:

`<CustomUserLayerPath>\<OSID_OSNAME>`

If you set the custom user layer path using GPO, use `%<USERNAME>%` and `%<USERDOMAIN>%` to prevent the GPO from expanding the paths.

- If `CustomUserLayerPath` is defined, it is used instead of any other path.
- If `CustomUserLayerPath` is undefined, `UserLayerSharePath`, which is inside the same key, is used.
- If `UserLayerSharePath` is undefined, the `StorageLocation` listed in the JSON for the App Layering appliance is used. You can edit the `UserLayerSharePath` in the management console, in the **System > User Layer Storage Location** setting.
- If there is no `StorageLocation` listed in the JSON from the App Layering appliance, then `RepositoryPath` is used. You can edit the `RepositoryPath` at the same registry location as `CustomUserLayerPath` and `UserLayerSharePath`.
- When `CustomUserLayerPath` is defined, the path where user layers are created is the expanded path, plus `\<OSID_OSNAME>`. All other paths are share paths, and they are appended to `\Users\<Domain_UserName>\<OSID_OSNAME>`.

If you use AD attributes, the attributes must be enclosed in hashes (for example, `#aAMAccountName#`). Custom AD attributes can be used to define organizational variables, such as locations or users. Attributes are case-sensitive.

Examples:

- `\server\share\#sAMAccountName#` stores the user settings in the UNC path
- `\server\share\JohnSmith` (if `#sAMAccountName#` resolves to JohnSmith for the current user)

Where user layers are created on the appliance

On the appliance's network file share, user layers are created in the `Users` folder. For example:

```
1 \MyServer\*MyShare*\Users
```

Each user has their directory within the `Users` directory. A user's directory is named as follows:

```
1  Users\\*DomainName_username*\\*OS-Layer-ID-in-hex*_*OS-Layer-name*\\*username*.vhf
```

For example:

- User's login name: **jdoe**
- User's Domain: **testdomain1**
- OS layer: MyOSLayer (ID is in hexadecimal format: *123456*)
- User layer is created in:

```
1  \MyServer\MyShare\Users\testdomain1_jdoe\123456_MyOSLayer\jdoe.vhd
```

Where users can access their user layer

When Full user layers are created, users can access the entire C:\ (subject to Windows rights and company security on the directories).

When Office 365 layers are created, the user layers directory is redirected to the Office 365 layer:

```
1  C:\user\\<username\>\Appdata\local\Microsoft\Outlook
```

Add a storage location

To add a storage location for an image's user layers:

1. Log in to the management console.
2. Select **System > User Layer Storage Locations**. A list of file shares is displayed, except for the appliance's main file share.
3. Select **Add Storage Location**, and enter a **Name** and **Network Path** for the new location.
4. Under **Assignments**, click **Add Groups**.
5. Expand the directory, select the desired users, and click **Save**.
6. Click **Confirm and Complete** to add the storage location.

Once the storage locations are added, you must set security on the user layer folders.

Configure security settings on user layer folders

You can specify more than one storage location for your user layers. For each storage location (including the default location) you need to create a \Users subfolder and secure that location.

A domain administrator must set the security on each user layer folder to the following values:

Setting name	Value	Apply to
Creator Owner	Modify/Delete subfolders and files*	Subfolders and Files only
Owner Rights	Modify	Subfolders and Files only
Users or group	Create Folder/Append Data; Traverse Folder/Execute File; List Folder/Read Data; Read Attributes	Selected Folder Only
System	Full Control	Selected Folder, Subfolders, and Files
Domain Admins, and selected Admin group	Full Control	Selected Folder, Subfolders, and Files

*On some servers, **Creator Owner** requires **Delete subfolders and files** (an advanced permission), so that App Layering can clean up after user layer repairs.

To configure security on user layer folders:

1. Log in to the management console.
2. Click **System > User Layer Storage Locations**. The file shares displayed are the storage locations defined for user layers. Say you've defined three storage locations so that you can manage storage for Group1 and Group2 separately from everyone else in the organization:
 - *Default location* - \\MyDefaultShare\UserLayerFolder\
 - *Group1* - \\MyGroup1\Share\UserLayerFolder\
 - *Group2* - \\MyGroup2\Share\UserLayerFolder\

Note: The appliance's main file share, which is used for storing OS, app, and platform layers, is *not* listed as a user layer storage location. For more about the App Layering file share, see [Setting up a file share](#).

3. Create a \Users subdirectory under each file share:

```
1  \\MyDefaultShare\UserLayerFolder\Users\  
2  
3  \\MyGroup1Share\UserLayerFolder\Users\  
4  
5  \\MyGroup2Share\UserLayerFolder\Users\
```

1. Apply the preceding list of security settings to each subdirectory under \Users.

Apply User personalization layer Studio policies to user layers

When user layers are enabled on a layered image, you can override the default repository path and layer size for the user layers by configuring the corresponding Citrix Studio Policies:

- **User Layer Repository Path:** Defines where on the network to access the user layers.
- **User Layer Size GB:** Defines how large to permit the user layer disks to grow.

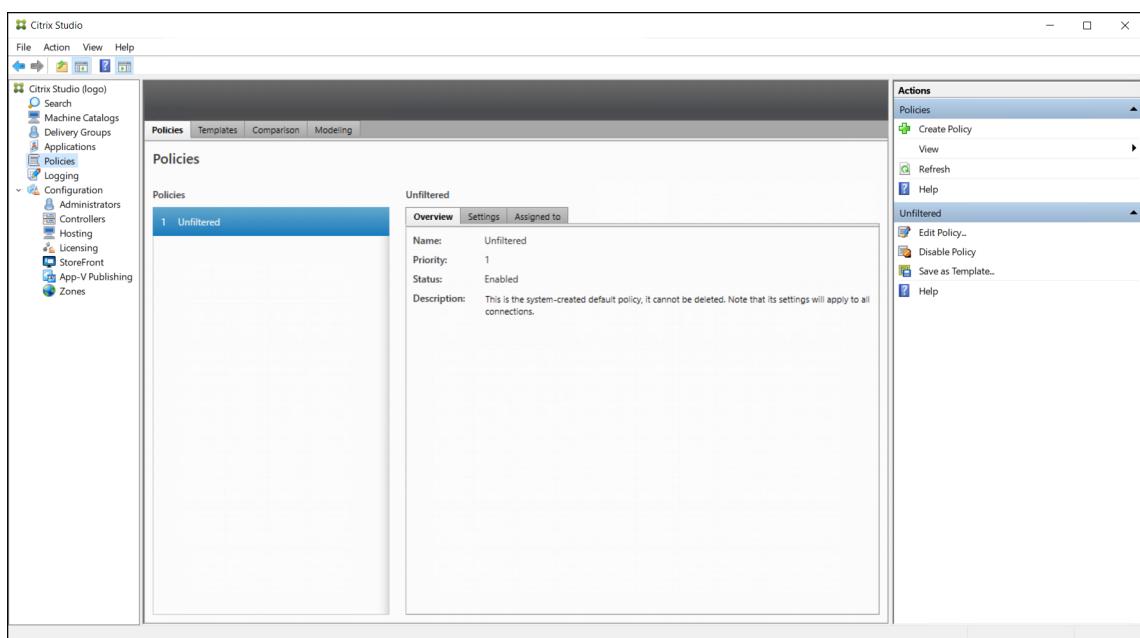
If the published image is running a supported version of the VDA, and these policies have been defined, the path and size defined in the policies are given the highest priority.

An increase to the assigned user layer size takes effect the next time the user logs in. A decrease in assigned user layer size does not affect existing user layers.

Define the Studio policies for an image's user layers

To configure the Citrix Studio policies for a layered image's user layers:

1. In Citrix Studio, select **Policies** in the navigation pane:

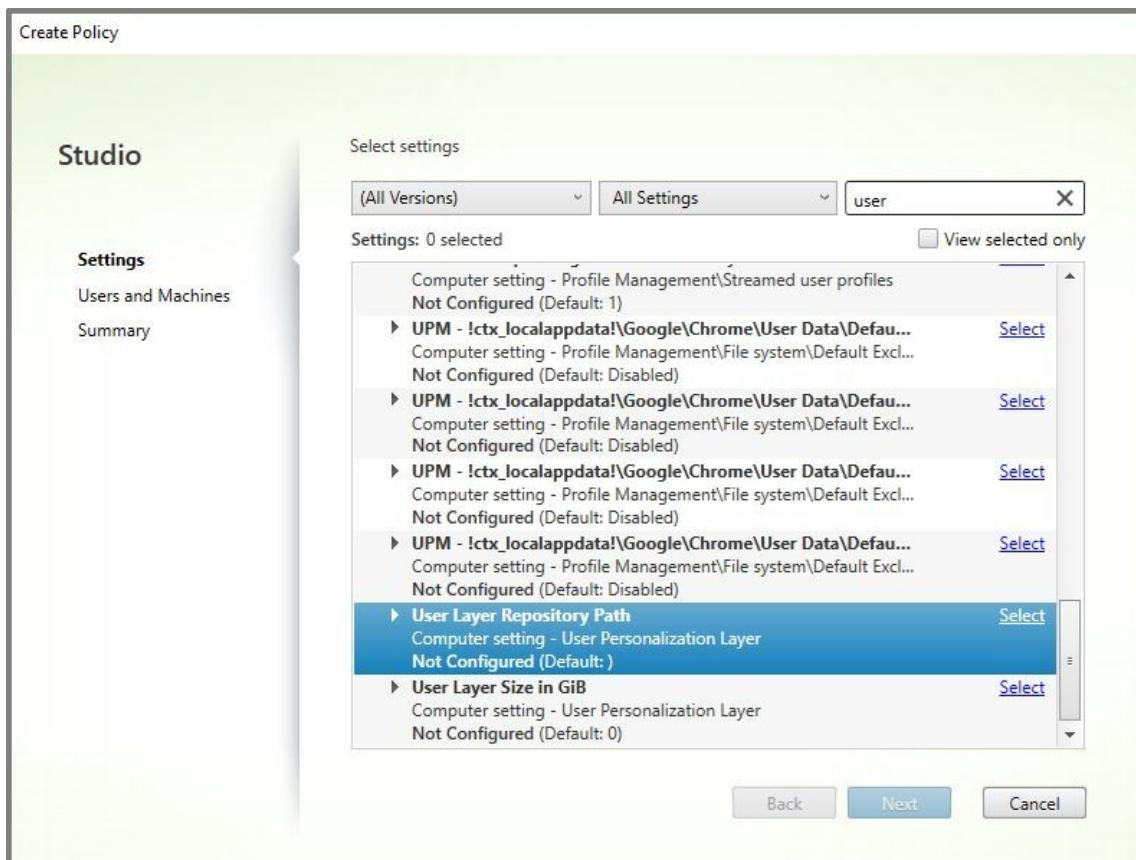


2. Select **Create Policy** in the Actions pane. The Create Policy window appears.
3. Type 'user layer' into the search field. The following two policies appear in the list of available policies:
 - User Layer Repository Path
 - User Layer Size GB

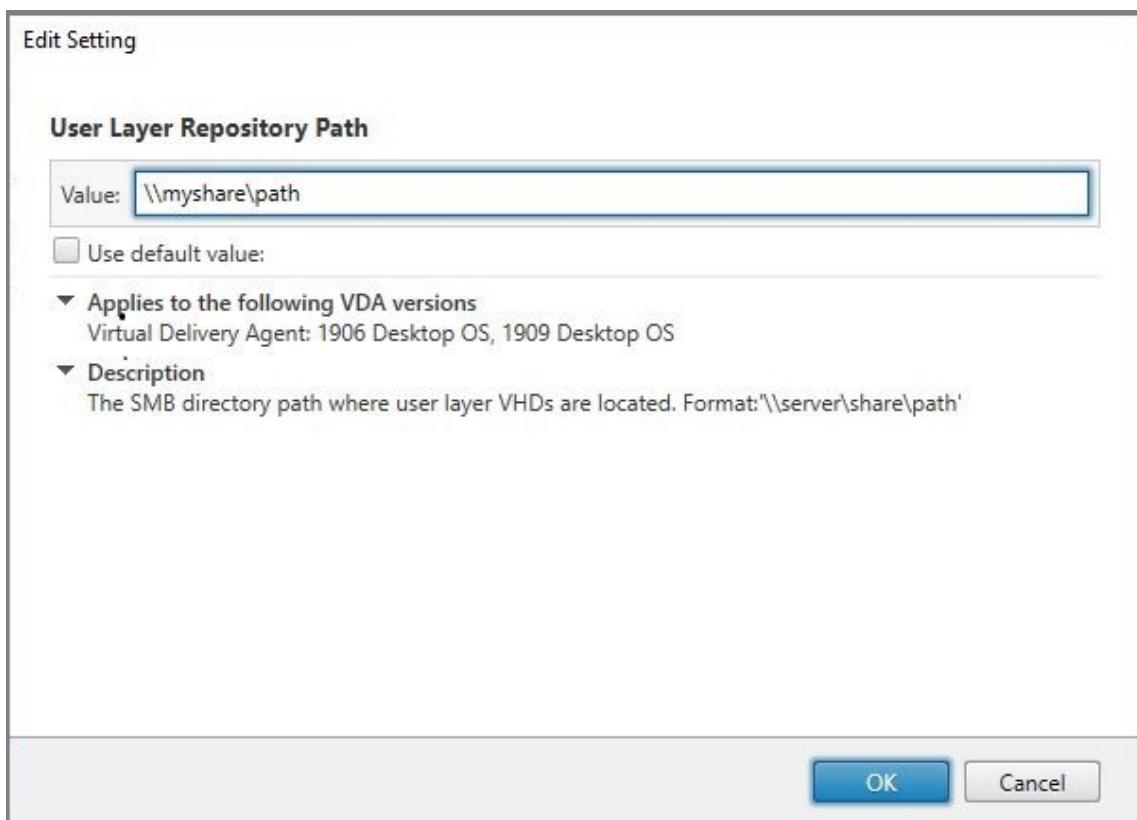
Note:

An increase to the assigned user layer size takes effect the next time the user logs in.
A decrease to assigned user layer size does not affect existing user layers.

