

Use Shift Toolkit to migrate or convert VMs

NetApp virtualization solutions

NetApp November 18, 2025

This PDF was generated from https://docs.netapp.com/us-en/netapp-solutions-virtualization/migration/shift-toolkit-overview.html on November 18, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Use Shift Toolkit to migrate or convert VMs	1
Learn about migrating VMs between virtualization environments using NetApp Shift Toolki	t 1
Use case	1
Toolkit Overview	1
Benefits of VM portability	2
How Shift toolkit works	2
Supported versions for the NetApp Shift Toolkit.	4
Supported VM guest operating systems	4
Supported ONTAP versions	4
Supported hypervisors	4
Install Shift Toolkit.	5
Prepare to install the NetApp Shift Toolkit for ONTAP storage	5
Requirements to install the NetApp Shift Toolkit	6
Install or upgrade the NetApp Shift Toolkit for ONTAP storage	10
Configure the NetApp Shift Toolkit	20
Run Shift toolkit	20
Shift Toolkit Configuration	22
Migrate VMs using Shift Toolkit	
Migrate VMs using the Shift Toolkit	
Migrate VMs from VMware ESXi to Microsoft Hyper-V using the Shift Toolkit	
Migrate VMs from Microsoft Hyper-V to VMware ESXi using the Shift Toolkit	63
Migrate VMs from VMware ESXi to Red Hat OpenShift Virtualization	84
Migrate VMs from VMware ESXi to Oracle Linux Virtualization Manager	100
Convert VMs using the Shift Toolkit	113
Overview	113
Convert to QCOW2 format	
Convert to VHDX format	119
Convert to RAW format	122
Monitor migration jobs with the Shift Toolkit dashboard	126
Job monitoring dashboard	
Understanding job status	
Configure advanced settings in the Shift Toolkit	127
Credential Security Service Provider (CredSSP)	
Logging and debugging	
Swagger	130
Script block and automation	134
Email notifications and alerts	
Abort and revert capabilities	
Migrate VMs from SAN environments for conversion with Shift Toolkit	
Requirements for SAN-based VMs	
Migration workflow for SAN environments	
Handling processor compatibility issues	
Next steps after converting or migrating VMs using the Shift Toolkit	139

Conclusion	. 139
Next Steps	. 139
Troubleshooting and Known Issues	. 139
Appendix	. 141

Use Shift Toolkit to migrate or convert VMs

Learn about migrating VMs between virtualization environments using NetApp Shift Toolkit

The NetApp Shift Toolkit is a standalone product designed to simplify and accelerate VM migrations across hypervisors, such as VMware ESXi, Microsoft Hyper-V, Oracle Linux Virtualization Manager, Redhat OpenShift, and others. It also supports disk-level conversions between various virtual disk formats.

Use case

Every organization is now seeing the benefit of having multi-hypervisor environment. With recent changes in the market, every organization is deciding on the best course(s) of action by weighing technical and commercial risks including migrating workload VMs to alternate hypervisors and focus on achieving business-defined objectives, and controlling vendor lock-in. This enables them to operate in an optimized fashion wrt licensing cost and extend IT budget on the right areas than spending for those unused cores on a specific hypervisor. However, the challenge has always been around migration time and the associated downtime.

With the NetApp Shift toolkit, migrating virtual machines (VMs) is no longer a concern. This standalone product enables fast and efficient migration of VMs from VMware ESXi to Microsoft Hyper-V. Additionally, it supports disk-level conversions between various virtual disk formats. Thanks to the out-of-the-box capabilities provided by ONTAP, these migrations can be incredibly swift, with minimal downtime. For example, converting a 1TB VMDK file typically takes a couple of hours, but with the Shift toolkit, it can be completed in seconds.

Toolkit Overview

The NetApp Shift toolkit is an easy-to-use, graphical user interface (GUI) solution that allows to migrate virtual machines (VMs) between different hypervisors and convert virtual disk formats. It utilizes NetApp FlexClone® technology to quickly convert VM hard disks. Additionally, the toolkit manages the creation and configuration of destination VMs.

Shift toolkit provides flexibility in a multi-hypervisor environment by supporting bidirectional conversion between the following hypervisors:

- VMware ESXi to Microsoft Hyper-V
- Microsoft Hyper-V to VMware ESXi
- VMWare ESXi to Oracle Linux Virtualization Manager (OLVM)
- VMWare ESXi to Red Hat OpenShift Virtualization

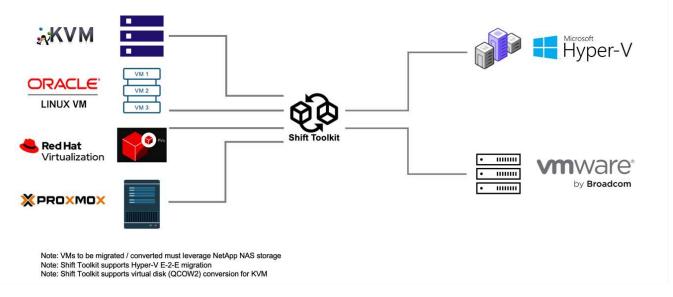
Shift toolkit supports disk-level conversions of virtual disks between hypervisors for the following disk formats:

- VMware ESX to Microsoft Hyper-V (virtual machine disk [VMDK] to virtual hard disk format [VHDX])
- Microsoft Hyper-V to VMware ESX (virtual hard disk format [VHDX] to virtual machine disk [VMDK])
- VMware ESX to KVM compatible hypervisors (VMDK to QCOW2)
- VMware ESX to KVM compatible hypervisors (VMDK to RAW)

Summarizing Conversion Options







Shift toolkit can be downloaded here and is available for Windows Systems only.

Benefits of VM portability

ONTAP is ideal for any hypervisor and in any hyperscalar. With FlexClone technology. VM portability in minutes is a reality than waiting for longer downtimes or settling down with pass through options.

Shift toolkit:

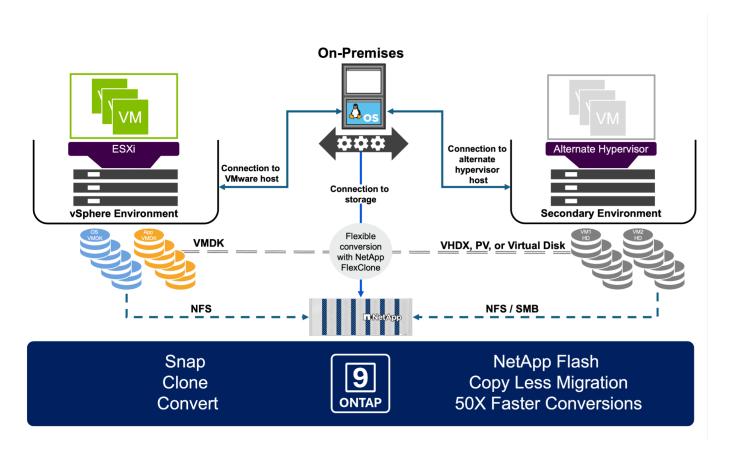
- · helps minimize downtime and enhances business productivity.
- offers choice and flexibility by reducing licensing costs, lock-in, and commitments to a single vendor.
- enables organizations looking to optimize VM licensing costs and extend IT budgets.
- reduces virtualization costs with VM portability and is offered free from NetApp.

How Shift toolkit works

At conversion time, Shift toolkit connects to VMware ESXi and Microsoft Hyper-V hosts and to shared NetApp storage. Shift toolkit leverages FlexClone to convert VM hard drives from one hypervisor to another by using three key NetApp technologies:

- Single volume and multiple protocols
 With NetApp ONTAP, multiple protocols can be easily used to access a single volume. For example,
 VMware ESXi can access a volume that is enabled with the Network File System (NFS) protocol, and
 Microsoft Hyper-V can access the same volume with the CIFS/SMB protocol.
- FlexClone technology
 FlexClone allows the rapid cloning of entire files or volumes with no data copy. Common blocks on the storage system are shared between multiple files or volumes. As a result, large VM disks can be cloned very quickly.
- VM disk conversion

The NetApp PowerShell Toolkit and Shift toolkit contain a large number of workflows that can be used to perform various actions on a NetApp storage controller. Included are PowerShell cmdlets that convert virtual disks to different formats. For example, VMware VMDK can be converted to Microsoft VHDX, and vice versa. These conversions are performed with FlexClone, which enables very rapid cloning and conversion of disk formats in one step.



Protocols and communication methods

Shift toolkit uses the following protocols during conversion or migration operations.

- HTTPS Used by the Shift toolkit to communicate with the Data ONTAP cluster.
- VI Java (openJDK), VMware PowerCLI Used to communicate with VMware ESXi.
- Windows PowerShell module Used to communicate with Microsoft Hyper-V.

The firewalls must enable traffic over the following ports:

Port	Protocol	Source	Destination	Purpose
443	TCP	Shift toolkit node	VMware vCenter	VMware Inventory
443	TCP	Shift toolkit node	VMware ESXi nodes	Invoke-vmscript dependency (routed)
443	TCP	Shift toolkit node	Target Hypervisor	Target inventory
443	TCP	Shift toolkit node	ONTAP system	ONTAP access
5985/5986	HTTP	Shift toolkit node	Hyper-V Hosts	WinRM

Supported versions for the NetApp Shift Toolkit

Verify that your Windows and Linux guest operating systems, ONTAP version, and hypervisors are supported by the NetApp Shift Toolkit.

Supported VM guest operating systems

The Shift Toolkit supports the following Windows and Linux guest operating systems for VM conversion.

Windows operating systems

- Windows 10
- Windows 11
- · Windows Server 2016
- Windows Server 2019
- Windows Server 2022
- Windows Server 2025

Linux operating systems

- CentOS Linux 7.x
- Alma Linux 7.x
- Red Hat Enterprise Linux 7.2 or later
- Red Hat Enterprise Linux 8.x
- Red Hat Enterprise Linux 9.x
- Ubuntu 2018
- Ubuntu 2022
- Ubuntu 2024
- Debian 12
- SUSE Linux Enterprise Server 12
- SUSE Linux Enterprise Server 15



CentOS Linux and Red Hat Enterprise Linux versions 5 and 6 are not supported.



Windows Server 2008 is not officially supported. However, the conversion process may work, and some customers have successfully converted Windows Server 2008 VMs. After migration, manually update the IP address, as the PowerShell version used for IP assignment automation is not compatible with Windows Server 2008.

Supported ONTAP versions

The Shift Toolkit supports ONTAP 9.14.1 or later.

Supported hypervisors

The Shift Toolkit supports the following hypervisor platforms for VM migration and conversion.



In the current release, end-to-end virtual machine migration is supported only with Hyper-V, VMware, OpenShift, Oracle Virtualization. For KVM destinations, only disk conversion is supported.

VMware vSphere

The Shift toolkit is validated against vSphere 7.0.3 or later.

Microsoft Hyper-V

The Shift toolkit is validated against the following Hyper-V roles:

- Hyper-V role running on Windows Server 2019
- Hyper-V role running on Windows Server 2022
- Hyper-V role running on Windows Server 2025

Red Hat OpenShift

The Shift toolkit is validated against Red Hat OpenShift and OpenShift Virtualization running 4.17 and later.

Oracle Linux Virtualization Manager

The Shift toolkit is validated against the following Oracle Linux Virtualization Manager versions:

- · Oracle Linux Virtualization Manager 4.5 or later
- The Oracle Linux Virtualization Manager host must have ovirt-engine-4.5.4-1.el8 RPM or later installed

KVM

For KVM destinations, the Shift Toolkit supports disk format conversion only (VMDK to QCOW2 or RAW). Hypervisor connection details are not required when selecting KVM from the destination dropdown. After conversion, use the QCOW2 disks to provision VMs on KVM-based platforms.

Install Shift Toolkit

Prepare to install the NetApp Shift Toolkit for ONTAP storage

Prepare to install the NetApp Shift Toolkit by ensuring your environment meets the prerequisites and selecting the appropriate installation package for your deployment scenario.

Before you begin

Verify that your environment meets the following requirements:

- Windows Server 2019, 2022, or 2025
- Dedicated virtual machine for the Shift Toolkit installation
- · Network connectivity between source and target environments
- Virtual machine running on VMware vSphere or Microsoft Hyper-V



Install the Shift Toolkit on a dedicated VM to enable management of multiple source and target hypervisors from a single server.