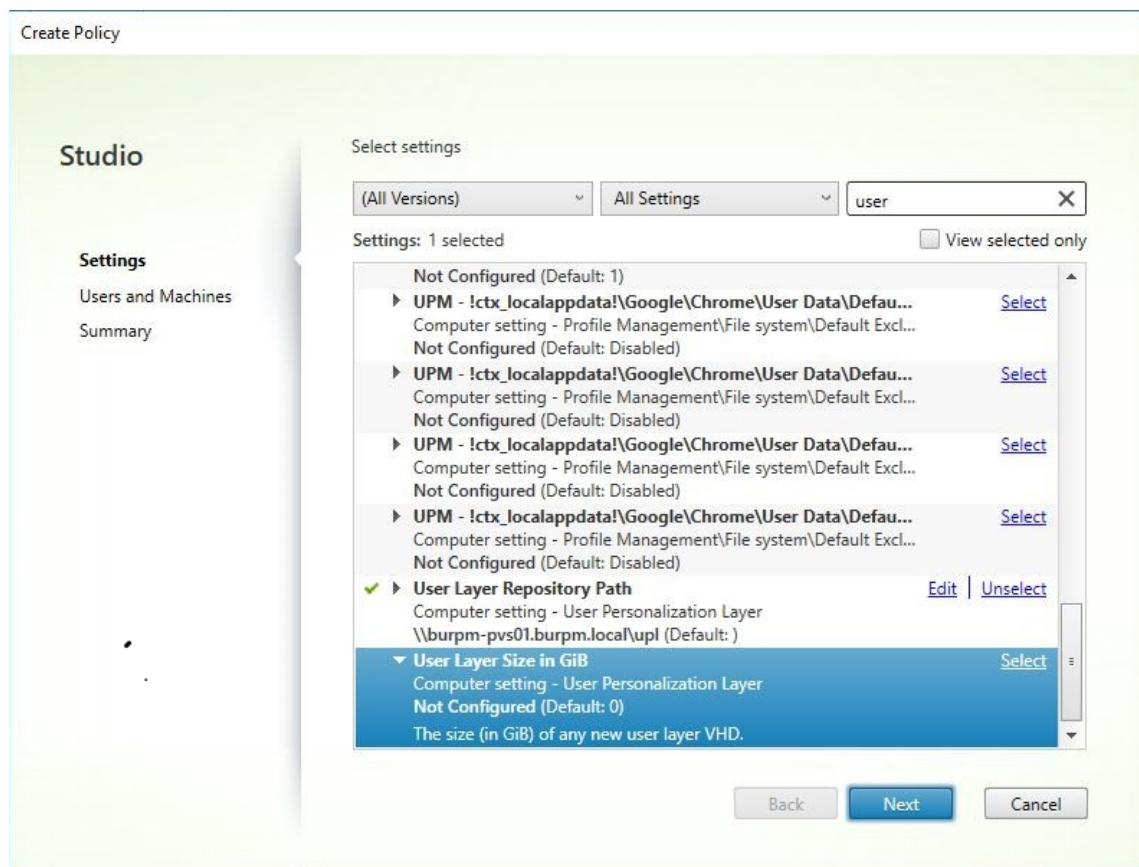
4. Click **Select** next to **User Layer Repository Path**. The Edit Setting window appears.



5. Enter a path in the format `\server name or address\folder name` in the **Value** field, Click **OK**:



6. Optional: Click **Select** next to User Layer Size in GB:



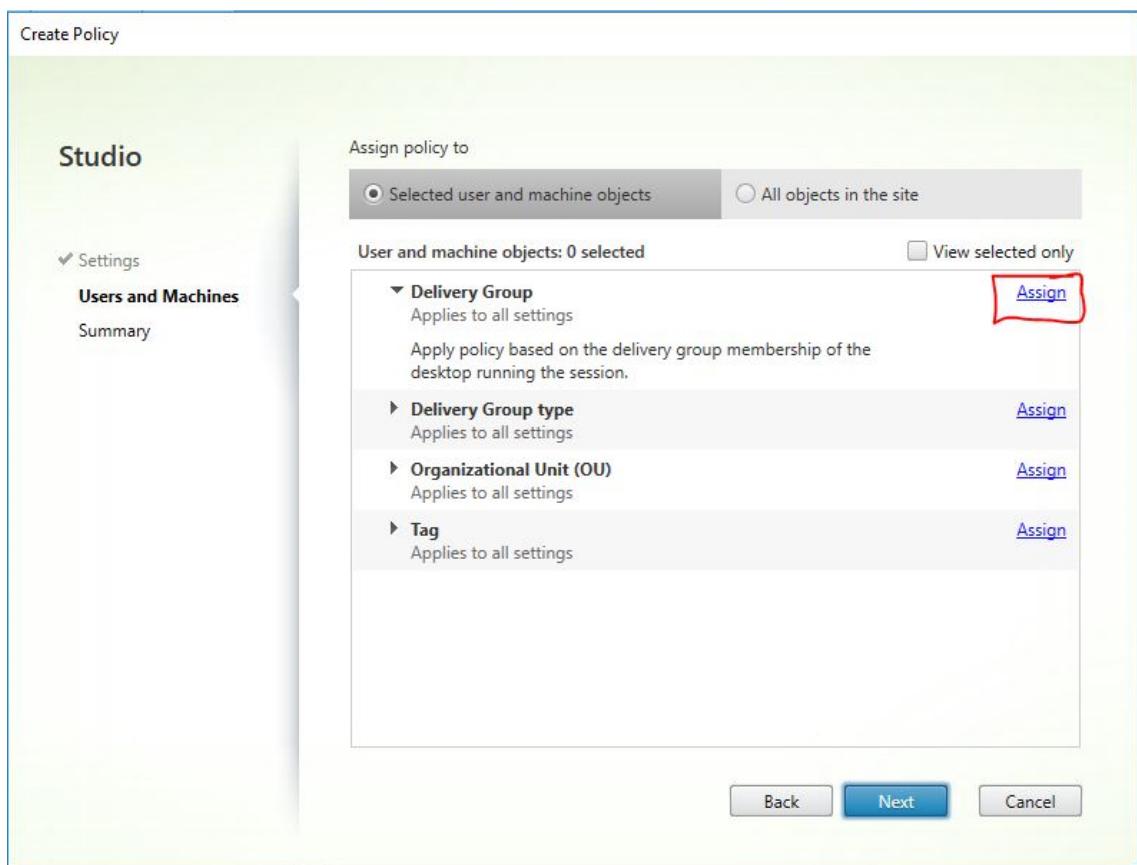
7. The **Edit Settings** window appears.

8. Optional: Change the default value of '**0**' to the **maximum size (in GB)** that the user layer can grow. Click **OK**.

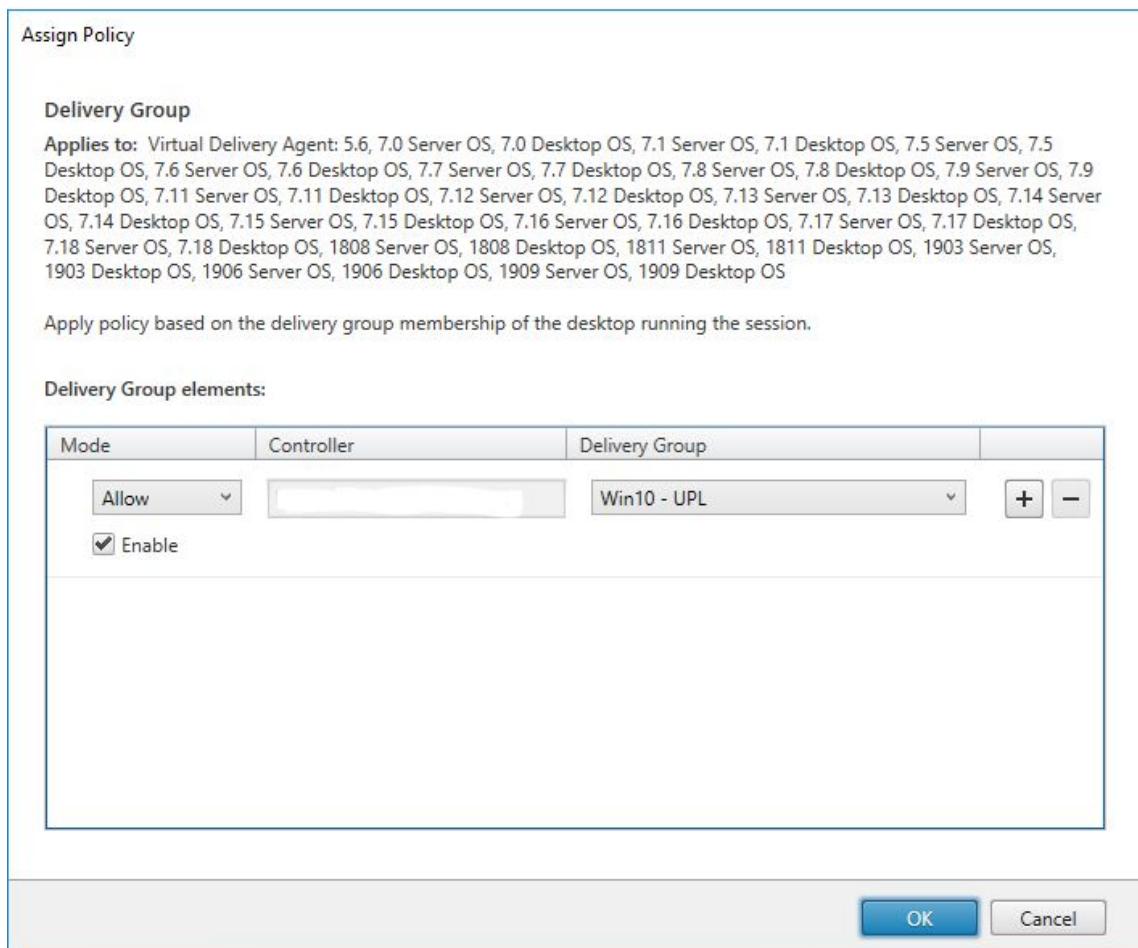
Note:

If you keep the default value, the maximum user layer size is 10 GB.

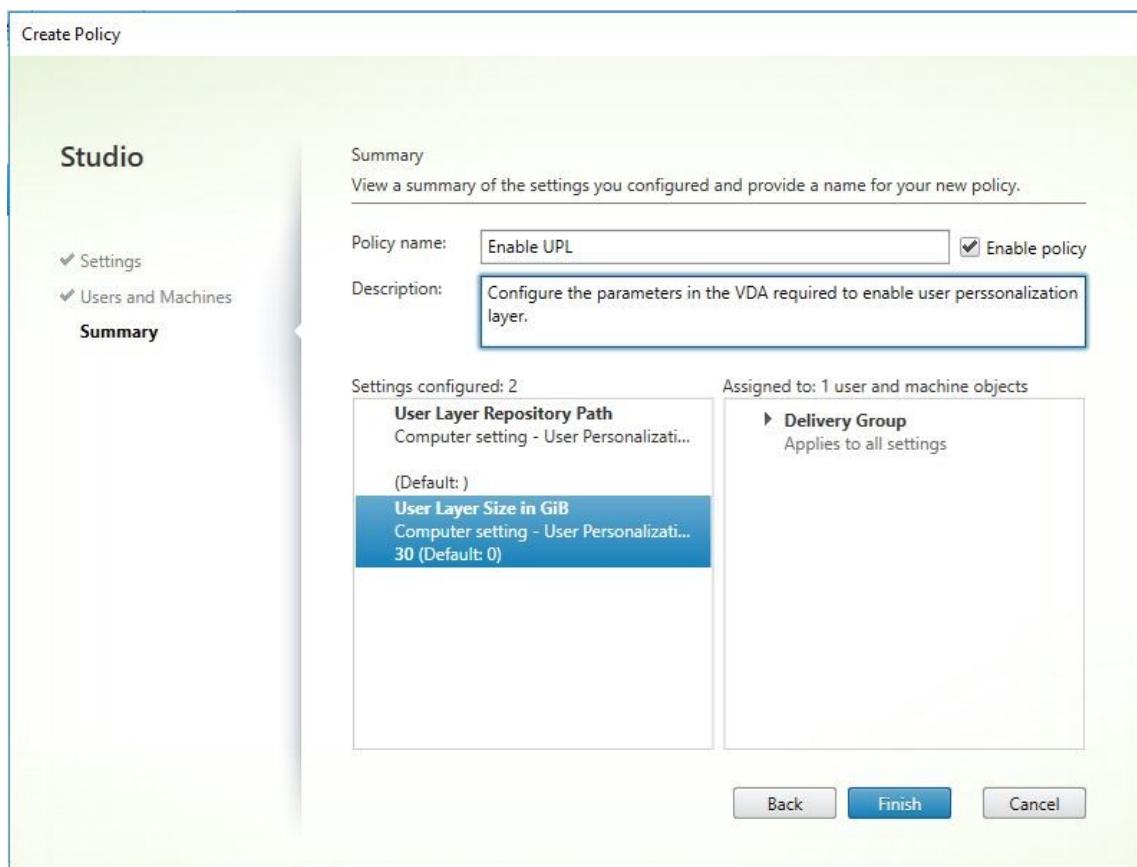
9. Click **Next** to configure Users and Machines. Click the **Delivery Group Assign** link highlighted in this image:



10. In the Delivery Group menu, select the delivery group created in the previous section. Click OK.



11. Enter a name for the policy. Click the checkbox to enable the policy, and click **Finish**.



Move existing user layers to a new storage location

Copy each user layer storage location to its new location:

1. Ensure that the user layer is not in use.

If a user logs in before you move their user layer, a new user layer is created. No data is lost, but if it happens, be sure to:

- Move the newly created user layer to the new directory.
- Preserve the user's ACLs.

2. Browse to the directory containing the user layer VHDX or VHD file.
3. Using the following command, copy each of the user layer VHDX or VHD files from the previous location to the new one.

```
1 xcopy Domain1\User1 Domain1_User1\ /O /X /E /H /K
```

1. Verify that all permissions are correct on the following directories, and files within them:

```
1  \\\Root\Engineering\Users
2
3  \\\Root\Engineering\Users\Domain1_User1\...
4
5  \\\Root\Engineering\Users\Domain2_User2\...
```

If you let users create user layers

If you choose to let users create user layers, you must manually clean up the original directories and files from your share.

User layer size

By default, the disk space allowed for an image's user layers is 10 GB per layer.

You can change the default user layer size by:

- Defining a quota for the user layer share
- Setting a Registry override

When using Office 365 user layers, the Outlook layer defaults to 10 GB, but Outlook sets the volume size based on the amount of free disk space. Outlook uses more or less space based on what is available on the layered image. The size reported is based on the layered image.

Order of precedence

When deploying user layers, the appliance uses the following order of precedence to determine the user layer size:

- Disk quota on user layer size set using either:
 - Microsoft File Server Resource Manager (FSRM)
 - Microsoft Quota Manager
- A Registry override for user layer disks:
(HKLM\SOFTWARE\Unidesk\ULayer\DefaultUserLayerSizeInGb)
- Default user layer size (10 GB)

Change the user layer size

Increasing to assigned user layer size will take effect the next time the user logs in. Decreasing to assigned user layer size does not affect existing user layers.

Define a disk quota for user layer disks

You can set a *hard* quota on the user layer disk size using either of Microsoft's quota tools:

- File Server Resource Manager (FSRM)
- Quota Manager

The quota must be set on the user layer directory, named **Users**.

Note:

Changing the quota (increasing or decreasing) only impacts new user layers. The maximum size of existing user layers was previously set and remain unchanged when the quota is updated.

Set maximum size registry overrides

You can override the default user layer max size using the registry on managed machines. The following registry keys are optional. You do not have to configure these keys for normal operation. If you need one of these keys, add it manually using a layer or a GPO/GPP.

Registry Root: HKLM\Software\Unidesk\Ulayer

Key	Type	Default Value	Description
UseQuotaIfAvailable	String	True; False	True to enable discovery and use of quotas. False to disable.
DefaultUserLayerSizeInGbDWord		User defined	The size of the user layer in GB (for example 5, 10, 23, ...) When not specified the default is 10.

Key	Type	Default Value	Description
QuotaQuerySleepMS	DWord	User defined	The number of milliseconds to wait after creating the directory for the user layer before checking to see if it has a quota. Some quota systems take time to apply the quota to the new directory, for example, FSRM. The default quota is 1000.

Repair a user layer

The user layer repair feature lets you remove an app and its files from a person's user layer. You can use this feature after delivering an app to users who have already installed the app locally on their user layer. The repair feature removes conflicting files whether you deliver the new app layer as part of the base image or as an elastic layer.

- **Example 1:** You create an app layer that includes the file, you.txt, and provide the app layer elastically to users. When a user changes the file, the changes are stored in their user layer. If their changes break the app, or the file was corrupted, the user layer repair feature lets you clean up the problem file by removing it from the user layer. The user then sees the file that is provided elastically as part of the app layer.
- **Example 2:** A user deletes an app that is assigned to them elastically. Because the user layer takes precedence, once the user's local copy of the app is deleted, the user no longer sees that version of the app. The user sees the app layer that is assigned as an elastic layer.
- **Example 3:** A user installs an application locally, and sometime later the administrator creates an app layer for the same application. The user layer repair feature removes any conflicting files installed by the app from the user layer so that the user then sees the version supplied in the app layer.

How user layer repair works

The appliance generates user layer repair JSON files that you can use to clean up or restore the user layer. You manually copy the JSON files to the user layers that need the repair.

If the repair upload folders do not exist on the network share they are created automatically. The appliance writes the repair JSON files to the following directories on the File share:

- 1 <StorageLocationShare>\Unidesk\Layers\App\Repair\
- 2 <StorageLocationShare>\Unidesk\Layers\App\PackageAppRules

The **Repair** directory contains the JSON files for each version of each layer that the appliance knows about. Whenever you finalize a new app layer or a version of it, the appliance generates and uploads the repair files.

The repair files for each layer include:

- 1 UserLayerRepair_LayerIdInDecimal_RevisionIdInDecimal.json
- 2 UserLayerRepair_<layer id>_<layer version>.json

To see the layer ID in the console, click the **Layer** tab, select the layer, and click the **i** icon. The layer ID is displayed along with other layer details.

The **PackageAppRules** directory contains the package app rules for each version of a user layer.

How long does it take to repair a layer?

The repair process time varies based on how large the layer is and how many objects need to be deleted.

A repair of a layer that needs to be mounted but has no actual operations to perform adds about 5 seconds to the login process. Login time is reduced to 2 seconds when the app layer is included in the image.

The time varies depending on the operations. For a typical app layer, it is less than 10 seconds, so 12–15 in total.

Repair a user layer

To repair the user layer for a user:

1. Identify the version of an app layer that must be repaired.
2. Locate the pre-generated **UserLayerRepair** files. If the files have not been generated, contact App Layering Support. Your Support engineer can generate the repair files manually for you.
3. Copy the user layer repair files directly to the user's VHDX or VHD location. The next time the user logs in, a repair operation occurs.

If the user layer repair task is completed, the **UserLayerRepair.JSON** file is removed.

Note:

If a JSON Rules file exists on the share and has been modified by the user, it is not overwritten. This allows users to modify those files as desired.

Log files for user layer repairs

The log file, ulayersvc.log, contains the output of the user layer repair executable.

```
1 C:\ProgramData\Unidesk\Logs\ulayersvc.log
```

Any changes made during the cleanup are logged there, along with any other changes that the service logs.

What happens if a repair fails?

In the event of a failure, the user receives a message that the repair has failed and that they must contact their admin. You can configure the message in the same place as the other storage location messages.

A repair failure can occur in the following cases:

- Bad UserLayerRepair.json formatting (unlikely, since the JSON files are generated).
- Cannot find a specified app layer's .VHD or in-image package_app_rules file.
- Failure to attach an app layer's VHD file.
- Unexpected (random) exceptions interrupting the repair process.

If any of these issues occur, the UserLayerRepair.JSON file is NOT removed, and processing of remaining JSON files stops.

To identify the exact reasons for the failure, review the user's ulayersvc.log file. You can then allow the repair to run again on subsequent logons. Assuming the cause of the failure is resolved, the repair eventually succeeds, and the UserLayerRepair.JSON files are removed.

Update layer

January 5, 2026

The steps for updating the software in an OS, platform, or app layer are virtually the same. You add a version to the layer, install the upgrade or patch on the packaging machine, verify, and then finalize the layer. Once updated, you deploy the new layer version, which varies based on the type of layer.

The platform layer is the highest priority layer and critical for the deployment of images, especially with regards to devices, such as your networks. Whenever you update the infrastructure software, you must add a new version to the platform layer.

You add a version to the platform layer using the new OS layer as the base. Once the packaging machine has started, shut down the machine for finalization. The platform layer gathers the critical components from the new OS layer version, and updates them in the platform so that they match the OS version.

Add a version to the layer

For example, to add a version to an OS layer:

1. In the Citrix App Layering™ Management Console, select **Layers > OS Layers**
2. Select an OS layer and click **Add Version** on the **Version Information** tab.
3. In Version Details:
 - a) For **Base Version**, select the version to use as the base for the new layer version. The default is the latest version.
 - b) Enter a name for the **New Version**. This can be the OS version or other identifying information.
4. Select a **Connector configuration** for the hypervisor where you create your layer.
5. Enter a file name for the Packaging Disk, and select the disk format to use if you are using the appliance's File Share, instead of a connector configuration. This disk is used for the packaging machine (the virtual machine) where you install the application.
6. Verify your settings and click **Add Version**. This runs a task to create an OS version. When the task completes, it shows a status of **Action Required**. When you select the task and click **View Details**, the following text displays:

“The Packaging Disk has been published. The virtual machine ‘<...>’ can be found in folder ‘<...>’ in data center ‘<...>’. Power on this virtual machine to install your application. When the installation is complete, power off the virtual machine before clicking **Finalize** on the Action bar.”

Next, you can deploy a packaging machine for this OS layer version.

Deploy a packaging machine to your hypervisor

The packaging machine is a virtual machine where you install the updates or applications to include in the layer. The packaging machine is a temporary virtual machine that is deleted once the OS layer has been finalized.

The task description contains directions to navigate to the location in your hypervisor where the packaging machine for this layer has been created.

1. To create the packaging machine in your hypervisor, begin with the expanded packaging disk task shown in step 2.
2. Log into your hypervisor client.
3. Back in the management console, use the instructions in the expanded packaging disk task to navigate to the packaging machine.

Install the OS update

1. Remote log into the packaging machine. Be sure to log in to the User account you used to create the OS.
2. Install any updates or applications you want to include in the new OS layer version, such as Windows Updates or antivirus applications.
3. If an application installation requires a system restart, restart it manually. The packaging machine does not restart automatically.
4. Make sure the packaging machine is in the state you want it to be for the user:
 - a) If the applications you install require any post-installation setup or application registration, complete those steps now.
 - b) Remove any settings, configurations, files, mapped drives, or applications that you do not want to include on the packaging machine.

Next, you shut down the packaging machine and verify that the layer is ready to finalize.

Note:

When you upgrade Windows 10 from one major version to another (1703 to 1709, for instance), the previous Windows installation is left in a C:\Windows.old folder. In App Layering, you must not delete this folder. Our software needs to copy our drivers and other files from Windows.old to Windows once the upgrade is completely finished. We will clean up Windows.old when you Finalize the OS layer.

Verify the Layer and shut down the packaging machine

The next step is to verify that the layer is ready to be finalized. To be ready for finalization, any required post-installation processing, for example, a reboot or a Microsoft [ngen](#) process, must complete.

To verify that any outstanding processes are complete, you can run the Shutdown For Finalize tool (icon below), which appears on the Packaging Machine's desktop.

To use the Shutdown For Finalize tool:

1. If you are not logged into the packaging machine, remote login as the user who created the machine.

2. Double-click the Shutdown For Finalize icon. A command line window displays messages detailing the layer verification process.
3. If there is an outstanding operation that must be completed before the layer can be finalized, you are prompted to complete it. For example, if a Microsoft `ngen` operation must complete, you can try to expedite the `ngen` operation, as detailed below.
4. Once any pending operations are complete, double-click the Shutdown For Finalize icon again. This shuts down the Packaging Machine.

The layer is ready to finalize.

If the connector configuration you selected is set to **Offload Compositing**, the layer is automatically finalized. Otherwise, the next step is to finalize the layer manually, as described in the next procedure.

Layer integrity messages

The following layer integrity messages tell you what queued operations must be completed before the layer is ready to finalize:

- A RunOnce script is outstanding - please check and reboot the Packaging Machine.
- A post-installation reboot is pending - please check and reboot the packaging machine.
- A Microsoft `ngen` operation is in progress in the background. - An MSI install operation is in progress - please check the packaging machine.
- A reboot is pending to update drivers on the boot disk - please check and reboot the packaging machine.
- A Microsoft `ngen` operation is needed.
- Software Center Client is configured to run, but the SMSCFG.INI is still present. To learn more about deploying SCCM in a virtual environment, see the Microsoft TechNet article, [Implementing SCCM in a XenDesktop VDI environment] (<https://social.technet.microsoft.com/wiki/contents/articles/23923.implementing-sccm-in-a-xendesktop-vdi-environment.aspx>).

For details about what the layer integrity messages mean and how to debug them, see [Debugging Layer Integrity Problems in Citrix App Layering 4.x and later](#).

You cannot bypass layer integrity messages by shutting down the machine. The App Layering software stops and returns you to the packaging machine until the processes have completed.

If a Microsoft `ngen` operation is in progress, you may be able to expedite it, as described in the next section.

Expedite Microsoft Ngen.exe operations, if necessary

Once all software updates have been installed, you must allow `Ngen.exe` to essentially recompile .NET byte code into native images and construct the registry entries to manage them.

`Ngen.exe` is the Microsoft Native Image Generator, which is part of the .NET system. Windows determines when to run `Ngen.exe` based on what software is being installed and what Windows detects in the configuration.

Important:

When `Ngen.exe` is running, you must let it complete. An interrupted `Ngen.exe` operation can leave you with non-functioning .NET assemblies or other problems in the .NET system.

Normally, `Ngen.exe` is a background operation that pauses when there is foreground activity. If you want to expedite an `Ngen.exe` operation, you can bring the task into the foreground to complete it as quickly as possible.

To bring the task into the foreground:

1. Open a command prompt as Administrator.
2. Go to the `Microsoft.NET\Framework` directory for the version currently in use:

```
cd C:\Windows\Microsoft.NET\FrameworkNN\vX.X.XXXX
```

3. **If using .NET Framework 3 or later**, enter the following `Ngen.exe` command to run all queued items. This command processes queued component installs before building assemblies.

```
ngen eqi 3
```

The `Ngen.exe` task moves to the foreground in the command prompt, and lists the assemblies being compiled. It is OK if you see compilation messages.

Ensure that all `Ngen.exe` processes have run to completion. You can use the **Task Manager** to see if an instance of `MSCORSVW.EXE` is running. If it is, allow it to complete, or run `ngen eqi 3`.

Caution:

Do not reboot to stop the task. Allow the task to complete!

4. **If using .NET Framework 2 or earlier**, enter the following `Ngen.exe` command to run the queued items.

```
ngen update /force
```

This brings the ngen task to the foreground in the command prompt, and lists the assemblies being compiled.

Note:

It's okay if you see **compilation failed** messages.

Look in the Task Manager to see if an instance of MSWORD.EXE is running. If it is, you must allow it to complete, or rerun 'ngen update /force'. Do not reboot to stop the task. Allow it to complete.

Check the status of an **Ngen.exe** operation by opening a command prompt as Administrator and running this command: **ngen queue status**

Caution:

Do not reboot to stop the task. Allow the task to complete!

5. When all operations are complete, shut down the virtual machine using the **Shutdown For Finalize** shortcut available on your desktop.

Finalize the layer manually

Once the packaging machine is created and any apps or updates installed, you can finalize the layer.

Note: When you finalize a new version of an OS layer, the system deletes the packaging machine so as not to incur more costs.

When a layer is ready to finalize:

1. Return to the management console.
2. In the Layers module, select the layer.
3. Select **Finalize** in the Action bar.
4. Monitor the Taskbar to verify that the action completes successfully and that the layer is deployable.

Export and import layers

January 5, 2026

The layer *Export and import* feature lets you export layers from your App Layering appliance, and import the layers into an appliance installed in another location in your environment. For example, you

can use this feature to move layers from an appliance in an on-premises environment to an appliance in a cloud environment. Or, from a proof-of-concept environment to production.

The appliance runs an analysis on each layer it imports to determine its suitability for elastic layering. The results of the analysis are included in the layer details.

Before you start

Before exporting or importing layers, please meet the following requirements, and review related considerations.

Requirements

To export and import layers you need the following:

- The appliance that contains the layers you want to export.
- The appliance to which you want to move the layers.
- A configured SMB Network File Share that is:
 - Reachable by both appliances.
 - Has enough space for all exported layers, plus the meta data and icons for the layers.
- Information required:
 - The path, user name, and password for the SMB Network File Share to which you want to export the layers.
 - Administrator name and password for the management console.

Considerations

Consider the following points when exporting and importing layers.

If file names on the File Share include special characters Since the File Share does not support some special characters, layer names with special characters are temporarily changed to underscores (_) on the File Share. All unsupported characters are changed.

For example, a layer named “Notepad++” becomes “Notepad__” when exported. When the layers are imported onto the new appliance, the original name is restored, and it again appears as “Notepad++”

Run no more than one import or export at a time Only one import or export should be run at a time.

The required OS layer must be included in your layer export You can copy exported layers to a different File Share for import, but be sure to move the required OS layer file along with the other layer files. The appliance imports the OS layer first, because the OS layer is required to import other layers.

Naming layers for export All layers are exported to the following directory on the File Share:

\network-file-share\Unidesk\Exported Layers\

The exported file names reflect the Layer Name, Layer Version Name, and Layer GUIDs.