Select an installation package

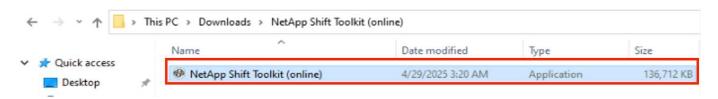
The Shift Toolkit is available in two installation packages from NetApp Toolchest:

Online installer (~130 MB)

- Downloads and installs prerequisites from the internet during installation
- · Requires internet connectivity during installation
- · Smaller package size for faster download

Offline installer (~1.2 GB)

- · Includes all prerequisites bundled within the package
- · Supports installation on VMs without internet access
- · No proxy configuration required
- Provides more control over the installation process





Use the offline installer for air-gapped environments or when you need complete control over the installation process without external dependencies.

Requirements to install the NetApp Shift Toolkit

Verify that your environment meets the hardware, connectivity, and ONTAP storage requirements before installing the Shift Toolkit.

Hardware requirements

Ensure the Shift Toolkit server meets the following minimum hardware requirements:

- CPU: 4 vCPUs
- Memory: 8 GB minimum
- **Disk space**: 100 GB minimum (900 MB available for installation)

Connectivity requirements

Verify the following connectivity requirements are met:

- · Shift Toolkit must be installed on a standalone Windows server (physical or virtual)
- The hypervisor and storage environment must be configured to allow the Shift Toolkit to interact with all components
- For Hyper-V migrations, the Shift server, ONTAP CIFS server, and Hyper-V servers must be on the same Windows Active Directory domain
- Multiple LIFs for CIFS and NFS are supported for use with Storage Virtual Machines (SVMs) during VM

conversions

 For CIFS operations, time settings must be synchronized between the Windows domain controller and the ONTAP storage controller

ONTAP storage configurations

Configure ONTAP storage components including SVMs, qtrees, and CIFS shares to support Shift Toolkit migrations.

Create a new SVM (recommended)

Although the Shift Toolkit permits the use of an existing SVM, NetApp recommends creating a dedicated SVM for migration operations.

Creating a new SVM provides the following benefits:

- Isolates migration operations from production workloads
- Ensures the SVM meets Shift Toolkit requirements without modifying production configurations
- Simplifies configuration for bi-directional migrations between VMware and Hyper-V

Use Storage vMotion to move VMs to a new designated NFSv3 datastore on the dedicated SVM without downtime. This approach ensures migrated VMs do not reside on the production SVM.

Use the ONTAP CLI, NetApp PowerShell Toolkit, or ONTAP System Manager to create the new SVM. For detailed steps, refer to the ONTAP documentation for provisioning a new SVM with both NFS and SMB protocols enabled.



For bi-directional migration between VMware and Hyper-V, enable both NFS and SMB protocols on the SVM and provisioned volumes.

Qtree requirements

Create qtrees on the volume that will host converted VMs. Qtrees segregate and store converted disk files based on the target hypervisor.

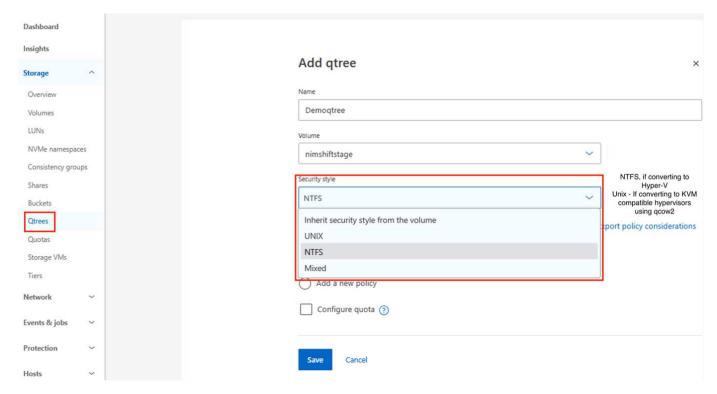
Security style by migration type:

- ESXi to Hyper-V: NTFS security style (stores converted VHDXs)
- Hyper-V to ESXi: UNIX security style (stores converted VMDKs)
- ESXi to OpenShift Virtualization (QCOW2): UNIX security style
- ESXi to OLVM (RAW or QCOW2): UNIX security style

Source Hypervisor	Target Hypervisor	Protocol	Source	Destination
VMware	Hyper-V	NFS	Volume (with NFS v3 and SMB)	Qtree with NTFS sec style within the same source volume
Hyper-V	VMware	SMB	Volume (with SMB and NFS v3)	Qtree with UNIX sec style within the same source volume
VMware	OpenShift	NFS	Volume (NFS v3)	New Cloned Volume as the PVC
VMware	Oracle Virtualization	NFS	Volume (NFS v3)	Qtree with UNIX sec style within the same source volume

The Shift Toolkit does not verify qtree security styles. Create qtrees with the appropriate security style for your target hypervisor and disk format.

For detailed steps, refer to Create a qtree in the ONTAP documentation.





The destination path must be on the same volume as the source VM.



For OpenShift Virtualization, converted QCOW2 files can optionally be placed directly on the volume without using a qtree. Use the Shift Toolkit GUI or APIs to perform this conversion.

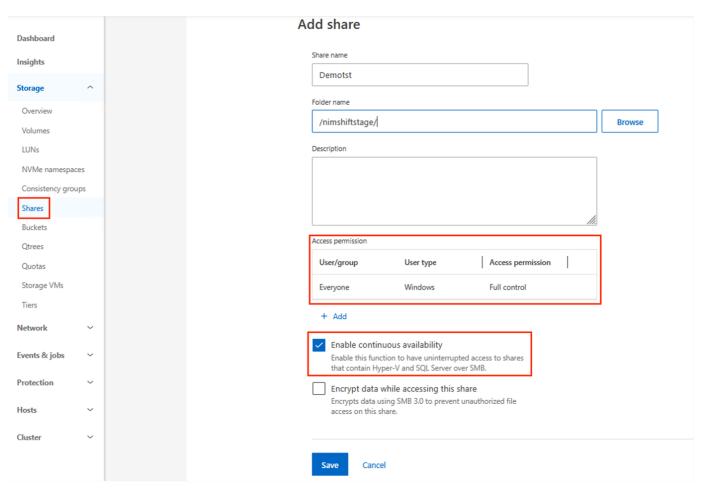
CIFS share requirements

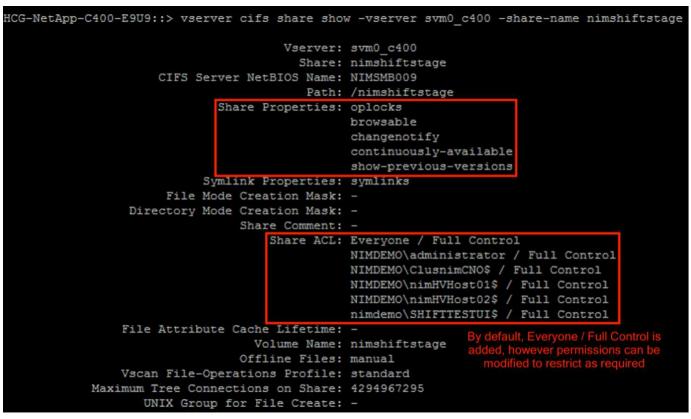
For Hyper-V migrations, create a CIFS share for storing converted VM data. Both the NFS share (source VMs) and CIFS share (converted VMs) must reside on the same volume.

Configure the CIFS share with the following properties:

- SMB 3.0 enabled (enabled by default)
- · Continuously available property enabled
- · Export policies for SMB disabled on the SVM
- Kerberos and NTLMv2 authentication permitted on the domain

For detailed steps, refer to Create an SMB share in the ONTAP documentation. Select the continuous availability property along with other default properties.







ONTAP creates the share with the Windows default share permission of Everyone / Full Control.

Install or upgrade the NetApp Shift Toolkit for ONTAP storage

Install or upgrade the NetApp Shift Toolkit after verifying that your environment meets the preparation and prerequisite requirements.

Install the Shift Toolkit

Download and run the installer to set up the Shift Toolkit on your Windows server.

Steps

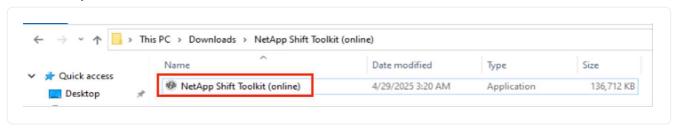
1. Download the Shift Toolkit package and unzip it.

Show example



2. Double-click the downloaded .exe file to initiate the Shift Toolkit installation.

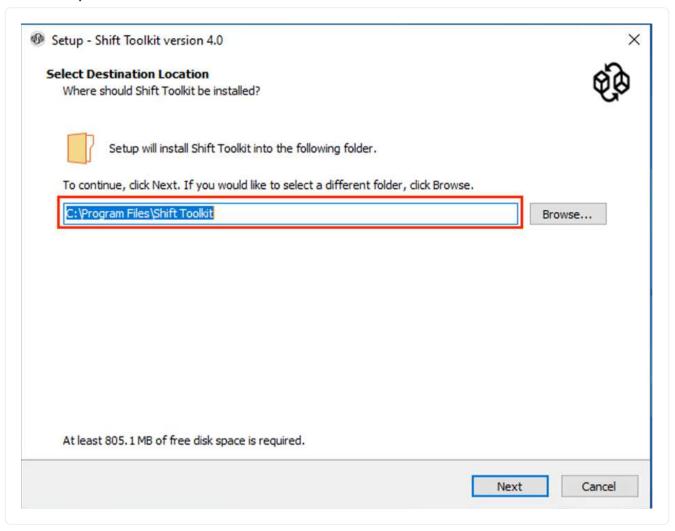
Show example



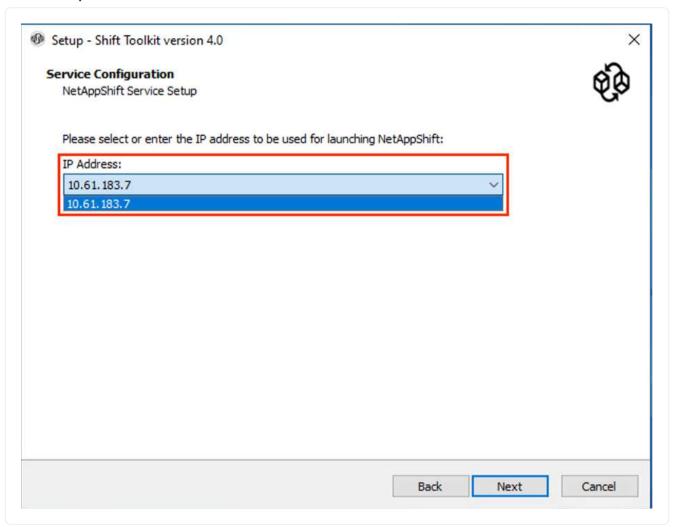


All pre-checks are performed during installation. If minimum requirements are not met, appropriate error or warning messages are displayed.

3. Select the installation location or use the default and click **Next**.



4. Select the IP address that will be used to access the Shift Toolkit UI.

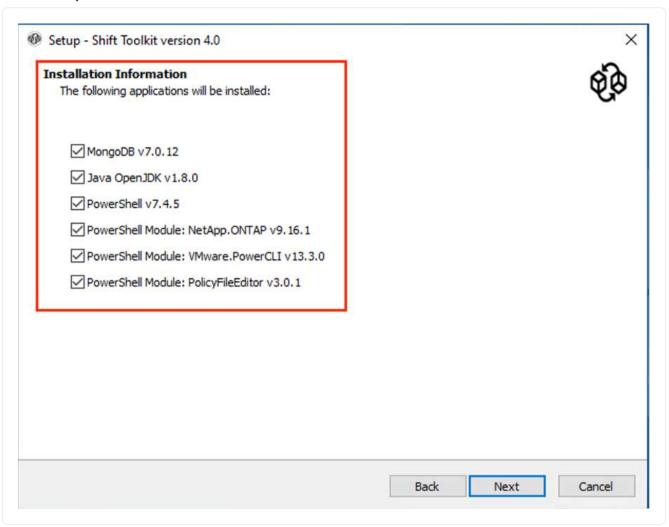


(i)

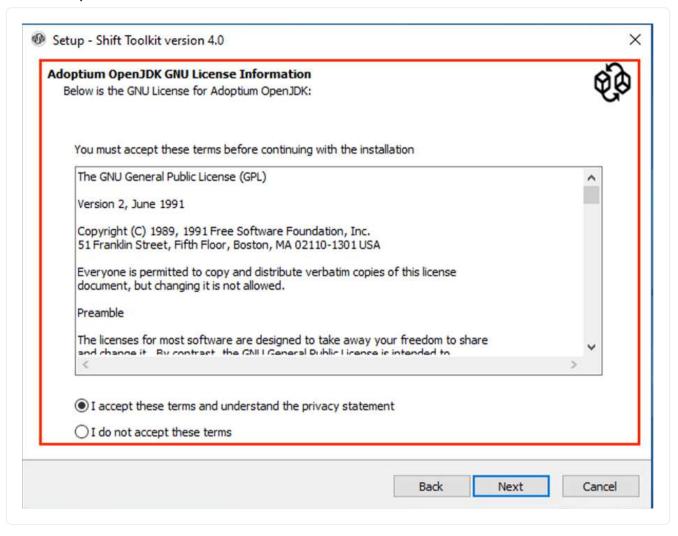
If the VM has multiple NICs, the setup process allows you to select the appropriate IP address from a dropdown.