Selecting layers When selecting layers, the icons not only indicate which layers are selected, they also indicate whether a folder contains layers that have already been exported or imported, depending on which operation you are engaged in.

When exporting layers:

- If an icon is partially selected, it means that the folder includes some layers that are not selected.
- If an icon is grayed out, it means that the layer cannot be selected for export, most likely because the layer has already been exported. Hover over the layer for a message about why it cannot be selected.

When importing layers:

- If an icon is partially selected, it means that the folder includes some layers that are not selected.
- If an icon is grayed out, it means that the layer cannot be selected for import, most likely because the layer has already been imported. Hover over the layer for a message about why it cannot be selected.

Searching layers The **Search** box allows you to find all layers and versions containing the search text. Searches are *not* case sensitive.

Remember, when you select a folder, the UI only selects the subset of layers that are visible. When a folder you select has a *partially* selected icon, it means that some of the selected layers have already been exported or imported. You can view those layers, which are normally hidden, by selecting the **Show versions which cannot be selected** check box above the list.

Search speed affected by the number of subdirectory levels When exporting a layer, if your sub-directories are several levels deep, our software scans every level under that directory, and therefore takes longer to display directories.

Deleting exported layers You can delete an exported layer from the database only when it is not included in an image template.

Export layers

1. In the App Layering management console, select **Layers > Export**.
2. Enter the path of the File Share where you want to export the layers. The App Layering software appends the following path to the Universal Naming Convention (UNC) you specify:
\Unidesk\Exported Layers
3. Enter the **Username** and **Password** for an administrator who has Full Access to the File.
4. Click **Connect**. If necessary, adjust the credentials until the test is successful. You cannot proceed until there is a connection to the File Share.
5. Choose the layers for export:
 - a) Under **Version Selection**, click **Edit Selection** to reveal folders of your OS layer versions, Platform layers, and App layers.
 - b) If you have already exported to the selected location, you can click the **Show versions which cannot be selected** check box to reveal the layers previously exported to this location.
 - c) Select the layers and versions to export, then click **Save**.
6. Click **Confirm and Complete** to verify the space available does not exceed the estimated file size. When space is sufficient, the **Export Layers** button becomes available. You can optionally enter a comment. A *Layer Export* task is created where you can track progress.
7. If you cancel the export before the OS layer is fully exported, be sure to re-export *all* of the layers again, including the OS layer.
8. If you create more layers after the initial export, run another export to the same location. Only the new layers are exported.

Warning

Do *not* attempt to edit or rename any of the exported files.

Import layers

When importing layers from one appliance into another, if two layers have the same name even though the contents of the layer are different, the layer that is imported has a “1” appended to the name. If other layers with the same name are imported, the “1” is incremented.

Note:

To import an app or platform layer, the OS layer must exist on the appliance, or be imported at the same time. You can import several layers at a time, and the OS layer is always processed before any dependent layers.

1. Log into the App Layering management console and select **Layers > Import**.

Note:

You can deselect individual layers in the folder, as explained in step #4 below.

2. Enter the path to the File Share where you exported layers. The following is appended to the URL you specify.

\Unidesk\Exported Layers

3. Enter the **Username** and **Password** for an administrator who has *Full Access* to the File Share.
4. Click **Connect**. If necessary, adjust the credentials until the test is successful.

The system compares the contents of the appliance with that of the selected File Share, and prepares to import the layers that have not yet been imported.

5. Choose the layers for import:

- a) Click **Edit Selection** to expand the OS layer that includes the layers you want to import, and select one of the subfolders to import. This selects every layer and version available for import in the folder.
- b) If you want to see the layers that have already been imported from this location, make them visible by clicking the check box **Show versions which cannot be selected**.
- c) If one of the folders includes layers that you do *not* want to import, deselect each of those layers.

6. Verify the layers to be imported, then click **Confirm and Complete**.

- Verify the layers queued up for import. Only layers that have *not yet been imported* from the File Share are listed.
- Verify that there is enough space on the appliance’s local storage for the layers. The system does not allow the import to proceed until there is enough space for the layers.

Important

If you cancel an OS layer import, all layer imports that rely on the OS layer are canceled.

7. When all settings are valid, click **Import Layers**. An Import task is created where you can track progress.

Once the layers are imported, an “Elastic Fit Analysis” is run on the layers, allowing you to see which layers can be elastically assigned.

Exclude files from layers (Advanced feature)

January 5, 2026

You can exclude specific files and folders from a composited layer to prevent files from persisting on a user’s desktop. For example, you can exclude antivirus software files and folders that must not persist for a desktop from one login to the next.

The exclusions you define are applied to a composited layer, once that is part of a published image. This feature isn’t enforced on a packaging machine, only on a published image where the layers have been composited. That means that you define the exclusions while creating the layer, include the layer in the image template, and then publish the image.

Default exclusions

The Gold Image tool updates maintain a folder of .txt files to introduce and accumulate default exclusions for the App Layer file system. OS layers must be updated with the latest tool versions to ensure the correct and full set of exclusions are in place.

The location for these default exclusions is `C:\Windows\Setup\Scripts\CitrixDefaultExclusions\`. Customers don’t need to do anything with this folder or its contents. Any new exclusions can be removed through an OS layer revision if they cause problems for a customer.

Note:

Future Gold Image tool installations will overwrite local changes made by customers, so reporting issues with any default exclusions is recommended.

This feature complements the user exclusions delivery method and follows the same format, restrictions, and usage as `c:\Program Files\Unidesk\Uniservice\UserExclusions\` files would.

The two new default exclusions files are `FsLogixExclusions.txt` and `GroupPolicyHistoryExclusions.txt`.

Limitation

Excluded files and folders on elastic layers aren't processed. Exclusions can only be processed when present in the image.

Specify files and folders to exclude

In the `C:\Program Files\Unidesk\Uniservice\UserExclusions\` folder, create one or more `.txt` files that specify paths to be excluded.

All valid paths to files and directories are excluded and then read from the image. All changes to those files and directories on the writable layer no longer persist.

If one of the files you create contains an invalid path, processing of that file stops and moves to the next `.txt` file within the `\UserExclusions` folder.

You can also use a `*` character to wildcard one directory for exclusion. For example, `C:\Users*\AppData\Local\Temp\`, where `*` indicates any user name. In this case, any user name that matches the rest of the path fits the exclusion rule, allowing the administrator to skip the user's `\Temp` directory for all users who use that image.

For each exclusion rule, you can only wildcard one directory (use one `*`) in a single path. You can't exclude multiple directories with one `*`. For example, using the rule `C:\Top*\Bottom\` excludes the files in directories `C:\Top\First\Bottom\`, `C:\Top\Second\Bottom\`, and so on. But files in the directory `C:\Top\First\Second\Bottom\` aren't excluded, because there are two directories between `\Top\` and `\Bottom\` rather than one.

There's no limit to the number of exclusion rules that you can set containing a wildcard (`*`).

Examples

Exclude a file:

```
1 c:\test\test.txt
```

Exclude a directory:

```
1 c:\test\
```

Restrictions

The following restrictions apply to exclusions.

Directory name

- Begin the path with C:\
- End with a Backslash (\)

Exclusions These top-level directories can't be excluded:

- C:\
- C:\Program Files\
- C:\Program Files (x86)\
- C:\ProgramData\
- C:\Windows\
- C:\Users\

The following characters and expressions aren't allowed in exclusions:

- No question marks (?)
- No regular expressions (no %x%)
- No forward slash (/)
- No network (\\)
- No path to a different directory (\..\)
- No quotation marks (")
- No colon (:) after C:\

Log

Log messages are available in:

```
1 C:\Program Files\Unidesk\Uniservice\Log\Log0.txt
```

Messages written to the log:

- User exclusion added: Includes the details about the file or directory.
- Failed to add user exclusion: Includes details about the unsupported exclusions.

Publish layered images

January 5, 2026

The Citrix App Layering™ service lets you publish *layered images* as disks compatible with your platform. You can use a layered image to provision virtual machines, as you would with any other image.

About layered images

Layered Images are bootable images composed from an OS layer, a Platform layer, and any number of App layers.

You publish layered images from an image template. The image template lets you specify the layers to include in the layered images. You can also specify the following:

- The connector configuration to use to access a location in your environment.
- Whether to enable Elastic layers for delivering applications to users when they log in.
- Whether to enable User layers, which persist users' application data and settings.

Layered images for provisioning systems

The way you specify which layers to include in a layered image is by saving the combination of layers you want for a particular group of users in an image template. You then use this template to publish a layered image to your chosen platform.

When you need to update the layered image, you simply edit the image template to add or remove layer assignments and publish a new version of the image.

Create or clone an image template

January 5, 2026

An image template stores the list of layers to include in the layered images that you publish. From a single template, you can publish as many layered images as you need to provision systems in a particular location.

Once you create an image template from scratch, you can clone the template to quickly create a set of templates that have the same settings.

When you upgrade the software in an App or Platform layer, you then update your image templates to use the new layer version. To deliver the new layer version to users, you can republish your layered images and use the updated images to provision your systems.

Requirements

To create an image template, use

- OS layer
- Platform layer (Optional)

Important:

If using the Platform layer, it must have the same hardware settings as the OS layer. You choose these settings when deploying the virtual machine for the OS and Platform layers.

- App layers (Optional)
 - You don't need to include App layers when you create an image template.
 - You can add App layers to an image template and then republish your layered images at any time.

You can create an image template without App layers. This is useful for testing your OS layer before using it to create App layers.

Create an image template from scratch

To create an image template:

1. In the App Layering management console, select the **Images** module, then click **Create Template**.
2. Enter a **Name** for the template and notes in the **Description** field (optional), so you can identify the template when choosing one for publishing a layered image.
3. Select one of the **Available OS Layers**. If there is more than one layer version, the most recent version is selected by default. You can select an older version from the drop-down menu.
4. (Optional) Select the **App Layers > Edit Selection** that you want to include in the layered images that you publish using this template.
5. (Optional) Select a Platform layer with the tools and hardware settings for the required image target.

6. Choose a **Connector Configuration** for the platform where you are creating this layer. If the configuration you need isn't listed, you must create the connector configuration from the **Connectors** page.
7. Edit the following fields, as needed:
 - **Layered Image Disk File name (Auto-populated)**: The name for the layered image disk is auto-populated from the template name.
 - **Defragmented Layered Image Disk**: When enabled, the layered image disk will be defragmented. This option is available for Offload Compositing connectors only.
 - **Layered Image Partition Size**: The default disk size of 100 GB is recommended.
 - **Layered Image Disk Format**: The default disk format is VHD, but you can also select VMDK or QCOW2. If you are publishing to the appliance's File Share instead of using a connector configuration, this setting allows you to choose a disk format compatible with the environment to which you are copying the disk.
 - **Sysprep**: The options available depend upon the hypervisor or Provisioning Service specified in your connector configuration:
 - Azure, Hyper-V, XenServer®, Nutanix, vSphere: Defaults to *Generalized Offline*. (For Azure, this is the only option.)
 - Machine creation, Citrix Provisioning™, View: *Not generalized* is the only option for Machine creation, Citrix Provisioning, and View running on any of the hypervisors.
 - File Share: Defaults to *Not generalized*, if using a File Share instead of a connector configuration.
 - **Elastic Layering**: Select the **Application Layering** option to activate Elastic Layering on this layered image.
 - **User Layer**: When enabled in System Settings, you can select **Full** User layers (Labs), **Office 365** (desktop), or Session Office 365 (server OS) option. Choose the **Full** option to save the settings and data for users independent of specific applications. Choose **Office 365** or **Session Office 365** to save the settings and data for Outlook 365 running on a desktop system or a session host.
8. Select **Confirm and Complete**, and enter any comments you would like for this layer.
9. Click **Create Template** to save your changes, or **Create Template and Publish** to save the template and then publish the layered images.

The new template icon appears in the Images module.

Clone an image template

You can create a copy of an image template by cloning it. Each clone is a stand-alone copy of the original. The audit history shows that the template was cloned “Created (Cloned),” and indicates which

image it was cloned from, “Cloned from *template-name*.”

The first clone is named the same as the original template, with “- Copy” appended to it. Each subsequent clone has “- Copy*N*” appended instead, where *N* is an incrementing sequence number. The sequence number is incremented to the first available number, rather than the number after the last one already in use. The maximum number is 1000.

To clone an image template:

1. In the App Layering management console, select the **Images** module.
2. In the Images module, select an image template, and click **Clone** on the Action bar or popup menu. A copy of the template is created with “- Copy (1)” appended to the name.

You can rename and edit the clone for your purposes.

Next step

You can now use the image template to [Publish layered images from image template](#)

Publish layered images from template

January 5, 2026

Layered images are virtual machines composited from the layers and settings specified in an image template. Using an image template, you can publish as many layered images as you need to a location in your Provisioning Service or hypervisor that you specify in the connector configuration.

When layered images are published, you can use scripts to perform important layer-specific steps. For example, you can activate Microsoft Office, which may need to be done before the virtual machine is used as a master disk for your deployment tools.

The mechanism used to run these scripts can vary, including our own `kmssetup.cmd` functionality, run-once support, or even running the scripts manually. After all of the scripts run or other manual steps are taken and the virtual machine is in the desired state, a guest OS shutdown is initiated either by the scripts or manually. If you use the `kmssetup.cmd` functionality, there is a documented process for initiating a shutdown after all layer scripts and other `kmssetup` functionality are complete.

To publish a layered image:

1. In the Images module, select one or more image templates that you want to publish.
2. On the **Action** menu, select **Publish Layered Image**.

3. On the **Confirm and Complete** page, select **Publish Layered Images**. For each Image Template, this starts a task called **Publishing Layered Image**.
4. Check the taskbar, and when the disk for this image is created, click the link in the task description to advance to the next stage.

The link leads you to a virtual machine whose state of creation depends upon the platform you are publishing to:

- **XenServer®, MS Hyper-V, Nutanix (or provisioning service on it):** The virtual machine is created, but off. Log in as a guest, and let the machine run any scripts specified in the connector configuration.
- **Azure (or Machine creation for Azure):** A window for creating a virtual machine is open, but incomplete. Enter the values required, finish creating the machine, and power it on. The machine runs any scripts specified in the connector configuration.
 - An Azure connector is used specifically for layer creation.
 - A Machine creation for Azure connector is used specifically for template publishing.
- **Google Cloud:** Create a VM from the Google Cloud Image on the Google Cloud console, and power it on as the guest operating system.
- **VMware Cloud:** Create a VM using a VMware Virtual Machine template, and power it on as the guest operating system.

Note:

A **VMware Virtual Machine template**, rather than a standard VM template, is required for the virtual machine's network to work correctly.

5. When the virtual machine is in the desired state, shut it down. The task status changes to **Done**.
6. Use the information in the task description to navigate to the image in your environment.

Possible error message

If you receive the following error and you want to log into the published image to make changes, it must be through the local user and **not** a domain user.

This system was not shut down properly. Please log off immediately and contact your system administrator.

Manage image template

January 5, 2026

Whenever you create a new layer or add a new version to an existing one, you can:

- Update the layers selected in your image template(s).
- Use the template(s) to publish new versions of your layered images.
- Use the new layered images to manually provision your systems.

You can change or delete a template without affecting any previously published layered images, because an image is *not* associated with the template used to create it.

Update image templates with a new layer version

When you add a new version to an app layer or an OS layer, you can quickly identify the image templates that include the layer, and select which templates to update with the new version.

1. In the App Layering management console, select **Layers**, and then select subtab for the type of layer you are updating.
2. Select the Layer you updated, then the new version of the layer you want to assign.
3. Click **Update Assignments**. Image templates that include this layer are listed.
4. Select the image templates to which you want to assign the layer or layer version.
5. Click **Save**.
6. Click **Confirm and Complete**.

Edit any image template setting

When you want to change the settings that you use to publish any of your layered images, you can edit the image template you originally used to publish the layered image(s) and publish a new version of the image(s).

1. In the App Layering management console, select **Images**.
2. Select the template you want to edit, and click **Edit Template**.
3. You can change the **Name**, **Description**, and **Icon** for the Image.
4. Select a different version of your chosen OS layer by expanding the layer and choosing a different one.
5. Add or remove app layers to include in the layered images that you publish using this template. If there is more than one version of a layer, you can choose a different version by expanding the layer and choosing a different one.

6. Change the location to which the Layered Image is published by selecting a different **Connector Configuration**.
7. Change the selected **Platform Layer**, if for example, you are publishing to a different environment.
8. Edit the **Layered Image Disk** details, for example, to enable elastic layering on the image.
9. Click **Confirm and Complete**, and enter comments for this layer.
10. Click **Save Template Changes**, or click **Save Template and Publish** to publish the layered images after saving the template.

Delete an image template

When you no longer need an image template, you can remove it from the management console.

An image template cannot be deleted while it is being used to publish layered images.

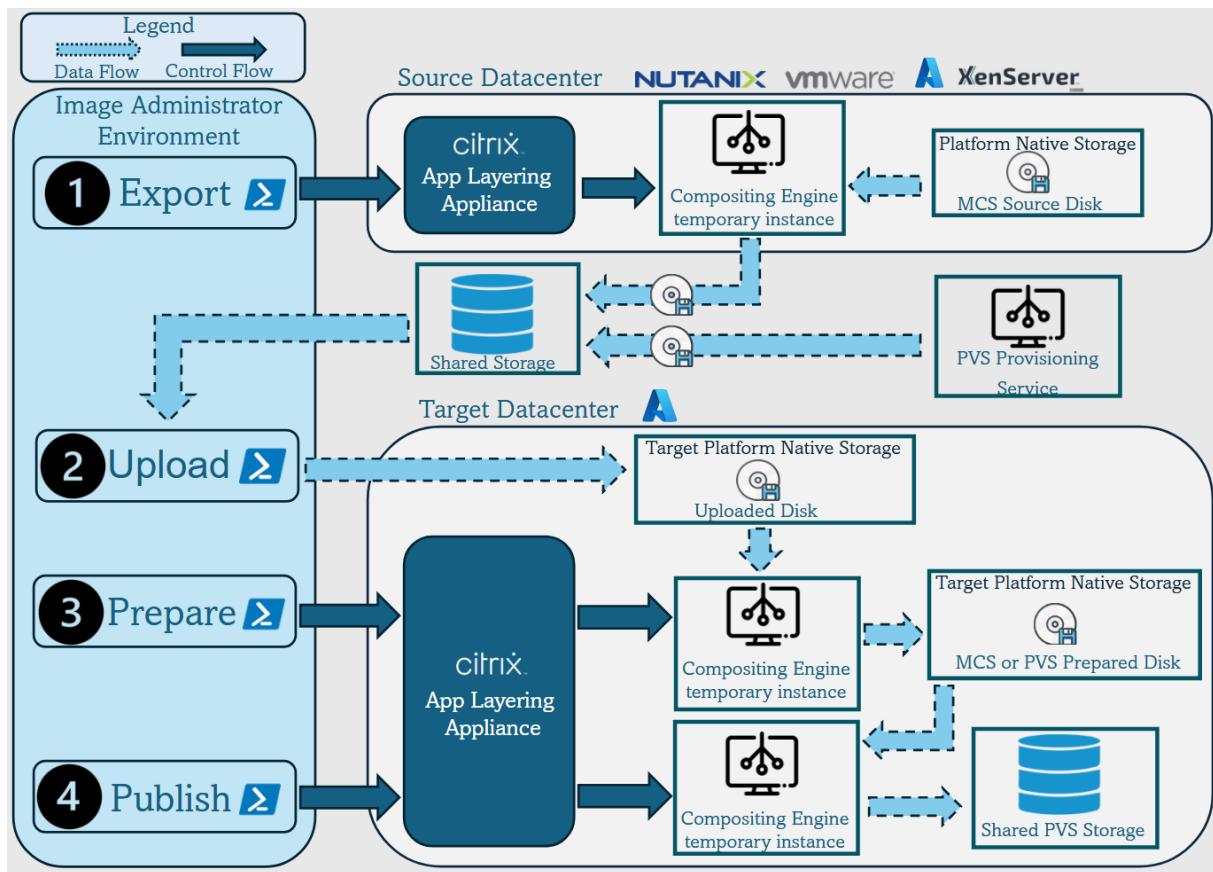
1. In the management console, select **Images**.
2. Select the template you want to delete, and click **Delete Template**.
3. Enter any comments you would like, and click **Delete Template**.

Image Portability

January 6, 2026

Overview

The Image Portability feature uses a multi-phase workflow to export, upload, and prepare a master catalog image from one compute platform to a different one (e.g. from VMware to Azure). Using the [Citrix.AppLayering PowerShell module](#), Image Portability exports the image from the source hypervisor platform, copying the image to a shared storage. Next, the exported image is uploaded to the target platform using the [Citrix.Image.Uploader PowerShell module](#). After being uploaded to the target platform, Image Portability prepares the image to be compatible with your destination platform. Once prepared, the image is ready for use as a new machine catalog with MCS, or to be published to a shared PVS storage.



These high-level workflows are based on the image's source and target provisioning configuration (Machine Creation or Citrix Provisioning™). The chosen workflow determines which Image Portability tasks are required.

Refer to the following table to understand which tasks are required for each of the supported IPS workflows.

Workflow (Source to Target)	Export	Upload	Prepare	Publish
MCS to MCS	✓	✓	✓	
PVS to MCS*		✓	✓	
PVS to PVS*		✓	✓	✓
MCS to PVS	✓	✓	✓	✓

*Assumes you have the original image as a Citrix Provisioning vDisk and do not need to export it directly out of the source platform hypervisor.

Note:

The [Citrix Provisioning](#) connector is currently not compatible with Image Portability tasks.

To publish an image for PVS, run the publish task using the [Windows File Share](#) connector configured with a share that corresponds to the PVS server's store location. This will create a VHD(x) containing the image within the PVS store. You can then register the disk with PVS manually, using the PVS console or PowerShell cmdlets.

Components

Image Portability includes the following components:

- App Layering appliance
- Compositing Engine VM
- PowerShell modules
 - Citrix.AppLayering
 - Citrix.Image.Uploader

App Layering appliance

The App Layering appliance runs in your environment (both on-premises and in your public cloud subscription) and acts as a controller for individual tasks. It creates and manages the Compositing Engine VMs. The appliance controlling IPS tasks must have bi-directional connectivity with the Compositing Engine (CE) in the target platform (e.g., an appliance in VMware needs to connect to an Azure CE, and the Azure CE needs to connect back to the ELM, typically via ExpressRoute). You will need to set up an [IPS-compatible connector](#) within the App Layering appliance to perform IPS tasks.

Note:

A user must be granted the 'Administrator' or 'Image Porter' role to use this feature. Refer to [Assign roles](#) for more information.

Compositing Engine

The Compositing Engine is the workhorse of the Image Portability feature. The Compositing Engine (CE) is a single VM created at the start of an Image Portability export or prepare task. These VMs are created in the environment defined by the connector configuration provided for the task. For example, when exporting a disk from vSphere, the CE is created in the vSphere datacenter. Likewise, when running a prepare task in Azure, the CE is created in Azure. The App Layering appliance attaches a copy or read-only version of your disk to the CE VM, and the required manipulations to the disk are

performed without modifying the original disk. Upon completion of the prepare or export task, the CE VM and all of its components are deleted.

PowerShell modules

Two Powershell modules are required to use the Image Portability feature:

- Citrix.AppLayering
- Citrix.Image.Uploader

Refer to the [Requirements documentation](#) for environment compatibility prior to installing both modules. For information on installation/uninstallation, refer to the [Powershell modules documentation](#).

Citrix.AppLayering The Citrix.AppLayering PowerShell module provides several cmdlets to perform the following:

- Connect and Disconnect to an App Layering appliance
- List existing [connector configurations](#)
- Monitor, wait for, and stop tasks
- Get details of the items associated with a task
- Start Export, Prepare, and Publish tasks
- Start Diagnostics tasks

For more information on the cmdlets and examples on how to use them, refer to the [developer documentation](#).

Citrix.Image.Uploader The Citrix.Image.Uploader Powershell module provides functionality for uploading images to the destination platform. For details and examples on how to use the module to automate the image upload process, refer to the [developer documentation](#).

Requirements

January 5, 2026

To get started with Image Portability, you must meet the following requirements.

Source Platforms

Image portability supports porting images from the following platforms to one of the supported destination platforms:

- VMware vSphere 8.0 and 9.0
- XenServer 8
- Nutanix AHV (Prism Element)
- Microsoft Azure

Destination Platforms

Image portability supports porting images to the following platforms:

- Microsoft Azure

A Citrix Machine Catalog image

IPS requires using images that have the following tested configurations:

- A supported Windows operating system
 - Windows Server 2016, 2019, 2022H2, or 2025
 - Windows 10 or 11
- Provisioned using Machine Creation Services or Citrix Provisioning
- Citrix Virtual Delivery Agent

App Layering appliance

Image Portability is available starting with version 25.11 of the [App Layering appliance](#). It is recommended that you have an App Layering appliance in each location that hosts images you plan to export or prepare using Image Portability.

Refer to the [Install appliance](#) documentation for additional information on setup.

An SMB (Windows) file share

You need a Windows SMB file share for storage of the output of export and publish tasks. This share must be accessible to the Compositing Engine VM, which is created in the environment where the image being exported resides. While this may be a Citrix Cloud Resource Location, with the new Image

Portability feature (IPS), the file share may also reside in a customer-managed data center or cloud tenant, independent of Citrix Cloud.

Refer to the [Windows File Share](#) connector documentation for additional information on setup.

An IPS-compatible Connector

Both export and prepare tasks require an IPS-compatible connector configured on the App Layering appliance. Refer to Connector documentation for your specific platform for additional information on setup. For example, our [Image Portability for Azure](#) Connector documentation outlines the requirements for configuring a connector needed to start a prepare task in your Azure environment.

A machine for running PowerShell scripts

The machine used to run the PowerShell cmdlets of the [Citrix.AppLayering](#) or [Citrix.Image.Uploader](#) PowerShell module must meet the following requirements:

- Must have Windows PowerShell version 5.1 (PowerShell 6 and above not supported)
- A Windows machine with .Net framework, minimum version 4.6.2
- Can connect to the Microsoft PowerShell Gallery to download the required PowerShell libraries

The machine will also have to meet additional requirements specific to the module being used:

Citrix.AppLayering

- A network connection must exist to the App Layering appliance in the environment where the image(s) reside

Citrix.Image.Uploader

- A fast network connection to the SMB file share
 - Can be the same machine that hosts the file share
- A fast network connection to the destination platform

Additional considerations

Depending on your version of Windows, you might need to disable TLS 1.0/1.1 support. Refer to the [Microsoft PowerShell Gallery TLS support documentation](#) for more information.

By default, PowerShell does not automatically authenticate through a proxy server. Make sure you've configured your PowerShell session to use your proxy server, per Microsoft, and your proxy vendor best practices.

Required platform permissions

Refer to your specific [Connector configuration](#) for additional requirements.

Networking

The App Layering appliance creates a worker VM called the compositing engine (CE) to perform image operations. Both the App Layering appliance and CE must be able to communicate with each other over port 443.

For tasks that involve files on an SMB share (for example, export tasks), the CE must be on a network with connectivity to the SMB share.

For prepare tasks, the operating system contained in the image is booted (on the CE) to do specialization and other tasks. If the image contains management or security agents that phone-in to a control server, these processes can interfere with the preparation process.

If the domain unjoin option is specified, network connectivity can affect the results. If the compositing engine VM can reach the Active Directory domain controller over the network, unjoin removes the computer account from the domain. This breaks the domain membership for the source VM from which the image was extracted.

It is recommended to isolate the network provided for the operation of other network resources. This can be done by selecting 'Network isolation' in the connector configuration. For an example, refer to our [Image Portability for Azure](#) documentation.

PowerShell Modules

January 6, 2026

Image Portability uses two PowerShell modules as part of its workflow:

- [Citrix.AppLayering](#)
- [Citrix.Image.Uploader](#)

Refer to the [Requirements](#) documentation for environment compatibility prior to installing both modules.

Citrix.AppLayering

The Citrix.AppLayering PowerShell module provides several cmdlets to perform the following:

- Connect and Disconnect to an App Layering appliance
- List existing [connector configurations](#)
- Monitor, wait for, and stop tasks
- Get details of the items associated with a task
- Start Export, Prepare, and Publish tasks
- Start Diagnostics tasks

For more information on the cmdlets and examples on how to use them, refer to the [developer documentation](#).

Citrix.Image.Uploader

The Citrix.Image.Uploader Powershell module provides functionality for uploading images to the destination platform. For details and examples on how to use the module to automate the image upload process, refer to the [developer documentation](#).

Citrix.AppLayering PowerShell module

January 6, 2026

This document goes over basic steps to install, update, and uninstall the Citrix.AppLayering PowerShell module. When working with Image Portability, this module provides functionality to perform export, prepare, and diagnostic tasks.

For more information on the associated cmdlets and examples on how to use them, refer to the [following documentation](#).

Install libraries and modules

Image Portability uses libraries from the Microsoft PowerShell Gallery to drive portability operations.