5. Review the required components that will be automatically downloaded and installed, and then click **Next**.

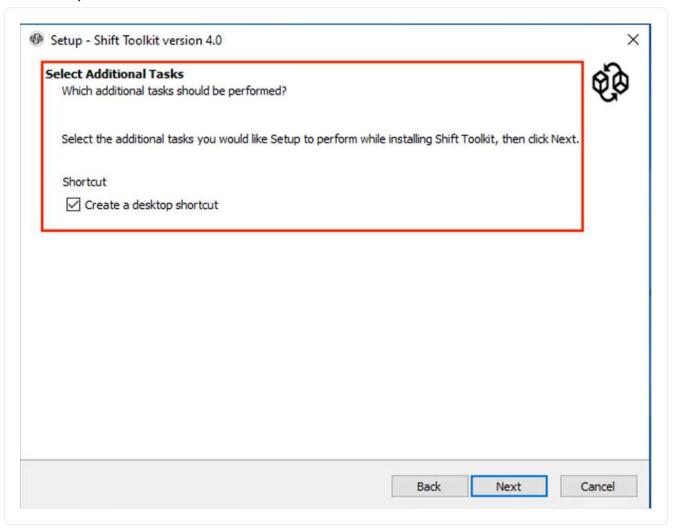
These mandatory components are required for proper Shift Toolkit operation:



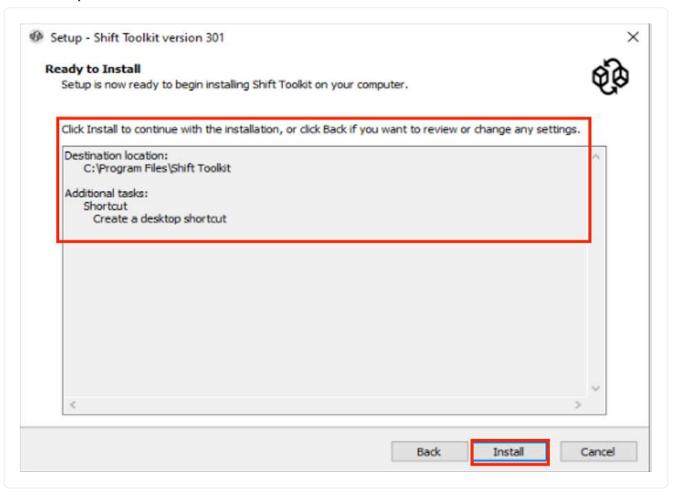
6. Review the Java OpenJDK GNU licensing information and click Next.



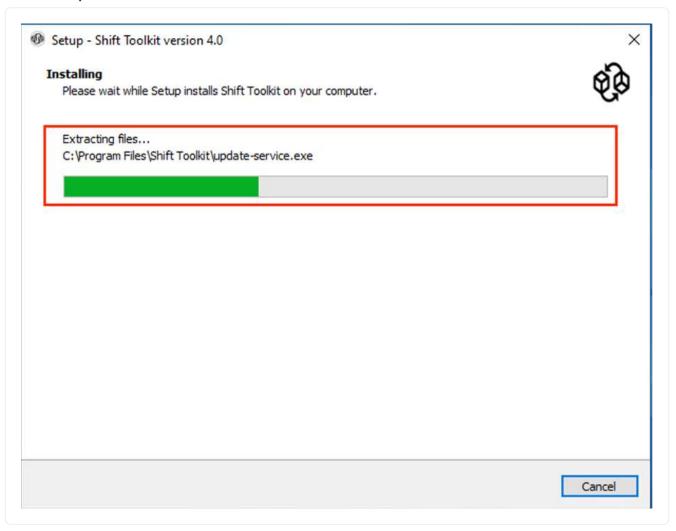
7. Keep the default setting for creating a desktop shortcut and click **Next**.



8. Click **Install** to begin the installation.

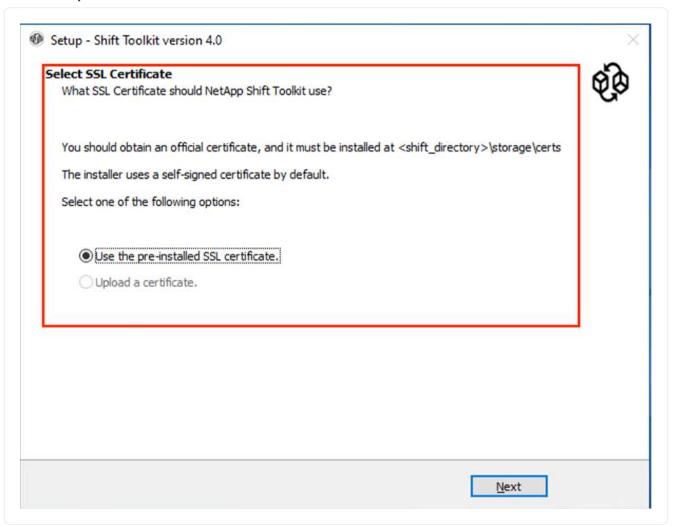


9. Wait for the installation to complete. The installer downloads and installs all required components. Click **Finish** when complete.



The installation can take 10-15 minutes.

10. Accept the self-signed certificate prompt and click **Next**.

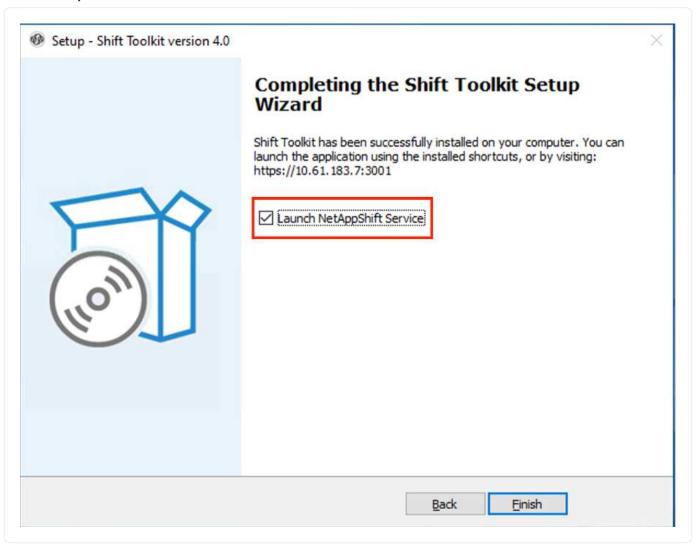




The self-signed certificate can be replaced with a third-party or CA-generated certificate. Replace the certificate in the certs folder located at <installation directory>\Storage\Certs.

Result

The Shift Toolkit installation is complete.



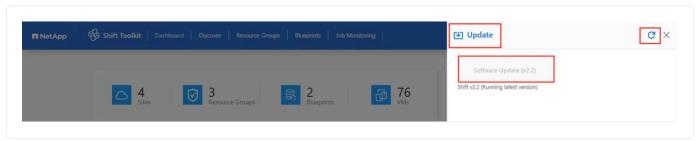
(i)

For VMs without internet access, the offline installer performs the same steps but installs components using packages included in the executable.

Upgrade the Shift Toolkit

Upgrades are fully automated and can be completed with a single click.

Show example



The Shift Toolkit updater service listens on port 3002 and performs the following steps:

1. Downloads the upgrade package

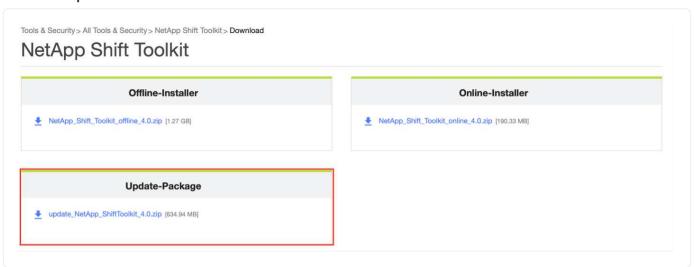
- 2. Stops the Shift Toolkit service
- 3. Extracts files and overwrites required files
- 4. Runs the update using the same IP address (retaining metadata)
- 5. Redirects the UI to the Shift Toolkit UI listening on port 3001

For deployments without internet connectivity

Manually download the upgrade package (filename starts with "update") from NetApp Toolchest and place it in the designated folder $C: \NetApp_Shift$.

Create this folder path if it doesn't exist. All other steps remain the same as the online upgrade procedure.

Show example



Configure the NetApp Shift Toolkit

Configure the Shift Toolkit to automate the migration or conversion of VMs) This process includes adding source and destination sites, configuring storage, grouping VMs into resource groups, creating migration blueprints, and scheduling migrations.

Run Shift toolkit

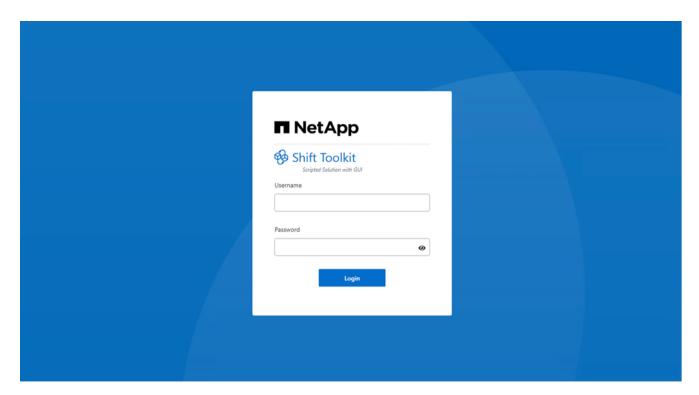
• Using the browser, access Shift toolkit UI by entering the http://<IP address specified during installation>:3001



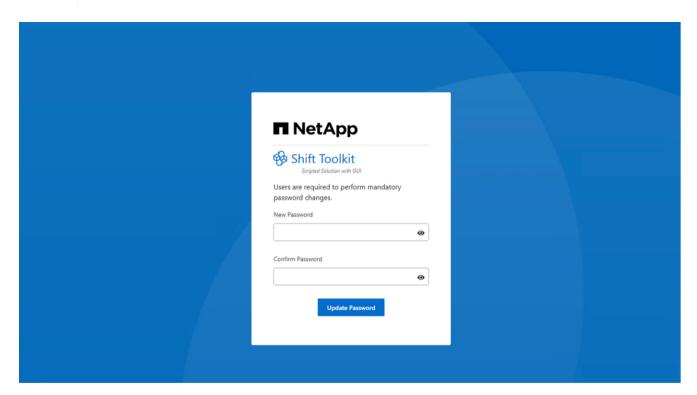
Use Google chrome or Internet explorer for best experience.

· Access the UI using default credentials as below:

Username: admin Password: admin



- Admins are required to change their credentials during the initial login.
- After the mandatory password change, the admin credential can also be changed using "Change Password" option after logging in to the GUI.



Once done, accept the legal EULA by clicking on "Accept and Continue"

LICENSE AGREEMENT FOR NETAPP, INC. TOOLS

This License Agreement ("Agreement") is a legal agreement between You either as an individual or as an authorized representative of a business entity (hereafter referred to as You and/or Your), and NetApp, Inc. (NetApp). NetApp is willing to license to You the NetApp software product accompanying this EULA, which includes, without limitation, computer software features, authorized updates and upgrades or other supplements to the software, images, music, text and/or animations incorporated into the software, media, printed materials, or online or electronic documentation, provided by NetApp or made available for download (collectively referred to as the Software). Your download, installation and/or use of the Software constitutes acceptance of all of the terms stated herein. If You do not agree with all of these terms, You must promptly return the Software to NetApp or the authorized reseller from which You obtained the Software.