Run the following PowerShell command to install the latest version of the Citrix.AppLayering module:

```
1 Install-Module -Name "Citrix.AppLayering" -Scope CurrentUser
```

Run the following PowerShell command to install a specific version of the Citrix.AppLayering module:

```
1 Install-Module -Name "Citrix.AppLayering" -RequiredVersion 1.0.0 -Scope CurrentUser
```

Confirm the module was successfully installed by running the command:

```
1 Get-InstalledModule -Name "Citrix.AppLayering"
```

This command returns an output similar to the following:

Name	Repository	Description
Citrix.AppLayering	PSGallery	Commands to Export and Prepare an image, and additional ways to interact with the App Layering appliance

Update modules to latest version

Run the following command to update the module to the latest version:

```
1 Update-Module -Name "Citrix.AppLayering" -Force
```

Run the following command to update to a specific version of the Citrix.AppLayering module:

```
1 Update-Module -Name "Citrix.AppLayering" -RequiredVersion 1.0.0
```

Note:

If the Citrix.AppLayering PowerShell module is updated to a version that is not supported by your current App Layering appliance some cmdlets may not function correctly. This will require one of the following:

- Install or update your Citrix.AppLayering Powershell module to the version specified by the failed Image Portability task
- [Upgrade](#) your App Layering appliance

Uninstall libraries and modules

Run the following command to uninstall the Citrix.AppLayering module:

```
1 Get-InstalledModule -Name "Citrix.AppLayering" | Uninstall-Module
```

Run the following command to uninstall a specific version of the Citrix.AppLayering module:

```
1 Uninstall-Module -Name "Citrix.AppLayering" -RequiredVersion 1.0.0
```

Basic usage and examples

Once the module is installed, you can view the available commands in the Citrix.AppLayering PowerShell module by running:

```
1 Get-Command -Module "Citrix.AppLayering"
```

To get detailed information about a specific command, including its description and usage examples, use:

```
1 Get-Help <Command-Name> -Full
```

For a complete list of current commands and example use cases, refer to the following [documentation](#).

Citrix.Image.Uploader PowerShell module

January 6, 2026

This document goes over basic steps to install, update, and uninstall the Citrix.Image.Uploader PowerShell module. This module provides functionality to upload a VHD(x) from the SMB file share to a designated target environment.

For more information on the associated cmdlets and examples on how to use them, refer to the [following documentation](#).

Install libraries and modules

Image Portability uses libraries from the Microsoft PowerShell Gallery to drive portability operations.

Run the following PowerShell command to install the latest version of the Citrix.Image.Uploader module:

```
1 Install-Module -Name "Citrix.Image.Uploader" -Scope CurrentUser
```

Run the following PowerShell command to install a specific version of the Citrix.Image.Uploader module:

```
1 Install-Module -Name "Citrix.Image.Uploader" -RequiredVersion 1.0.0 -  
Scope CurrentUser
```

Confirm the module was successfully installed by running the command:

```
1 Get-InstalledModule -Name "Citrix.Image.Uploader"
```

This command returns an output similar to the following:

Name	Repository	Description
Citrix.Image.Uploader	PSGallery	Commands to Upload a VHD(x) to an Azure Storage Account, AWS, or GCP and Get information about a VHD(x)

Update modules to latest version

Run the following command to update the module to the latest version:

```
1 Update-Module -Name "Citrix.Image.Uploader" -Force
```

Run the following command to update to a specific version of the Citrix.Image.Uploader module:

```
1 Update-Module -Name "Citrix.Image.Uploader" -RequiredVersion 1.0.0
```

Install platform-specific third-party components

The Image Portability PowerShell module does not install third-party dependencies. Hence, you can limit installation to only the platforms you're targeting. If you're using one of the following platforms, follow the relevant instructions for the installation of platform dependencies.

Note:

Installing third-party components is only required for use with our Citrix.Image.Uploader PowerShell module. These are **not** required for running export or prepare tasks.

Azure

If you're uploading images to Azure, download and install the [Azure command-line utilities](#), then run these commands to install the required Azure PowerShell modules:

```
1 Install-Module -Name Az.Accounts -Scope CurrentUser -AllowClobber -Force
2
3 Install-Module -Name Az.Compute -Scope CurrentUser -AllowClobber -Force
```

Uninstall libraries and modules

Run the following command to uninstall the module:

```
1 Get-InstalledModule -Name "Citrix.Image.Uploader" | Uninstall-Module
```

Run the following command to uninstall a specific version of the Citrix.Image.Uploader module:

```
1 Uninstall-Module -Name "Citrix.Image.Uploader" -RequiredVersion 1.0.0
```

Note:

Third-party scripts and components aren't automatically removed when uninstalling IPS modules.

Basic usage and examples

Once the module is installed, you can view the available commands in the Citrix.Image.Uploader PowerShell module by running:

```
1 Get-Command -Module "Citrix.Image.Uploader"
```

To get detailed information about a specific command, including its description and usage examples, use:

```
1 Get-Help <Command-Name> -Full
```

For a complete list of current commands and example use cases, refer to the following [documentation](#).

Manage

October 23, 2024

This section explains how to manage the App Layering service, including:

- [System settings](#)

- [Storage](#)
- [Appliance settings](#)
- [App Layering services](#)
- [Users](#)
- [Firewall ports](#)
- [Directory service](#)
- [Migrate Appliance](#)
- [Upgrade](#)

System settings

March 14, 2024

You can specify settings for the following system configuration parameters by clicking the **Edit** button of each option, making your changes, and then clicking the **Save** button.

This section describes each appliance setting.

Monitoring and Storage

The following services run on the App Layering appliance:

- Management service
- Layering service
- BITS server service

For more details, go to [App Layering services](#).

Directory Services

You can configure the appliance to connect to a directory service, for example, Active Directory. When you connect to your directory service, you create one or more Directory Junctions to access specific domains or OUs. The appliance does not modify the directory service you connect to. The software caches the attributes for each directory service entry so that if the connection to the directory service is lost temporarily, the software can use the cached information for management tasks.

For more details, go to [Directory service](#).

User Layer Storage Locations

The appliance's local storage is a layer repository where the appliance creates, composites, and stores layers and layered images. To check the amount of free space in the appliance's local storage, you can see how much disk space is used in the System module of the management appliance.

For more details, go to [Storage](#).

Network File Share

The Network File Share is used to:

- Package layers using the Network File Share, rather than a connector for your hypervisor.
- Publish layered Images to the Network File Share, rather than using a connector for your publishing platform.
- Serve Elastic Layers.
- Upgrade the App Layering software.

Configure Network File Share

To configure the Network File Share, specify the following values:

- SMB File Share Path
- User Name and password

Test Network File Share

Then, test the connection to the file share by clicking **Test SMB File Share**. The test returns a message stating either **Success** or **Failed to mount network file share path**. You can enter a comment describing your changes.

HTTP certificate settings

Displays the currently set security certificate. Use the **Upload and Generate** buttons to upload an existing certificate or to generate a new one. Optionally, enter a comment that describes the changes you made.

Trusted certificates

When you enter the URL for the application, you are automatically redirected to a secure connection. If you specify HTTPS as the protocol in the URL, and the application does not include a security certificate from a Certificate Authority, you are prompted to bypass the security warnings the first time you access the application.

To eliminate the security warning, upload a trusted certificate that you create.

Requirements The requirements for the trusted certificate are:

- It must be a Privacy Enhanced Mail (PEM) certificate.
- It must include both the certificate and the key.
- It does not include a passphrase.

Create a CSR You need to generate a CSR file to give to the Certificate Provider for a certificate request. Since ELM is based on CentOS Linux, OpenSSL is included. Use the [OpenSSL CSR Wizard](#) from DigiCert to generate the required OpenSSL command. You can then use Putty to log on to the console and paste in the **OpenSSL** command, which generates the CSR. Refer to the **OpenSSL** commands in the [OpenSSL Quick Reference Guide](#).

Uploading the certificate

1. Create the PEM certificate.
2. Log in to the App Layering CacheCloud Infrastructure Management utility.
3. Select **Configuration > Upload SSL Certificate** in the left pane.
4. Browse to the self-signed PEM certificate file and click **Upload**.
5. Restart the Management Appliance.

Notification settings

You can configure automatic email notification settings for yourself or other users.

Set up email notifications from the appliance

To set up email notifications, complete the following fields. All fields are required.

1. In the Mail Server field, enter the name of your email server or SMTP relay server.

2. In the Mail Server Port field, enter the number of the port that the email server uses for communication.
3. In the User Name field, enter the user name for the email account you want to use for sending notifications. For example, username@domain.com.
4. In the Password field, enter the password for the email account.
5. In the From field, enter an email address to identify the source of the email message. For example, if you enter myaddress@mycompany.com, the email message displays the following in the From box of the received notification:

App Layering Manager [myaddress@mycompany.com]
6. In the Recipient List box, enter the email addresses that must receive notifications. Use a comma or semicolon to separate the email addresses.
7. Click Test Email Configuration to verify that the settings for the email server and account work correctly. If the test succeeds, the software displays a success message and sends the recipients a confirmation email.
8. Enter a comment, if necessary, and click **Save** to save the email settings. Any comments you enter appear in the Information view Audit History.

Security and retention settings

- Specify the number of minutes of inactivity before the management console logs you out.
- Specify the number of days that the appliance must retain completed Tasks before deleting them.
- Specify the number of days that the appliance must retain audit log files. After that time elapses, the software begins to overwrite the audit log.
- Specify the maximum disk space to use for all logs (in MB) and the number of days that the log files must be retained.
- Optionally, enter a comment that describes the changes you made.

About

This section displays more information about the Enterprise Layer Manager (ELM), such as **ELM Version**, the Hypervisor being used, and where to go for support.

Storage

September 6, 2025

The appliance’s local storage is a *layer repository* where the appliance creates, composites, and stores layers and layered images. To check the amount of free space in the appliance’s local storage, you can see how much disk space is used in the System module of the management appliance.

1. Log into the management console and select **System > Manage Appliance**.
2. In the Services table, the Local Storage for the layering service shows how much space is used and how much is free.

Note

- Disk space is shown in 1024-based Gigabytes, not metric.
- Free space is updated every time a layering service job completes. If you want to make sure the page has been refreshed, click the Refresh icon just above the Manage Appliance subtab.
- When creating a layer or adding a Version to it, extra space is temporarily required to build the Packaging Disk. You can calculate the amount of space needed during layer creation by adding the following layer sizes:
 - The size of the OS layer version you’re using.
 - The size of the writable disk you want for the app layer.
 - The size of any prerequisite layers (if you have any).

Add a disk to locally attached storage

When you install the appliance, it comes equipped with an additional 300 GB data disk that is used as a layer repository. You can expand the appliance’s local storage by adding another disk to it.

After adding a disk to the appliance virtual machine using your hypervisor console, do the following steps:

1. Log into your management console.
2. Select **System > Manage Appliance**.
3. Select **Expand Storage**. The Disk Selection tab is displayed of disks that are attached to the system and are *not* part of the layer repository.
4. Select the check box for each disk that you want to use to expand the layer repository. If a check box is grayed out and a yellow icon with an ! (exclamation point) is displayed, it means that the attached disk is not eligible for use (for example, if the disk is not blank). Once the attached disk is blank and unpartitioned, you will be able to use it to expand the appliance’s local storage.

5. On the Confirm and Complete tab, click **Expand Storage**.
6. According to best practice, once the disk has been added, reboot the appliance so that the disk becomes active.

Add space to an existing disk in locally attached storage

If it is not possible to add a new disk, you can add space to an existing local storage disk as follows.

1. Log into your hypervisor's management console, and follow the normal procedure to increase the size of the local storage disk. (You may have more than one of these disks, and can expand each one of them.)
2. Log into the management console and select **System > Manage Appliance**.
3. Select Expand Storage. A list of expanded disks is displayed. (You might also see attached disks that are not yet part of the layer repository, but you can ignore those.)
4. Notice that the New Size of the disk you expanded is larger than the Current Size.
5. Select the check box for the disk that you want to expand to the New Size.
6. On the Confirm and Complete tab, click Expand Storage.

Add storage locations for user layers

When you enable user layers on a layered image, the data and settings for each user are persisted between sessions.

When deploying with user layers enabled, you can add storage locations for those layers, rather than allowing user data to be saved on the appliance's main file share.

When configuring user layer storage locations:

- You can assign groups of users to each location.
- The first storage location added to the appliance becomes the default location for user layers not associated with any other storage location.
- Storage locations are listed in priority order.
- If a user belongs to more than one group and those groups are assigned to different storage locations, the person's user layer will be stored in the highest priority storage location. Once the person's user layer is saved to the highest priority location, if you change the priority order of the storage locations that the user is assigned to, data saved up until that point will remain in the previously highest priority location. To preserve the person's user layer, you *must* copy the their user layer to the new highest priority location.

Create user layer storage locations

To add a storage location:

1. Log into the management console.
2. Select **System > User Layer Storage Locations**.
3. Select **Add Storage Location**. A list is displayed of file shares, except for the appliance's main file share.
4. Select **Add Storage Location**, and enter a Name and Network Path for the new location.
5. On the User Layer Assignments tab, expand the directory tree and select the check box(es) for one or more groups to add to the new storage location.
6. On the Confirm and Complete tab, click **Add Storage Location**.

Next, you must set security on the user layer folders.

Configure security on user layer folders

You can specify more than one storage location for your user layers. For each storage location (including the default location) you need to create a /Users subfolder and secure that location.

The security on each user layer folder must be set to the following values by a domain administrator:

Setting name	Value	Apply to
Creator Owner	Modify	Subfolders and Files only
Owner Rights	Modify	Subfolders and Files only
Users or group	Create Folder/Append Data, Traverse Folder/Execute File, List Folder/Read Data, Read Attributes	Selected Folder Only
System	Full Control	Selected Folder, Subfolders and Files
Domain Admins, and selected Admin group	Full Control	Selected Folder, Subfolders and Files

Set security on the user layer folders

1. Log into the management console.
2. Select **System > User Layer Storage Locations**. The file shares displayed are the storage locations defined for user layers. For example, say you've defined three storage locations so that you can more easily manage storage for Group1 and Group2 separate from everyone else in the organization:
 - Default location - \MyDefaultShare\UserLayerFolder\

- Group1 - \MyGroup1\Share\UserLayerFolder\
- Group2 - \MyGroup2\Share\UserLayerFolder\

Note: The appliance's main file share, which is used for storing OS, app, and platform layers, is not listed as a user layer storage location.

3. Create a \Users subdirectory under each file share:

- \MyDefaultShare\UserLayerFolder\Users\
- \MyGroup1Share\UserLayerFolder\Users\
- \MyGroup2Share\UserLayerFolder\Users\

4. Apply the security settings listed above to each /Users subdirectory.

Appliance settings

September 6, 2025

The Citrix App Layering™ appliance is a virtual appliance that uses Enterprise Layer Manager (ELM) technology. The appliance coordinates communication and manages copies of your layers and image templates.

Based on CentOS, the appliance hosts the management console. The console lets you create and manage layers. It also lets you publish layered images using those layers.

You can log into the Appliance Configuration utility to modify the following administrator settings:

- Password
- Network address
- NTP servers
- Time Zone

Note

Appliance settings are not available for editing in Azure.

Before you start

Make sure that:

- The App Layering appliance is running in your hypervisor.
- You have the password for an account with administrator privileges

Log into the appliance using an account with administrator privileges

Using either your hypervisor console or SSH, log into the appliance as administrator. (The first time you log in, you use the default password, **Unidesk1**) The Appliance Configuration utility opens.

Configure networking (includes Static IP Address option)

You can change the appliance's IP address and DNS servers. The default **DNS** settings are retrieved using Dynamic Host Configuration Protocol (DHCP).

If DHCP is not available and you select **Static**, you are prompted to enter the **IP addresses** for your DNS servers.

Note:

If you change your appliance's IP address, you must [Manually register the App Layering agent with the appliance](#) so that the appliance can communicate with the agent.

To change the appliance IP address:

1. Using either your hypervisor console or SSH, log into the appliance as administrator. (The first time you log in, use the default password, **Unidesk1**) The Appliance Configuration utility opens.
2. At the Action prompt, enter **C** (for Configure Networking), and press Return.
3. At the next prompt, type **D** for Dynamic (DHCP) or **S** for Static.
If you choose **Static**, you are prompted for the IP address and Subnet mask. Also enter the default addresses for the Gateway and DNS server.
4. When prompted, enter **Y** to save settings.
5. At the Action prompt, enter **Q** to quit.
6. Restart the appliance.

Synchronize the system clock with NTP servers

You can synchronize the system clock on the appliance by configuring NTP servers. You can specify:

- How many NTP servers you need, with 6 being the maximum.
- Add and remove NTP servers, as needed.

Where possible your existing servers are used as defaults.

1. Using either your hypervisor console or SSH, log into the appliance as administrator. (The first time you log in, you use the default password, **Unidesk1**) The Appliance Configuration utility opens.

2. At the Action prompt, enter N for NTP servers change, and press Return. A list of your current NTP servers is displayed.
3. At the prompt, specify how many NTP servers you need by typing a number from 0 to 6.
 - 0 - All servers are removed (you receive a warning).
 - 1–6 - You are prompted to accept or replace each of the current servers.
4. For each server, press Enter to accept the current value. Or, enter a new server address (Example: 3.pool.ntp.org). Once the last address is entered, an NTP Server Summary is displayed.
5. Enter S to save the settings.
6. At the Action prompt, enter Q to quit.
7. Restart the appliance.

Change the Time Zone

1. Using either your hypervisor console or SSH, log into the appliance as administrator. (The first time you log in, you use the default password, **Unidesk1**) The Appliance Configuration utility opens.
2. At the Action prompt, enter T for Timezone change, and press Return. The current time zone is displayed.
3. Press Enter to display available timezones. The time zones listed are in alphabetical order, starting with the
4. Advance through the timezone codes until you see yours:
 - Enter - Advances one line at a time.
 - Page Up Page Down - Displays the next or previous screen full of choices.
 - Or search the timezones:
Type Slash (/) and part of the name you are looking for.
5. When your timezone is displayed, press Q to get to the prompt.
6. Type the number for your timezone. The timezone you entered is displayed.
7. Press Enter to complete the change.
8. At the Action prompt, enter Q to quit.
9. Restart the appliance.

App Layering services

May 30, 2019

The following services run on the App Layering appliance:

- Management service

- Layering service
- BITS server service

The services are displayed on the System tab.

Management service

The App Layering appliance uses the Management service to communicate with the following servers, agents, and platform software:

- Active Directory
- Windows file servers
- Network time servers
- Unix file servers
- DHCP server
- App Layering agents
- Your hypervisor and provisioning service

The firewall ports for each of the above components must be open. For details, see [Firewall ports](#).

Layering service

The Layering service manages your layers and image templates. The appliance keeps a set of master layers on its local storage, and Elastic layers on the appliance's File Share.

Note: Elastic layers are copies of App layers stored on the appliance's File Share. Elastic layers are delivered upon user login, rather than as part of the layered image.

On the System module, the Layering service shows the amount of available space on each of your File Shares.

BITS Server service

The App Layering appliance copies files to and from the appliance using Microsoft's BITS Server service and the location specified in the connector configuration.

Firewall ports for the BITS Server service

Open the firewall ports for the BITS Server service. For details, see [Firewall ports](#).

BITS Server logs

BITS Server produces its own logs, which you can find here:

/var/log/Unidesk

The logs are based on the log4net configuration settings, which you can find in the following location:

source\BitsServer\Citrix.AppLayering.BitsServer\log4net.config

If you want to change the Uploads folder location

You can mount another volume to use for BITS uploads by configuring the path to the volume for both the BITS server and the Hyper-V connector.

1. Mount the new volume.
2. Edit the Uploads folder location for BITS Server in the following json file:

`/var/aspnetcore/bits-server/appsettings.Production.json`

Change default UploadFolder settings of /mnt/repository/Uploads to the new location, for example /mnt/test/Uploads.

3. Locate the connector configuration file:

`/usr/local/lib/node_modules/unidesk-hyperv-connector/config.json`

Change the default upload folder (called fileUploadFolderPath) from /mnt/repository/Uploads to the new location:

4. Restart both services:

```
1 systemctl restart kestrel-bits-server
2 systemctl restart unidesk-hyperv-connector
```

5. Update the group and permissions on the new Uploads folder:

```
1 chmod 770 /mnt/test/Uploads
2 chmod g+s /mnt/test/Uploads
3 chgrp apache /mnt/test/Uploads
```

Directory service

September 6, 2025

You can configure the appliance to connect to a directory service, for example, Active Directory. When you connect to your directory service, you will create one or more Directory Junctions to access specific domains or OUs. The appliance does *not* modify the directory service you connect to. The software caches the attributes for each directory service entry, so that if the connection to the directory service is lost temporarily, the software can use the cached information for management tasks.

When creating a Directory Junction, you use the following industry standard acronyms:

- OU - Organizational Unit
- DC - Domain Component

About connecting the appliance to a directory service

In the Management Console, the **System > Directory Services** displays information on Users and Groups.

Supported protocols

When binding to a directory service, the App Layering appliance is compatible with the following secure socket and transport layer protocols:

- Secure Socket Layer:
 - SSL 3.0
- Transport Layer Security:
 - TLS 1.1
 - TLS 1.2

What happens when you add Directory Junctions

Each Directory Junction that you create specifies a starting node in the directory tree. A new directory junction cannot include users who are already members of another junction, and junctions cannot be nested.

If you add a Parent Directory Junction, all of its children are migrated to that junction. All imported Users and Groups will be moved to the Parent, along with all Elastic Assignments. After being moved, the Child Directory Junctions are deleted.

**If you're creating several Distinguished Names

**

The system compares the Domain Component first—the portions of the Distinguished Name that start with “**DC=**”. Please be aware that in Distinguished Names, order matters. For example, **DC=A,DC=B** is different than **DC=B,DC=A**. The system adds separate Directory Junctions if their DC components differ, or if their DC components match and the remaining components do not overlap. Directory Junctions are merged if their DC components match and their other components are related.

User attributes are imported from the directory service

The App Layering software imports and caches user and group attributes from your directory service when:

- You assign administrator privileges to a user.
- The values of the attributes change in the directory service.

The attributes that the software caches are read only. All changes to the attributes for directory service users come from the directory server.

Imported attributes are synchronized regularly

The software synchronizes the information it caches for directory service users with the directory service every 12 hours. If the software discovers that a user is no longer an object in the directory service, it classifies the user as abandoned (you can view this information in the Information view for the user).

Create a directory junction

1. Click **System > Directory Services**.

2. Click **Add Directory Junction**.

3. Specify the details for the directory server:

- **Server address** - The name for the server that you use for the directory service (IP Address or DNS Name).
- **Port** - Specify the port number for communicating with the directory server.
- **Use SSL** - Click to enable Secure Sockets Layer (SSL) communication.
If certificate errors occur, a list of these errors displays. If you are sure it is safe to ignore them, click **Accept and Continue**.
- **Bind Distinguished Name (DN)** - To determine the correct syntax for the Bind DN or user name, see the documentation for your directory. The following examples show some of the ways you can specify a user for the directory service:

- domain\username
- username@domain.com.

- **Bind Password** - Type the password.
- **Base Distinguished Name** - Specify where the software starts searching for users and groups in the remote directory service.
- **Directory Junction Name** - The name of the folder that you see in the tree view. You can use any name, including the name of a domain in your directory service tree.

4. Click **Confirm and Complete**.

Migrate Appliance

November 6, 2024

The Migration feature serves two main purposes

1. As of June 30th, 2024, CentOS Linux 7 has reached End of Life (EOL). To ensure optimal performance and security, it is highly recommended that any ELM running on release 2403 or earlier migrate off of CentOS 7. Utilize the migration feature with the latest version of release 2403 on a Source ELM in order to migrate onto a Target ELM running release 2409 or later.
2. The migration feature allows you to copy or migrate resources between ELMs, even if they are on different hypervisors.

Note:

The migration feature will be available to migrate off any ELM deployed on a supported hypervisor. However, at this time, GCP will not be supported as a Target ELM destination.

Before you migrate

Notes

- It is suggested that a snapshot/backup of both the Source and Target ELM be captured before migration.
- This process is **non-destructive** to the Source ELM layer data.
- The configuration of the Source ELM will be modified during a **Move** migration. See [Phases of migration](#) section for additional details.
- A custom certificate imported into the source ELM will not be migrated to the target ELM. You will need to recreate the certificate on the target ELM after completing the migration.

Permissions required

- Requires the Administrator role.
- No additional permissions required.

Resources required

- Source ELMs using a release earlier than 2403 must be [upgraded](#) to release 2403.
 - If you are running release 2403 and do not see the migration feature in the UI, ensure you have applied the latest 2403 upgrade which includes the migration feature.
 - The download for the latest version of release 2403 can be found [here](#).
- Source ELMs using a release of 2409 or later require no additional upgrading.
- Target ELM must use release 2409 or later.
 - See [downloads](#) for the latest version of release 2409 and any later releases available.
 - Target ELM, at minimum, must have the same amount of storage space as the Source ELM.
 - Target ELM must not have any layers created before migration.
 - Source ELM and Target ELM must have point-to-point network connectivity

Note:

A fast, local network connection will ensure the best experience with the migration feature.

Types of migration

There are two migration types available: **Copy** and **Move**.

Copy	Move
All layers, templates, and connectors are copied over.	All layers, templates, and connectors are copied over.
All previously registered agents will be registered on the Target ELM.	All previously registered agents will be registered on the Target ELM.
Source ELM will remain unchanged and can continue to be used.	Source ELM will no longer have an associated network file share.
Target ELM will require a new network file share , as two different ELMs cannot utilize the same network file share.	Target ELM will have the network file shares from the Source ELM set up.
	All Source ELM connector caches will be cleaned up.

Note:

In a Copy Migration, the Target ELM network file share should not be the same as the Source ELM. ELMs are not designed to share the same network file share. Doing so would cause issues for both the Source and Target ELMs.

Phases of migration

The migration process is broken up into two different phases: **Phase One** and **Phase Two**.

Phase One

- Layer Phase: Layers are scanned and copied over from the Source ELM to the Target ELM.
- Layers copied over to the Target ELM will be visible in the UI following a successful migration.
- Upon completion of Phase One, the user will be prompted with an option to **Proceed** and move into Phase Two.

Note:

Depending on the number and size of layers in the Source ELM this phase can take anywhere from several minutes to several hours.

Phase Two

- Database Phase: The database is transferred over from the Source ELM to the Target ELM.

Note:

This phase is relatively quick and will only take a few minutes.

Following the completion of a **Copy**, the user can continue to utilize the Source ELM and can configure the Target ELM for usage when ready.

Following the completion of a **Move**, the Target ELM will contain transferred resources and configurations after a successful migration.

Migrating an appliance

With the exception of Step 1, all of the following steps should be performed on the Source ELM.

1. Deploy a new Target ELM following the documentation for release 2409 or later.
2. The migrate feature can be accessed under **User Menu > Migrate Appliance**.

3. Supply the following information for the Target ELM:

- IP/FQDN
- Username (root user)
- Password (root user)
- Migrate Reason
 - Copy
 - The Source ELM will maintain ownership of the current resources.
 - The Target ELM will require new resources and configuration setup.
 - Move
 - The Target ELM will own/manage the current resources upon completion.
 - The Source ELM will have the configured network file share removed upon completion.
 - The Source ELM can be re-used and will require a new network file share.

4. Select **Confirm and Complete**.

5. Select **Migrate**.

- Once a migration has started the UI will prevent a user from starting a second migration until the current one has been completed or cancelled.

6. On the left-hand menu select **Tasks**.

7. In the active tasks table select the current running migration task and then select **View Details**.

8. View the current status of the migration.

- The Source ELM will continue to be available for usage during this time.
- During Phase One, new layers created at any point will be copied over to the Target ELM.

9. After Phase One has completed, user confirmation is required on the Source ELM UI to begin Phase Two of the migration process.

UI access during a migration

During a migration there are varying levels of UI access for both the Source and Target ELMs depending on the phase a migration is in.

Phase One

- Source ELM - Everything is enabled.
- Target ELM - Expand Storage and Export Logs are enabled.

Phase Two

- Source ELM - Expand Storage and Export Logs are enabled.
- Target ELM - UI is inaccessible until migration has completed.

Recovering from a cancelled or failed migration

Note:

The Source ELM will always remain unchanged and accessible until a Move migration has successfully completed.

Cancelling or failing a migration can lead to varying states for the Target ELM. Below are various states the Target ELM can end up in, and the required course of action to resolve any issues.

Cancelling or failing a migration

- Phase One
 - Target ELM will store copied layers in the layer repository.
- Phase Two
 - A user attempting to cancel the migration will be presented with a pop-up confirming the cancellation and inform the user of the non-recoverable state the Target ELM will be left in.
 - Target ELM will be in a **non-recoverable** state.

Recovering from a cancelled or failed migration

- Phase One
 - Starting a new migration will resume the migration process from where it left off. No additional steps required.
- Phase Two
 - Target ELM should be reverted to a snapshot taken prior to migration and retried.

Note:

A known cause of a failed migration is due to lack of storage on the Target ELM. Expanding storage space will fix the issue, and an active migration should continue from the point it left off.

Upgrade

September 6, 2025

Upgrading from release 2403 or earlier

If you are using a version of App Layering that is release 2403 or earlier, you **must** utilize the [migration feature](#) in order to access release 2409 or later. There is no in-place upgrade for release 2403 to release 2409. In order to upgrade, a new appliance running 2409 or later must be deployed, and then the migration feature must be used to migrate the configuration from the 2403 appliance to the newly deployed appliance. To access the migration feature, you will first need to confirm you are upgraded to the latest version of release 2403. The download for the latest version of release 2403 can be found [here](#), or can be upgraded through the appliance.