- 1. License Grant. Subject to payment of applicable fees, if any, and the limitations and restrictions set forth herein, NetApp and its licensors grant to You a non-exclusive, non-transferable, worldwide, limited, royally-free license, without right of sublicense, to install and use the Software, in accordance with the terms contained in the user documentation accompanying the Software (the Documentation). Use of the Software outside the scope of the Documentation is unauthorized and shall constitute a material breach of this EULA. NetApp's licensors shall be a direct and intended third party beneficiary of this EULA and may enforce their rights directly against You in the event of Your breach of this EULA. The Software is licensed, not sold, to You.
- 2. License Restrictions.

2.1 Standard Use Restrictions. Subject to any express restrictions contained within the Documentation, You shall not nor shall You allow any third party to: (a) decompile, disassemble, decrypt, extract, or otherwise reverse engineer or attempt to reconstruct or discover any source code or underlying ideas, algorithms, or file formats of, or of any components used in the Software by any means whatever; (b) remove or conceal any product identification, copyright, patent or other notices contained in or on the Software or Documentation; (c) electronically transmit the Software from one computer to another or over a network; (d) use any locked or restricted feature, function, service, application, protocol, operation, or capability without first purchasing the applicable license(s) and/or obtaining a valid license enablement key from NetApp, even if such feature, function, service, application, protocol, operation or capability is technically achievable without a key; (e) sell, lease, rent, lend, sublicense, distribute or otherwise transfer in whole or in part the Software or the applicable license enablement key to another party or to a different storage controller or cluster; or (f) modify the Software, incorporate it into or with other software, or create a derivative work of any part of the Software. You revise the software will breach this EULA, and such derivative work is and shall be owned entirely by NetApp or its licensors. You hereby assign and agree to assign to NetApp or the licensor of the Software all right, title and interest in and to said derivative work is and shall be owned entirely by NetApp or its licensors. You assign and agree to use in the design, construction, operation or maintenance of any nuclear facility, aircraft operation, air traffic control or life support system. NetApp disclaims any express or implied warranty of fitness for such uses. If You use the Software for such applications, You agree to indemnify, defend and hold NetApp and its licensors harmless from all claims, actions,

only for the current version of the Software available for download. If NetApp, at its sole option, supplies updates to You, the updates will be considered part of the Software, and subject to the terms of this EULA.

3. Intellectual Property Rights. The Software is protected by intellectual property and copyright laws and treaties worldwide and may contain trade secrets of NetApp or its licensors, who have and

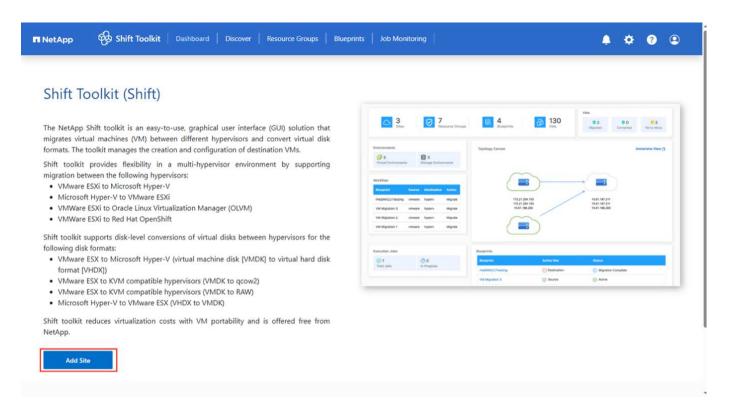
Accept and Continue

Shift Toolkit Configuration

Once the storage and connectivity to both the source and destination hypervisors have been configured properly, begin configuring Shift toolkit to automate the migration or conversion of the virtual machines to appropriate format, leveraging the FlexClone functionality.

Add Sites

The first step is to discover and add the source and then the target Hypervisor details (both hypervisors and storage) to Shift toolkit. Open Shift toolkit in a supported browser and use the default username and the password and click on "Add Sites".





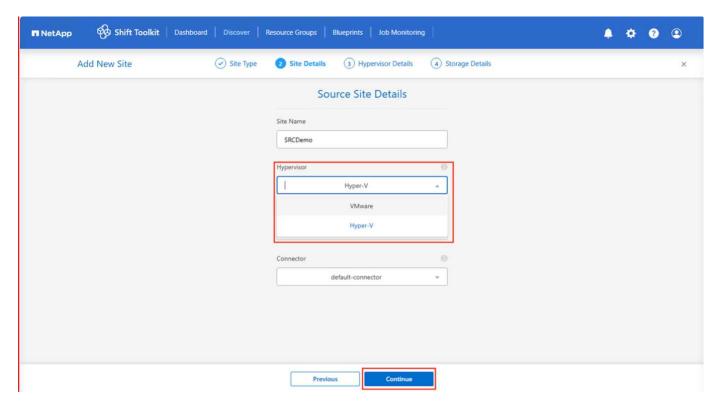
Sites can also be added using Discover option.

Add the following platforms:

Source

- · Source Site Details
 - · Site Name Provide a name for the site
 - Hypervisor Select VMware or Hyper-V as the source
 - · Site Location Select the default option
 - · Connector Select the default selection

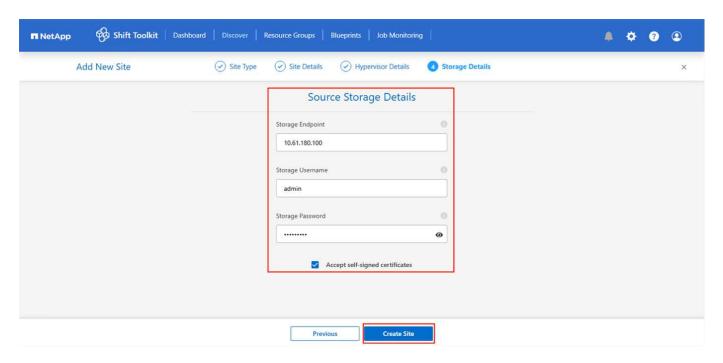
Once filled, click Continue.



- If the selection is VMware, enter the Source vCenter details.
 - Endpoint Enter the IP address or FQDN of the vCenter server
 - Username username to access the vCenter (in UPN format: username@domain.com)
 - vCenter Password Password to access vCenter for performing inventory of the resources.
 - vCenter SSL Thumbprint (optional)
- If the selection is Hyper-V, enter the Source Hyper-V details.
 - Endpoint Enter the IP address or FQDN of the standalone hosts or the failover cluster endpoint.
 - Hyper-V Username username to access the Hyper-V (in Down-level login (domain\username) or UPN format)
 - Hyper-V Password Password to access Hyper-V for performing inventory of the resources.

Select "Accept Self signed certificate" and click Continue.

ONTAP Storage system credentials



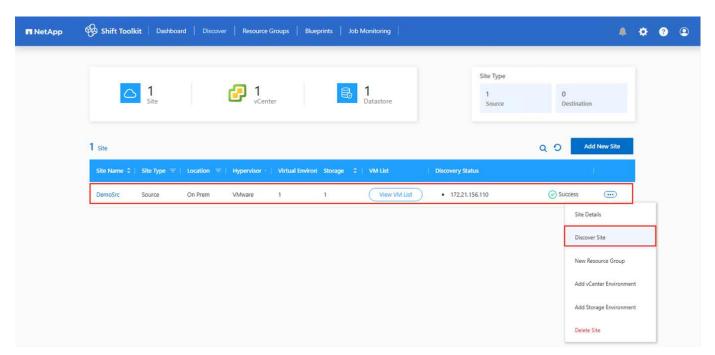
Once added, Shift toolkit will perform an automatic discovery and display the VMs along with the relevant metadata information. Shift toolkit will automatically detect the networks and vLANs used by the VMs and will populate them.



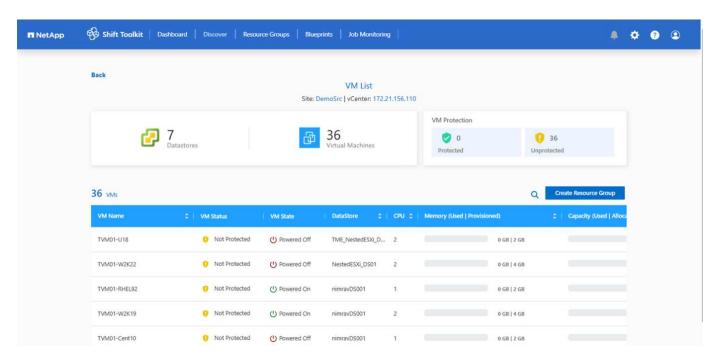
If any modifications are made to the source site, ensure to run the discovery to fetch the latest information. This can be done by clicking on 3 dots against the site name and click on "Discover Site".



The VM inventory is auto-refreshed every 24 hours.

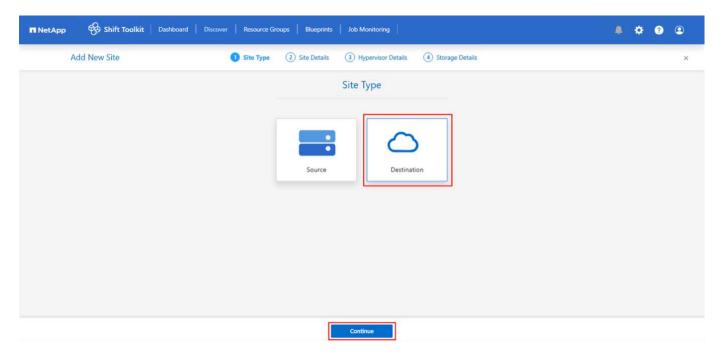


To view the discovery data for a specific source hypervisor, go to the dashboard, click on "View VM List" against the appropriate site name. The page will display the VM inventory along with the VM attributes.



Next step is to add the destination hypervisor. To add, click on "Add New Site" and select "Destination".

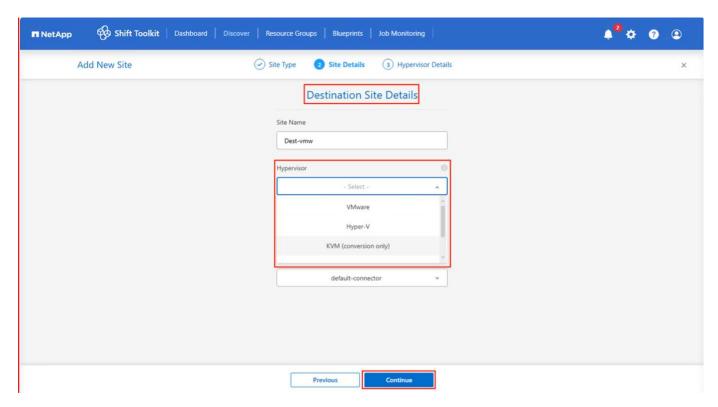
Destination



- · Destination Site Details
 - Site Name Provide a name for the site
 - Hypervisor Choose the appropriate target platform from the following options:
 - VMware
 - Hyper-V
 - OpenShift
 - OLVM

- KVM (conversion only)
- Site Location Select the default option
- Connector Select the default selection

Once filled, click Continue.



Based on the hypervisor selection, fill in the necessary details.

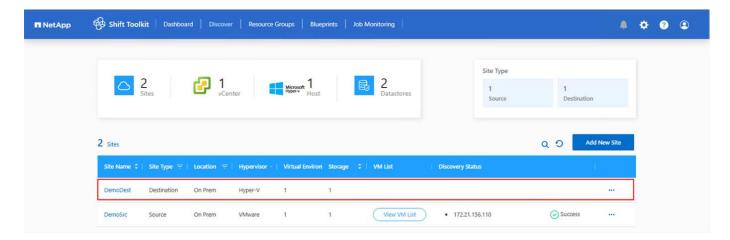
- · Destination hypervisor details
 - Respective hypervisor manager Endpoint IP address or FQDN
 - Username username to access (in UPN format: username@domain.com or domain\administrator)
 Password Password to access for performing inventory of the resources.

Select "Accept Self signed certificate".

· Once done, click "Create Site"



The source and destination storage system should be the same as the disk format conversion happens at the volume level and within the same volume.



Next step is to group the required VMs into their migration groups as resource groups.

Resource Groupings

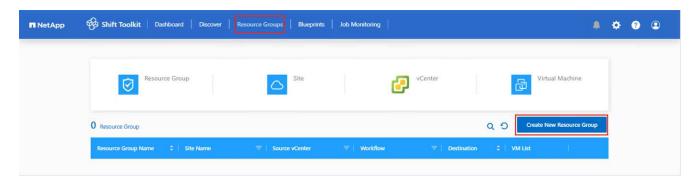
Once the platforms have been added, group the VMs you want to migrate or convert into resource groups. Shift toolkit resource groups allow you to group set of dependent VMs into logical groups that contain their boot orders and boot delays.



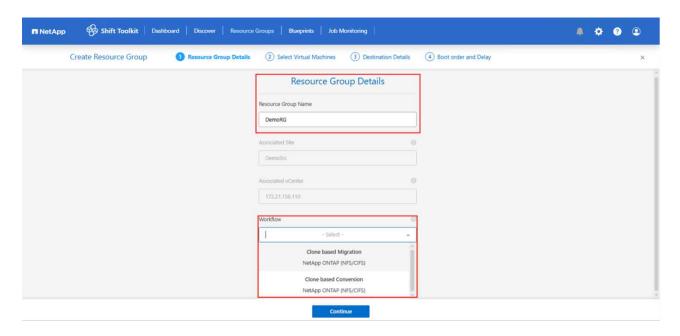
Ensure the Qtrees are provisioned (as mentioned in the pre-requisite section) before creating the resource groups.

To start creating resource groups, click on the "Create New Resource Group" menu item.

1. Access Resource groups, click on "Create New Resource Group".



- 2. On the "New resource group", select the Source site from the dropdown and click "Create"
- 3. Provide Resource Group Details and select the workflow. The workflow provides two options
 - a. Clone based Migration performs end to end migration of the VM from source hypervisor to destination hypervisor.
 - b. Clone based Conversion Performs conversion of the disk format to the selected hypervisor type.

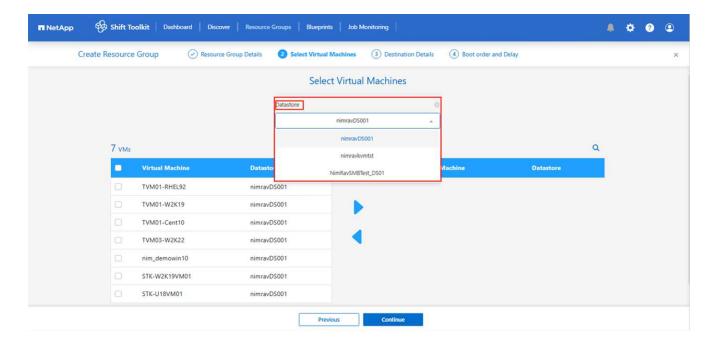


- Click on "Continue"
- 5. Select appropriate VMs using the search option. The default filter option is "Datastore".
 - (1)

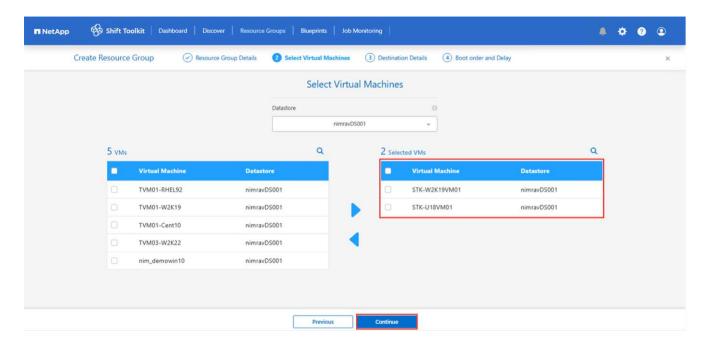
Move the VMs to convert or migrate to a designated datastore on a newly created ONTAP SVM before conversion. This helps isolating the production NFS datastore and the designated datastore can be used for staging the virtual machines.



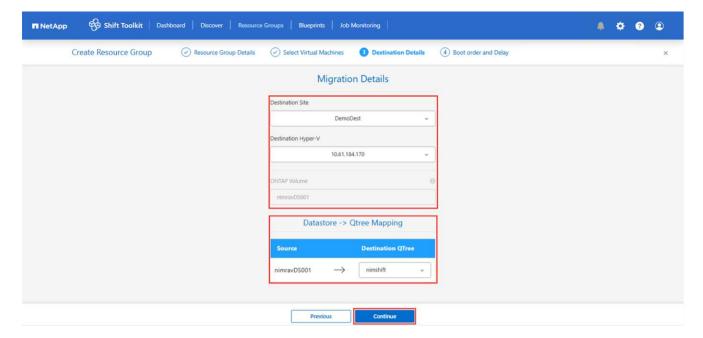
For OpenShift environment, VMDKs should be relocated to their corresponding volumes to replicate the PVC (Persistent Volume Claim) structure using ONTAP NAS storage driver. In future releases, additional enhancements will be incorporated to leverage ONTAP NAS economy driver.



The datastore dropdown in this context will only show NFSv3 datastores. NFSv4 datastores will not be displayed.



6. Update the migration details by selecting "Destination Site", Destination Hypervisor entry" and Datastore to Qtree or storage class mapping.





Make sure that the destination path (where the converted VMs are stored) is set to a qtree when converting VMs. Set the destination path to the appropriate qtree.



Multiple qtrees can be created and used for storing the converted VM disks accordingly.

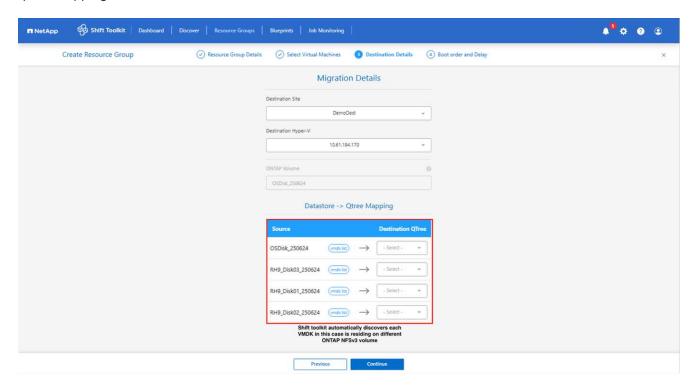


Virtual machines can be running on spanned datastores, and Shit toolkit will automatically detect them, however a qtree should be mapped for each volume.

Ability to migrate VMs with spanned VMDKs across multiple volumes

The Shift toolkit UI automatically selects all the spanned volumes that is part of a VM or VMs that are selected for that specific RG. This will list all the volumes in the RG page where we do the datastore –

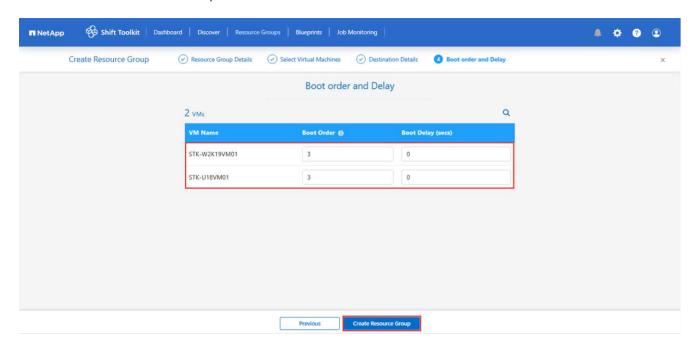
qtree mapping.



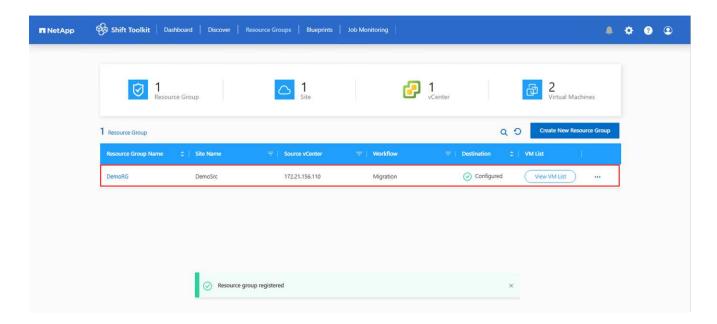
7. Select the Boot Order and Boot delay (secs) for all the selected VMs. Set the order of power on sequence by selecting each virtual machine and setting up the priority for it. 3 is the default value for all virtual machines.

Options are as follows:

- 1 The first virtual machine to power on
- 3 Default
- 5 The last virtual machine to power on



8. Click on "Create Resource Group".





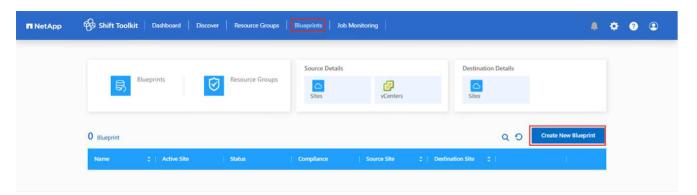
In the event of the need to modify the resource group so as to add or remove virtual machines, use this option against the resource group name and select "Edit Resource Group".

Blueprints

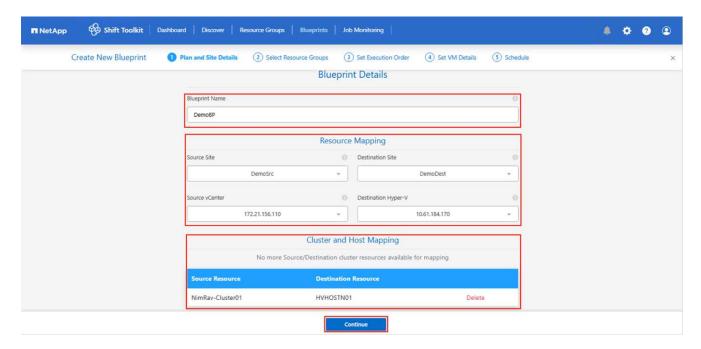
To migrate or convert virtual machines, a plan is necessary. Select the source and destination hypervisor platforms from the drop down and pick the resource groups to be included in this blueprint, along with the grouping of how applications should be powered on (i.e. domain controllers, then tier-1, then tier-2, etc). These are often called as migration plans as well. To define the blueprint, navigate to the "Blueprints" tab and click on "Create New Blueprint".

To start creating blueprint, click on the "Create New Blueprint".

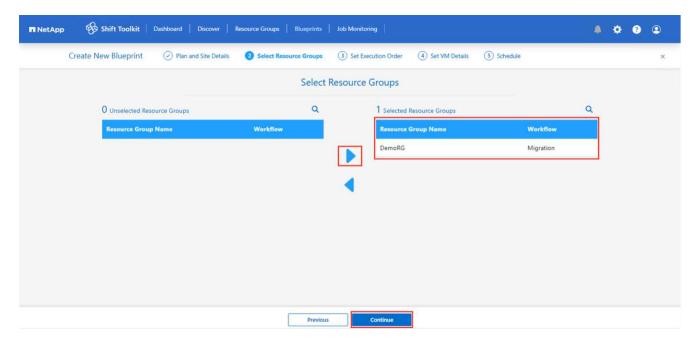
1. Access Blueprints, click on "Create New Blueprint".



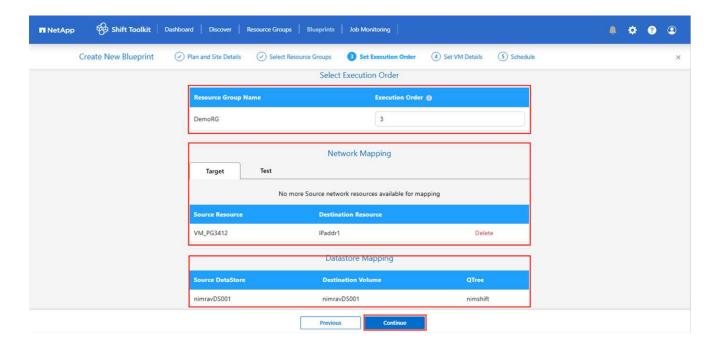
- 2. On the "New Blueprint", provide a name for plan and add necessary host mappings by selecting Source Site > associated vCenter, Destination Site and the associated hypervisor.
- Once mappings are done, select the cluster and host mapping.
 In the example below, Hyper-V is shown as the target. The hypervisor option displayed will vary based on the selected source site.



4. Select Resource Group Details and click on "Continue"

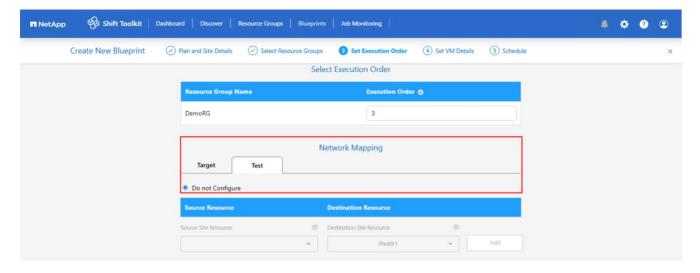


- 5. Set Execution Order for Resource Group. This option enables to select the sequence of operations when multiple resource groups exist.
- 6. After completing the previous steps, select Network Mapping and assign it to the appropriate network map. Ensure that the virtual switches, network profiles, or operators are already provisioned on the target hypervisor.