You can refer to the [Migrate Appliance](#) documentation for additional details. At a high level, these are the steps required to get your appliance running on release 2409 or later:

1. [Download](#) and deploy the latest available release, [configure](#) your appliance, and run it on your network.
2. Take a snapshot of BOTH your current 2403 appliance and the newly deployed appliances.
3. Upgrade your 2403 appliance to the latest available version, which will include the migration feature.
4. On the 2403 appliance, in the top-right user menu, click **Migrate Appliance**.
5. There are two migration modes: **Copy** and **Move**:
 - a) Copy will replicate the layer data and configuration to the new appliance, but not move or change any of the configuration in the old appliance.
 - The new appliance will require manual configuration before initial use.
 - b) Move will re-map the network file shares to the new appliance, and remove them from the old appliance.
 - It will leave the old appliance without a configured network file share, and the new appliance will be ready for use.
6. There are two phases to a migration, and user confirmation is required between Phase One and Phase Two in order to complete the migration process to the new appliance.
7. A migration is a task, and can be monitored and managed via the task list in the UI.

Upgrading from release 2409 or later

For the latest fixes and features, including compatibility with other software packages that you use, we encourage you to stay current with the App Layering upgrades. The upgrade process is partially au-

tomated, in that the appliance periodically checks for the latest package. It downloads new packages, verifies, and extracts the files. Users receive a message the next time they log in, and administrators with proper permissions can start the upgrade.

Besides upgrading the App Layering appliance, also expect to upgrade the:

- App Layering agent, if it is installed on your hypervisor and provisioning servers.
- Published layered images (requires provisioning your servers with the images).

This article explains how to complete each of the associated upgrades.

Before you upgrade

- Verify that a network file share is configured.
- Back up the appliance.

Verify that a network file share is configured You can confirm the share by logging on to the appliance and navigating to **System > Network File Share**. After ensuring you configured the file share, you can upgrade the appliance.

Back up the appliance Take a snapshot or checkpoint of the appliance.

Upgrade the appliance

App Layering upgrades are partially automated. The appliance periodically checks for upgrades and downloads the latest one, as long as the correct permissions, and other requirements, are in place.

The **Upgrade** folder includes the appliance upgrade, the agent upgrade, and a folder of tools to use on your OS layer.

What happens if an upgrade is found If an upgrade is found during the automated check, the latest available compressed folder is downloaded to your appliance. The appliance verifies the download, and extracts the files in the background.

Meanwhile, every user receives one of the following messages the next time they log in:

- **Start Upgrade:** Run the App Layering appliance software upgrade (Administrators only).
- **Close:** Dismisses the message, so you can manually start the upgrade later using **User tab > Upgrade Appliance**.

For more detail about what the appliance checks for, see [How the upgrade checks work, in detail](#).

Who can start the upgrade *Only an administrator can start the upgrade.*

If your firewall prevents automatic downloads of the upgrade package If firewall requirements prevent automatic downloading, download the upgrade package from the Citrix® download site. Copy the package to the network fileshare where the appliance can access it.

1. Navigate to the .zip file and extract the files.
2. Log on to the management console, select **User** and then click **Upgrade Appliance**.
3. Verify the upgrade path, and click **Upgrade**. The upgrade process starts and opens a Statuspage in the browser.
4. When the upgrade is complete, the status changes to **Upgrade Status: Complete**. Refresh the webpage to return to the management console.
5. Verify that the upgrade was successful by clicking the **About** link in the management console to confirm the version number.

Upgrade the App Layering agent (if installed)

Next, upgrade the App Layering agent. This component enables an appliance or packaging machine to run PowerShell commands locally. You can expect to find the agent installed in the following locations, if they exist in your environment:

- Hyper-V servers that App Layering uses for layer creation or image publishing.
- Citrix Provisioning™ servers where you publish layered images.
- Servers where a connector configuration runs PowerShell scripts. To view any existing connector configurations, open the management console and select **System > Connectors**.

To upgrade the App Layering agent in those locations:

1. Copy the agent upgrade file to the servers where the agent is installed.
2. Double-click the agent upgrade file, and follow the instructions for upgrading the agent.

Upgrade your published layered images

App Layering upgrades include driver updates, new features, and the bug fixes documented in the [What's new](#). Once you upgrade the appliance, upgrade your published images.

To upgrade your published images, you select each of your image templates, verify the settings, and use it to [publish](#) new versions of the layered images. Use the new images to provision your systems.

To apply the upgrade to your published layered images:

1. Log into the layering management console.

2. Select the Images tab. Your image templates are displayed.
3. If you want to verify or update a template's settings, select the image template, and click **Edit Template**. Edit the settings and click **Save Template and Publish**.
4. If the image template does *not* need editing, simply select the template, click **Publish Layered Image**.
5. Once published, use the new layered images to provision your systems.

Background: How the upgrade checks work, in detail

When the App Layering appliance checks for an upgrade:

- **If an update is not available:** Nothing happens. Another check is made at the next scheduled interval.
- **If an update is available, but there is no network file share configured:** The user receives a message that there is an upgrade available. It asks you to finish configuring the network file share.
- **If an update is available:** A job is started to “Download Upgrade Media.” You can check for progress through the following tasks:

- Downloading the upgrade media to local storage.
- Ensuring the checksum of the successfully downloaded upgrade package is correct.
- Extracting the downloaded upgrade package to the configured Network File Share. The download is extracted to the appliance's file share:

Location: *NetworkFileShare*AppLayeringVersion**

Example: *\MyServer\ApplayeringFileShare\4.0.8*

- If extraction is successful, the next time any user logs in they will be notified that an upgrade is available.
- If at any time during this process an error requires Administrator intervention, the job fails with an error. For example:
 - Out of space on local storage.
 - Out of space on the network file share.
 - Invalid files found.

Note:

If a job fails, it retries at the next check interval, regardless of whether the issue has been resolved.

- **If another update is found before a previously downloaded one is installed** - The new upgrade is downloaded, and once successfully completed, becomes “Upgrade Available.”
- **If one upgrade is downloading when another is made available** - The running download is aborted and a new download is started. All files related to the in-progress download are deleted.

Optional: How to check for available upgrades manually

Automated upgrade checks always pull the latest version, but you can manually check for updates.

To manually check for an update:

1. Log into the management console.
2. Click the **User** tab and then the **Upgrade Appliance** action. The latest version displays in the **Upgrade Disk** field.

If you are logged into a desktop as Administrator, and an Upgrade badge is displayed:

1. Click **Start Upgrade**. The download information is displayed in the **Upgrade Disk** field.
2. As Administrator, you can select a different **Upgrade Disk**.

Use the detailed steps in the earlier section, [Upgrade the appliance](#).

Users

September 6, 2025

This section explains how to [manage your users and user groups](#).

For more information on assigning user roles, see [Assign roles](#).

For more information on connecting to a directory service, see [Connect to a directory service](#).

Users and groups

September 6, 2025

When you first install the App Layering appliance and log onto the management console, there is a built-in administrator account that you can use to get started. This administrator has the rights to perform all App Layering operations. You can edit this user's properties, including the name, password, and contact info. Be sure to change the password for this built-in administrator account as part of installing and configuring the appliance.

Local users and groups

The OS layer preserves any local users or groups that you add, but app layers, platform layers, user layers, and elastic layers do not. For example, users and groups that you add or change while installing

an application on an app layer, platform layer, or user layer don't persist. Either add the local user or administrator to the OS layer before installing the application, or consider installing the application on the OS layer.

AD users and groups

Other than the built-in administrator account, all users and groups are actually AD users and groups imported via one or more directory junctions. Once your directory junctions are created, you can assign roles to each user. You can see the roles assigned to a user in the User Details. For configuration details, see [Connect to a directory service](#).

Roles

Roles determine the App Layering modules a user can manage. When you assign roles to Directory Service users and groups, they can use their Directory Service credentials to log into the management console. For more information on assigning user roles, see [Assign roles](#).

Firewall ports

September 6, 2025

The App Layering appliance communicates with your hypervisor, provisioning service, and the App Layering agent. This article details the ports that the appliance uses to communicate both internally with other App Layering-related services, and externally with servers, such as NTP servers. Be sure to open the necessary ports in your firewall before you install the App Layering appliance.

During App Layering installation, you open ports that the appliance uses to interact with services on the virtual server where it is hosted. If there is a firewall between the App Layering appliance and the machine on which you are running the App Layering agent or one of the App Layering connectors, you must manually open the port in the firewall used for that purpose. If during installation you changed any of the ports from the default setting, be sure to open the correct port.

The App Layering appliance uses the TCP/IP protocol, and IPv4 is required. There are three main classes of communication:

- Accessing and managing the appliance.
- Talking to other App Layering agent service.
- Talking directly to hypervisors that don't require the agent.

Note:

The App Layering appliance must be connected to a network file share.

Admin user

By default, App Layering uses the following ports in your firewall for the Admin User to interact with the Management console on the App Layering appliance virtual machine.

App Layering appliance

The connector services for the various hypervisors and provisioning services listed below all run on the App Layering appliance.

App Layering			
Destination	Activity	Protocol	Ports
Appliance	Management console	TCP	80, 443
Legacy Azure connector service	Communication	TCP	3000 (HTTP), 3500 (HTTPS)
BITS Server	Disk upload	TCP	3015 (HTTP), 3515 (HTTPS)
Citrix Provisioning™ connector service	Communication	TCP	3009 (HTTP), 3509 (HTTPS)
Google connector service - Google Cloud	Communication	TCP	3016 (HTTP), 3516 (HTTPS)
Hyper-V connector service	Communication	TCP	3012 (HTTP), 3512 (HTTPS)

Internal connections

By default, the App Layering service uses the following ports in your firewall for internal connections between the appliance and each of its destinations.

In the table, the following shorthand is used:

- **Appliance** - The App Layering virtual appliance.
- **Agent** - refers to the App Layering agent.
- **Admin user** - A management console user who is assigned the App Layering Admin role.

- **Compositing machine** - A virtual machine used to create and update layers using the App Layering [compositing engine](#), including:
 - Virtual machine created when you use a connector with **Offload compositing** enabled to create a layer, add a version to a layer, or publish a layered image.
 - Virtual machine in which the `ImportOsLayer.ps1` script runs to import the OS image as a new OS layer.

App Layering Source	App Layering Destination	Activity	Protocol	Ports
Agent	Appliance	Initial registration	TCP	443
Appliance	Agent	Communication	TCP	8016
Agent	Appliance	Log deliveries from agent	TCP	8787
Appliance	vCenter, ESXI hosts	Communication with datastore via ESXI host	TCP	443
Appliance	Active directory	LDAP	TCP	389, 636
Appliance	Compositing machine	Communication	TCP	443
Compositing machine	Appliance	Communication	TCP	443
Compositing machine	Appliance	Layer disk access via iSCSI	TCP	3260
Admin user	Appliance	Legacy Azure connector communication	TCP	3000 (HTTP), 3500 (HTTPS)
Appliance	Azure	Communication	TCP	443
Admin user	Appliance	Citrix Provisioning connector communication	TCP	3009 (HTTP), 3509 (HTTPS)
Agent on Citrix Provisioning server	Appliance	Disk download	TCP	3009 (HTTP), 3509 (HTTPS)
Admin user	Appliance	Hyper-V connector communication	TCP	3012 (HTTP), 3512 (HTTPS)

App Layering

App Layering Source	App Layering Destination	Activity	Protocol	Ports
Agent on Hyper-V server	Appliance	Disk download	TCP	3012 (HTTP), 3512 (HTTPS)
Agent on Hyper-V server	Appliance	Disk upload	TCP	3015 (HTTP), 3515 (HTTPS)
Appliance	vSphere	Communication	TCP	443
Appliance	XenServer®	Communication	TCP	443
Appliance	Prism	Disk upload	TCP	2222
Appliance	Prism	Communication	TCP	9440
Source Appliance	Target Appliance	Appliance migration	TCP	22000
Target Appliance	Source Appliance	Appliance migration	TCP	22000

External connection

By default, use the following port in your firewall for external connections between the App Layering appliance and the destination listed below.

Note:

These URLs are only accessible by the appliance using the credentials defined for it. Attempting to browse these sites results in an error message.

App Layering Destination	Activity	Protocol	Ports
< https://applayeringwebapi.azurewebsites.net >	API Access	TCP	443
< http://alcdn.citrix.com >	Download upgrade media	TCP	80

Key ports

Basic appliance management and access (always required)

- HTTP - Port 80
- HTTPS - Port 443
- SSH - Port 22

Servers

- Active Directory server - Port 389 - LDAP protocol
- Active Directory server - Port 636 - LDAPS protocol
- Active Directory server - Port 53 - DNS protocol
- Windows file servers, SMB - Port 445 - SMB protocol
- Network time servers - Port 123 - NTP protocol
- Unix file servers - Port 2049 - NFS protocol
- DHCP server, DHCP - Port 67 - UDP protocol
- App Layering appliance - Port 68 - DHCP protocol

App Layering agent

The agent uses the following ports for communications with itself and the appliance.

- Appliance to agent server:
 - Commands from appliance/SOAP - Port 8016
- Agent server to appliance:
 - Registration - Port 443 HTTPS
 - Log export - Port 8787
 - Citrix Provisioning disk download - Ports 3009 HTTP, 3509 HTTPS
 - Hyper-V disk download - Ports 3012 HTTP, 3512 HTTPS
 - Hyper-V disk upload - Ports 3015 HTTP, 3515 HTTPS

Connectors to hypervisors and provisioning services

Connectors on the appliance allow the appliance to communicate directly with the supported hypervisors and provisioning services using the following ports.

- XenServer - Port 443

- Citrix Provisioning - Port 8016 (App Layering agent)
- Google Cloud - Port 443
- Microsoft Azure management - Port 443
- Microsoft Hyper-V - Port 8016 (App Layering agent)
- Nutanix AHV - Port 2222, 9440
- VMware vSphere - Port 443 (Virtual Center, and ESX hosts for disk transfers)



© 2025 Cloud Software Group, Inc. All rights reserved. This document is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of Cloud Software Group, Inc. This and other products of Cloud Software Group may be covered by registered patents. For details, please refer to the Virtual Patent Marking document located at <https://www.cloud.com/legal>. Citrix, the Citrix logo, NetScaler, and the NetScaler logo and other marks appearing herein are either registered trademarks or trademarks of Cloud Software Group, Inc. and/or its subsidiaries in the United States and/or other countries. Other marks are the property of their respective owner(s) and are mentioned for identification purposes only. Please refer to Cloud SG's Trademark Guidelines and Third Party Trademark Notices (<https://www.cloud.com/legal>) for more information.