For test migration, "Do no configure Network" is the default selection and Shift toolkit does not perform IP address assignment. Once the disk is converted and virtual machine is bought on respective hypervisor side, manually assign the bubble network switches to avoid any colliding with production network.



7. Based on the selection of VMs, storage mappings will be automatically selected.



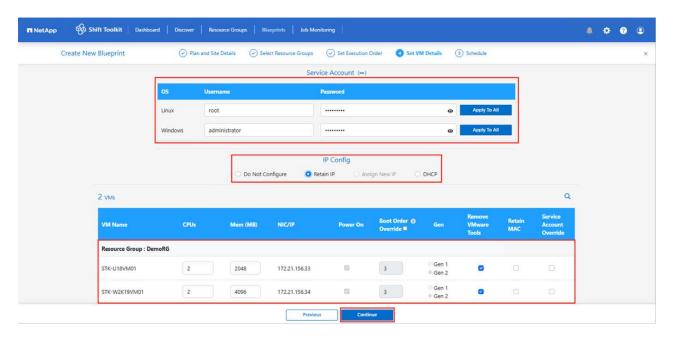
Make sure the qtree is provisioned beforehand and the necessary permissions are assigned so the virtual machine can be created and powered ON.

NOTE: In case of OpenShift, the PVCs are created using Trident CSI and there is no need to pre-create qtrees.

- 8. Under VM details, provide service account and valid user credentials for each OS type. This is used to connect to the virtual machine to create and run certain scripts that are necessary for removing VMware tools and backing up IP configuration details.
 - a. For Windows based OS, it is recommended to use a user with local administrator privileges. Domain credential can also be used, however ensure there is a user profile existing on the VM before conversion, otherwise domain credentials won't work as it would look for domain authentication when

there is no network connected.

b. In case of Linux distribution-based guest VMs, provide a user that can execute sudo commands without password meaning the user should be part of the sudoers list or added as a new configuration file to the /etc/sudoers.d/ folder.



In the recent release, the Shift Toolkit introduced greater flexibility in virtual machine preparation. By default, the toolkit automates VM preparation by deploying OS-specific scripts to:

- * Remove VMware Tools
- * Back up IP settings for reassignment based on the selected Blueprint

+

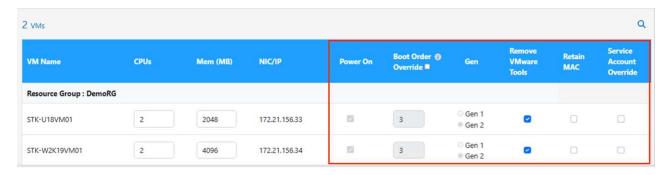
With the new enhancement, users can now override the default prepareVM tasks, enabling to execute custom scripts for manual VM preparation including IP assignment. This provides more control for environments with unique configuration or compliance requirements.

- 9. Again under VM details, select the relevant IP config option. By default, "Do not configure" is selected.
 - a. To migrate VMs with the same IPs from the source system, select "Retain IP".
 - b. To migrate VMs using static IPs in the source system and to assign DHCP on the target VMs, then select "DHCP".

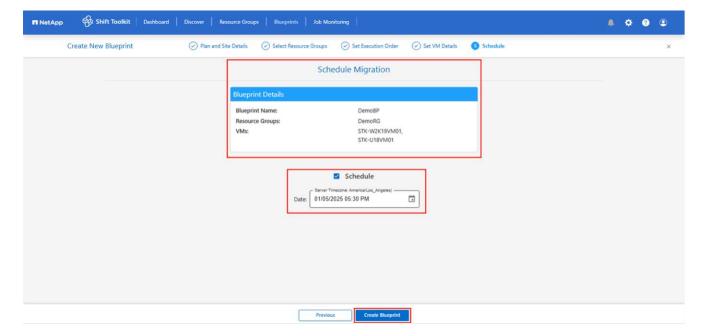
Make sure the following requirements are met for this functionality to work:

- Ensure the VMs are powered on during the prepareVM phase and up to the scheduled migration time.
- For VMware VMs, ensure that VMware Tools are installed.
- For Hyper-v as the source hypervisor, ensure Integration services is enabled and configured.
- For OLVM and OpenShift as the target hypervisor, ensure to mount the virtIO ISO file to the Windows VMs.
- Ensure the preparation script is run on the source VM by an account with administrator privileges on windows OS and with sudo privileges with no password option on Linux based distribution OS to create cron jobs.
- 10. The next step is VM configuration.

- Optionally resize the VMs CPU/RAM parameters which can be very helpful for resizing purposes.
- Boot Order override: Also modify the Boot Order and Boot delay (secs) for all the selected VMs across
 the resource groups. This is an additional option to modify the boot order if any changes required from
 what was selected during Resource group boot order selection. By default, the boot order selected
 during resource group selection is used, however any modifications can be done at this stage.
- Power ON: Uncheck this option if workflow should not power ON the virtual machine. Default option is ON meaning the VM will be powered ON.
- Remove VMware tools: Shift toolkit removes VMware tools after the conversion. This option is selected
 by default. This is an be unselected if the plan is to execute customer's own customized scripts.
- Generation: Shift toolkit uses the following rule of thumb and defaults to the appropriate one- Gen1 > BIOS and Gen2 > EFI. No selection is possible for this option.
- Retain MAC: The MAC address of the respective VMs can be retained to overcome licensing challenges for those applications relying on MAC.
- Service Account override: This option allows to specify a separate service account if the global one cannot be used.



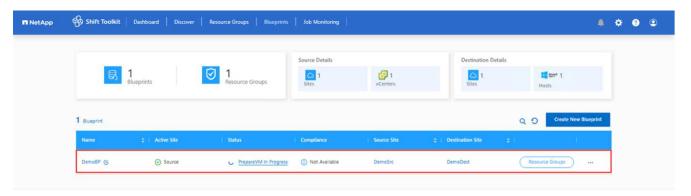
- 11. Click "Continue".
- 12. In the next step, schedule the migration by selecting the checkbox to set the date and time. Make sure all the virtual machines (VMs) are prepared and powered off before the scheduled date. Once done, click on "Create Blueprint".





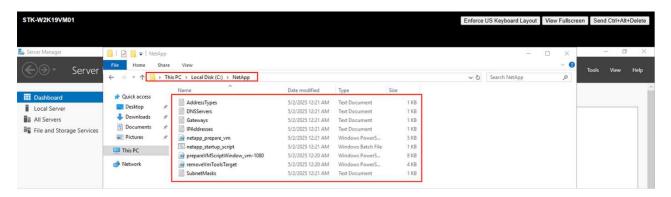
While scheduling, choose a date that is at least 30 minutes ahead of the current Shift VM time. This is to ensure the workflow gets enough time to prepare the VMs within the resource group.

13. Once the blueprint is created, a prepareVM job is initiated and it automatically runs scripts on the source VMs to prepare them for migration



This job runs a script using invoke-VMScript method to copy the necessary scripts for removing VMware tools and backing up network configuration details, including IP address, routes, and DNS information, which will be used to maintain the same settings on the target VM.

a. For Windows-based operating systems, the default location where the preparation scripts are stored is the "C:\NetApp" folder.



b. For Linux-based VMs, the default location where the preparation scripts are stored is /NetApp and the /opt directory.



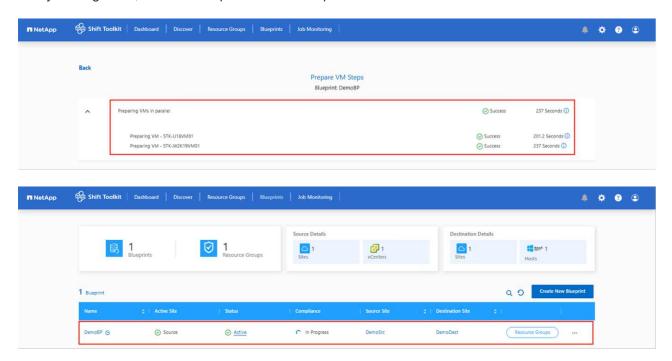


For a Linux source VM running CentOS or Red Hat, Shift toolkit is intelligent to automatically install the necessary Hyper-V drivers. These drivers must be present in the source VM before the disk conversion to ensure the VM can boot successfully after the conversion.



For detailed information, refer to System stuck in dracut after the migration of a RHEL VM to hyper-v.

Once the prepareVM job completes successfully (as shown in the screenshot below), the VMs are ready for migration, and the blueprint status will update to "Active."



Migration will now happen at the set time or can be started manually by clicking on Migrate option.

Migrate VMs using Shift Toolkit

Migrate VMs using the Shift Toolkit

Use the Shift Toolkit to migrate VMs between virtualization platforms. The process involves preparing the VMs, converting disk formats, and configuring network settings on the target environment.

Supported migrations

The Shift Toolkit provides flexibility in multi-hypervisor environments by supporting bidirectional migration between the following hypervisors:

- VMware ESXi to Microsoft Hyper-V
- Microsoft Hyper-V to VMware ESXi
- VMware ESXi to Oracle Linux Virtualization Manager (OLVM)
- VMware ESXi to Red Hat OpenShift Virtualization

Migration workflow

After creating a blueprint, you can initiate the migration process. During migration, the Shift Toolkit performs a series of steps to convert disk formats and create virtual machines on the target host as defined in the

blueprint.

The Shift Toolkit performs the following steps during migration:

- 1. Delete existing snapshots for all VMs in the blueprint
- 2. Trigger VM snapshots for the blueprint at the source
- 3. Trigger volume snapshot before disk conversion
- 4. Clone and convert VMDK to VHDx format for all VMs
- 5. Power on VMs in the protection group at the target
- 6. Register the networks on each VM
- 7. Remove VMware Tools and assign IP addresses using trigger scripts or cron jobs depending on the OS type

Network tips and considerations

Consider the following network requirements and behaviors when planning your migration. The Shift Toolkit automatically copies network settings from source VMs and reapplies them to migrated VMs, but interface naming and network adapter configurations may vary between Windows and Linux systems.

General requirement

• Ensure static IP addresses are available and not assigned to another VM

Windows VMs

- The prepare script copies network configuration details (IP address space, gateway address, DNS servers)
- The trigger script reapplies network settings during migration for single or multiple NICs based on blueprint mapping
- After migration, Windows Device Manager may display old network adapter information from pre-migration, which doesn't affect the new adapter or cause IP conflicts
- Upgrade to v4.0 to automatically remove orphaned network devices from the registry and Device Manager

Linux VMs

- The prepare script copies network configuration details (IP address space, routes, DNS servers, network device names)
- The script identifies the Linux distribution's networking type and applies IP settings accordingly
- The network reassignment script is configured as a cron job using crontab and triggered on boot
- The script reapplies network settings for single or multiple NICs based on blueprint mapping

Interface naming

- Converted VMs may have interface names like eth0 or ensp0 instead of the source interface names (for example, ens192 or ens33)
- The script updates network configuration details to match new interface names
- If predictable names are used with proper udev matching rules and the interface name is retained on the target hypervisor, the script skips network configuration, removes VMware Tools, and reboots the VM



The Shift Toolkit allows overriding network preparation, enabling administrators to run custom scripts for IP assignment or other configurations.

Supported networking mechanisms

- NetworkManager
- Netplan
- ifconfig
- wicked

The Shift Toolkit retains IP addresses as specified in the blueprint.

Migration phases

The following are the migration phases you will follow to migrate VMs using the Shift Toolkit.

- 1. **Prepare VM**: Prepare VMs for migration and verify that all prerequisites are completed.
- 2. Migrate and validate: After preparation is complete, migrate VMware VMs to the target hypervisor.

After migration completes, verify that VMs boot successfully and data has migrated properly.

3. **Test the migration**: Test migration simulates the migration by converting the VMDK to the appropriate format and creating VMs using the converted virtual disk file on the qtree.

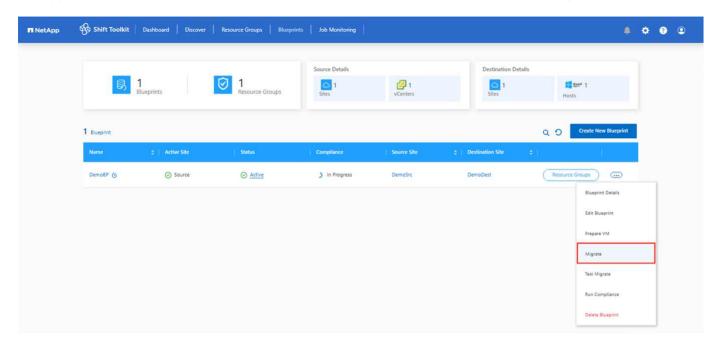
Test migration does not include network mapping configuration, which should be performed manually to a test network.



The Shift Toolkit does not alter the source VM except for copying scripts needed for VM preparation. This allows for swift rollback in case of conversion failures.

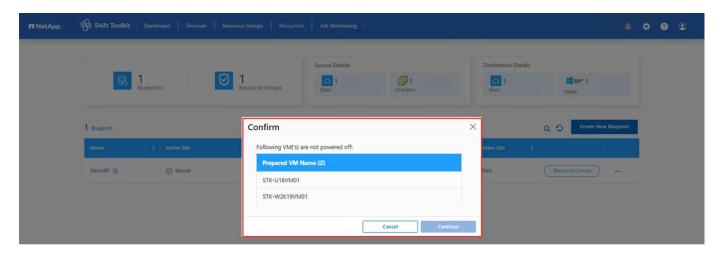
Execute a migration

To trigger the migration workflow with the configuration specified in the blueprint, click Migrate.



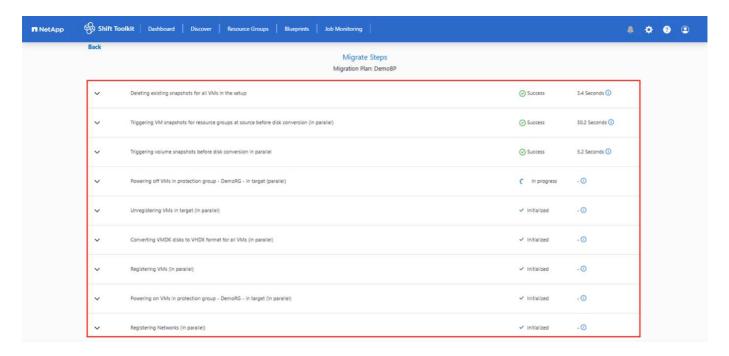
Once initiated, the workflow activates and the conversion process follows the outlined steps to register the VM. If VMs within the blueprint are not powered off, the Shift Toolkit prompts for a graceful shutdown before

proceeding.

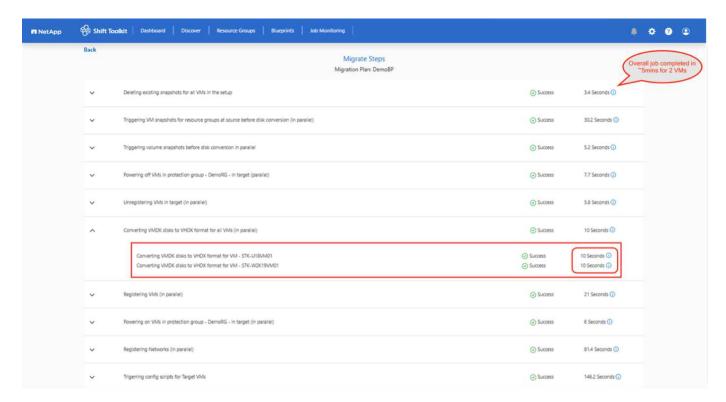


(i)

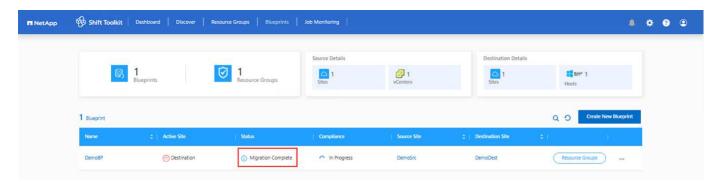
NetApp recommends triggering no more than ten conversions in parallel from the same source to the same destination.



The conversion of VMDK to any file format completes in seconds, making this the fastest option available. This approach helps reduce VM downtime during migration.



Once the job completes, the blueprint status changes to "Migration Complete".



Migrate VMs from VMware ESXi to Microsoft Hyper-V using the Shift Toolkit

Migrate VMs from VMware ESXi to Microsoft Hyper-V using the Shift Toolkit by preparing VMs, converting disk formats, and configuring the target environment.

The Shift Toolkit enables VM migration between virtualization platforms through disk format conversion and network reconfiguration on the destination environment.

Before you begin

Verify that the following prerequisites are met before starting the migration.

Hyper-V requirements

- Hyper-V hosts configured as standalone hosts or failover cluster
- · Hyper-V user account with administrator privileges
- · Hyper-V hosts are network reachable with up-to-date DNS entries
- Virtual switches configured with appropriate trunking

- Virtual switch type "External" for network selection
- · NFS share (for VMs to be converted) and destination share (for converted VMs) on the same volume
- SMB Constrained Delegation configured using Enable-SmbDelegation to avoid access denied errors
- SMB 3.0 enabled (default)
- · Continuously available property enabled for SMB shares
- Export policies for SMB disabled on the storage virtual machine (SVM)



SCVMM is not a supported endpoint for migration in the current release.

• The Hyper-V FCI and host discovery relies on DNS resolution. Ensure hostnames are resolvable from Shift Toolkit VM. If resolution fails, update the host file (C:\Windows\System32\drivers\etc\hosts) and retry the discovery operation.

VMware requirements

- VM VMDKs are placed on NFSv3 volume (all VMDKs for a given VM should be part of the same volume)
- VMware tools are running on guest VMs
- VMs to be migrated are in a RUNNING state for preparation
- VMs must be powered off before triggering migration
- VMware tools removal happens on the destination hypervisor once VMs are powered on

Guest VM requirements

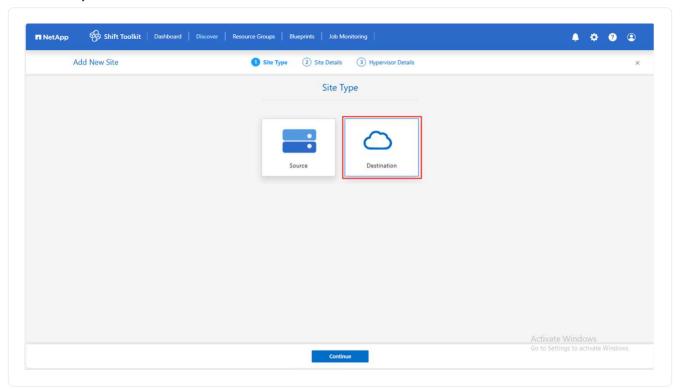
- For Windows VMs: Use local administrator credentials (domain credentials can also be used, however ensure a user profile exists on the VM before conversion)
- For Linux VMs: Use a user with permissions to execute sudo commands without password prompt (user should be part of the sudoers list or added to /etc/sudoers.d/ folder)

Step 1: Add the destination site (Hyper-V)

Add the destination Hyper-V environment to the Shift Toolkit.

Steps

1. Click Add New Site and select Destination.



2. Enter the destination site details:

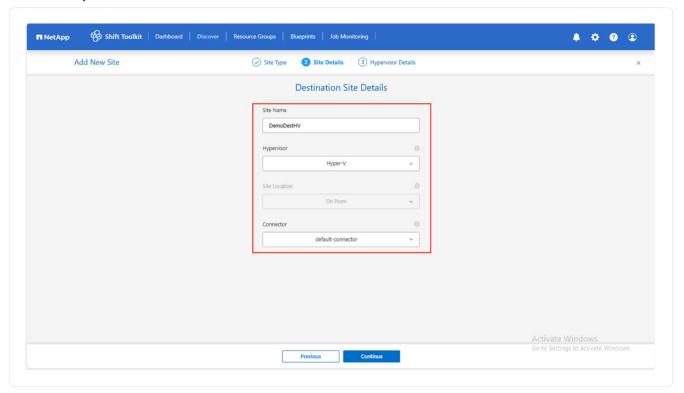
• Site Name: Provide a name for the site

• **Hypervisor**: Select Hyper-V as the target

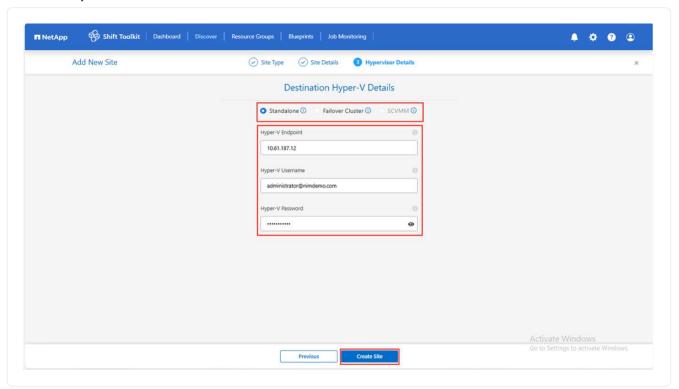
• Site Location: Select the default option

• Connector: Select the default selection

3. Click Continue.

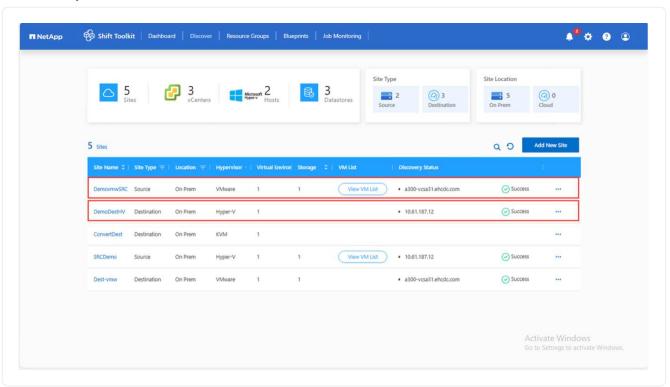


- 4. Enter the destination Hyper-V details:
 - Hyper-V Standalone or failover cluster manager: IP address or FQDN
 - **Username**: Username to access (in UPN format: username@domain.com or domain\administrator)
 - Password: Password to access Hyper-V host or FCI instance for performing inventory of the resources
- 5. Select **Accept Self signed certificate** and click **Continue**.



6. Click Create Site.

Show example





The source and destination storage system should be the same as the disk format conversion happens at the volume level and within the same volume.

Step 2: Create resource groups

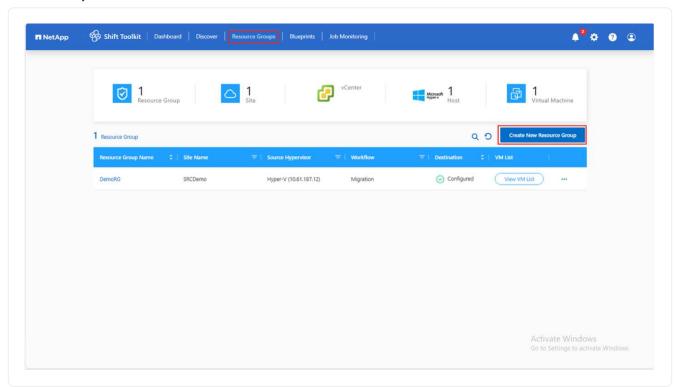
Organize VMs into resource groups to preserve boot order and boot delay configurations.

Before you begin

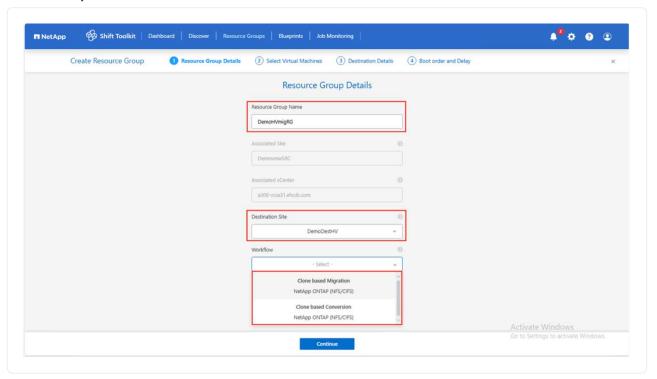
- Ensure gtrees are provisioned as specified in the prerequisites
- Move VMs to a designated datastore on a newly created ONTAP SVM before conversion to isolate production NFS datastores from the staging area

Steps

1. Navigate to Resource Groups and click Create New Resource Group.



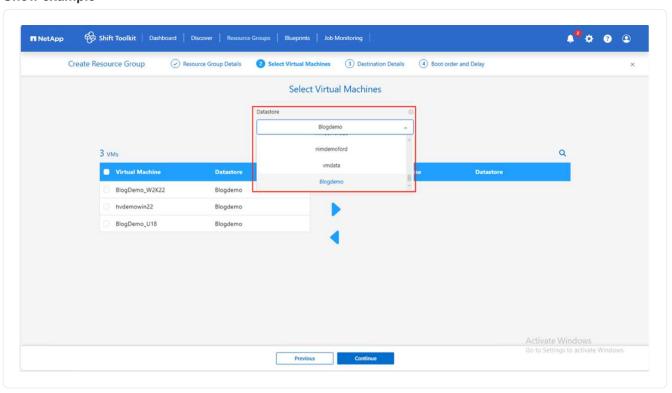
- 2. Select the **Source site** from the dropdown and click **Create**.
- 3. Provide resource group details and select the workflow:
 - Clone based Migration: Performs end-to-end migration from source to destination hypervisor
 - Clone based Conversion: Converts disk format to the selected hypervisor type

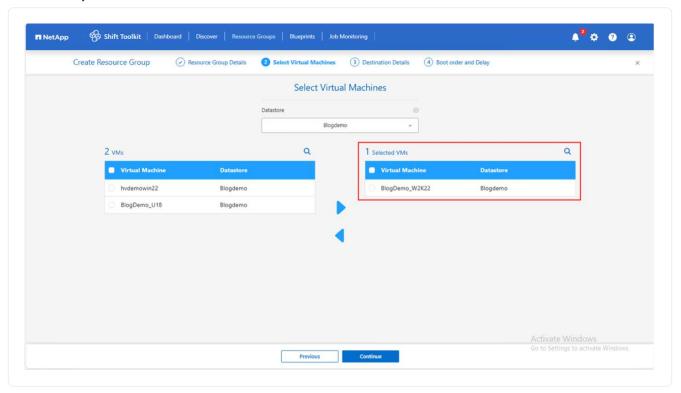


- 4. Click Continue.
- 5. Select VMs using the search option (default filter is "Datastore").

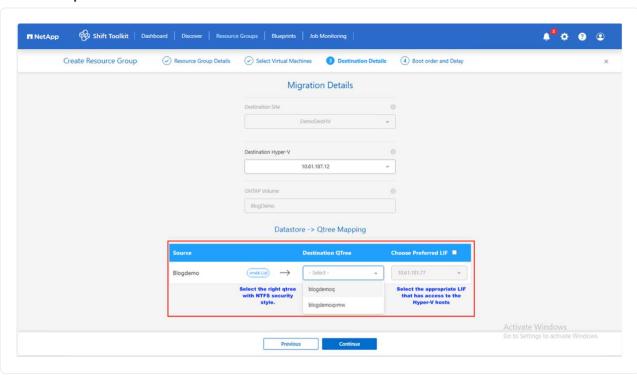


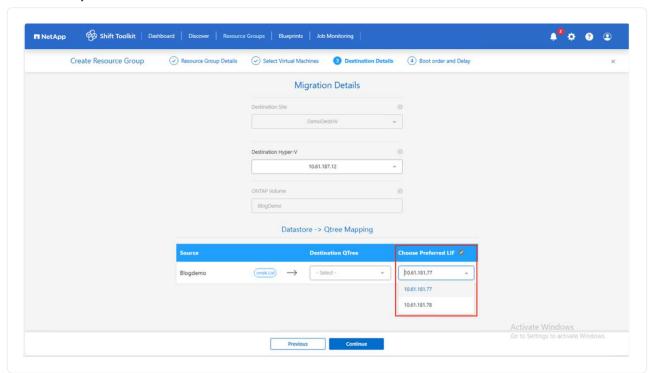
The datastore dropdown only shows NFSv3 datastores. NFSv4 datastores are not displayed.





- 6. Update migration details:
 - Select Destination Site
 - Select Destination Hyper-V entry
 - Configure Datastore to Qtree mapping

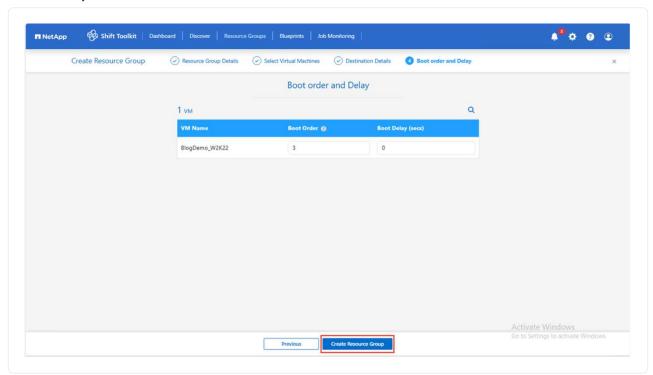






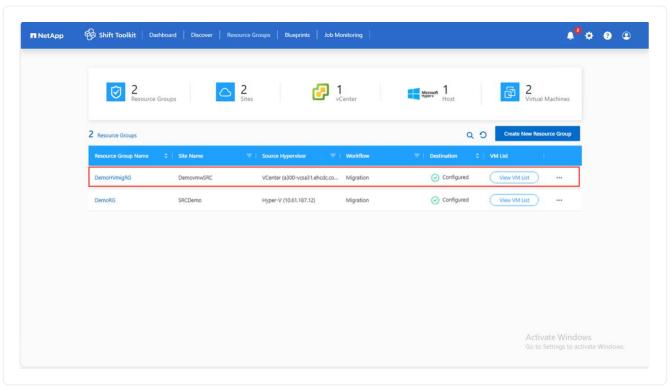
Ensure the destination path (where the converted VMs are stored) is set to a qtree when converting VMs from ESXi to Hyper-V. Multiple qtrees can be created and used for storing converted VM disks.

- 7. Configure boot order and boot delay for all selected VMs:
 - 1: First VM to power on
 - 3: Default
 - **5**: Last VM to power on



8. Click Create Resource Group.

Show example



Result

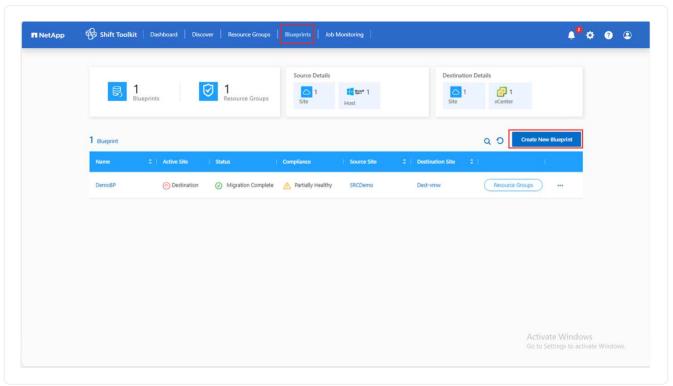
The resource group is created and ready for blueprint configuration.

Step 3: Create a migration blueprint

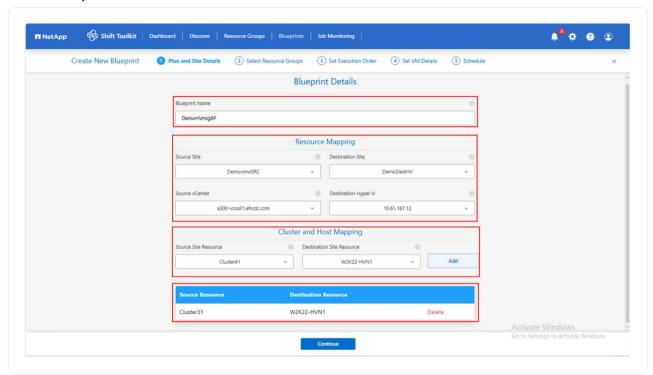
Create a blueprint to define the migration plan, including platform mappings, network configuration, and VM settings.

Steps

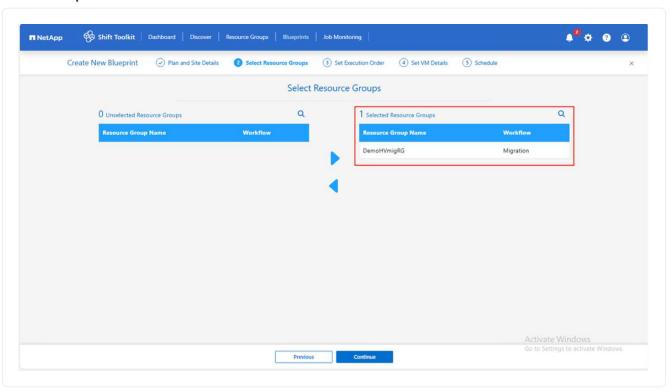
1. Navigate to Blueprints and click Create New Blueprint.



- 2. Provide a name for the blueprint and configure host mappings:
 - Select Source Site and associated vCenter
 - Select **Destination Site** and associated Hyper-V target
 - · Configure cluster and host mapping



3. Select resource group details and click Continue.

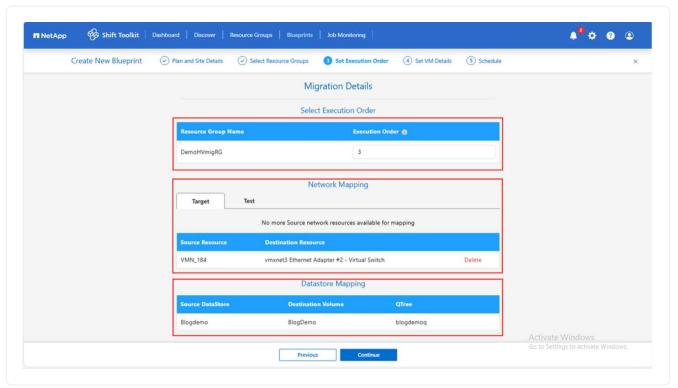


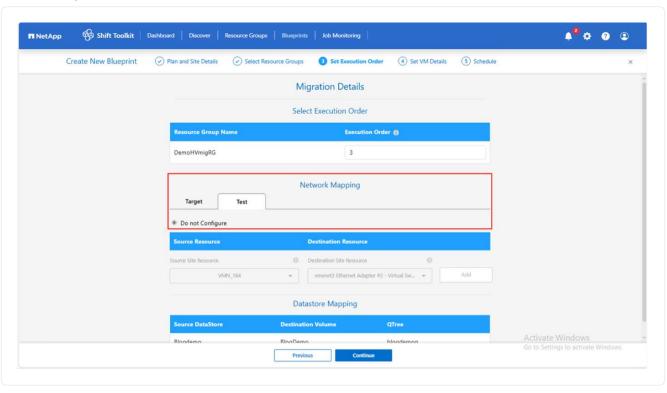
- 4. Set execution order for resource groups if multiple groups exist.
- 5. Configure network mapping to appropriate virtual switches.



Virtual switches should already be provisioned within Hyper-V. On Hyper-V side, the virtual switch type "External" is the only supported option for network selection. For test migration, select "Do not configure Network" to avoid production network conflicts; manually assign network settings after conversion.

Show example





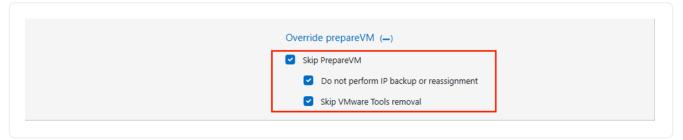
6. Review storage mappings (automatically selected based on VM selection).



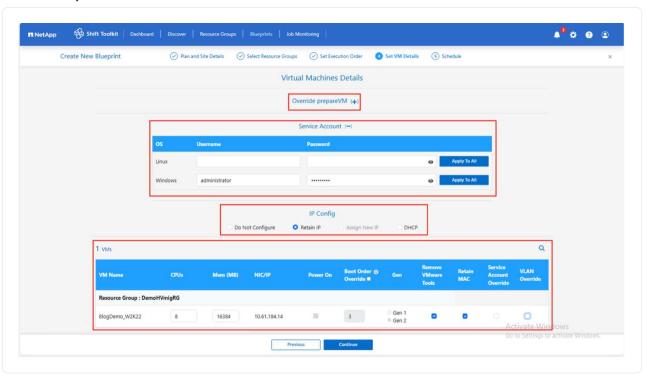
Ensure the qtree is provisioned beforehand and necessary permissions are assigned so the virtual machine can be created and powered on from SMB share.

7. Configure the prepareVM override option if needed. This option is useful when you need to skip VM preparation by the Shift Toolkit and instead perform those tasks using custom scripts. It also enables customization of the IP address to meet specific environment requirements.

Show example



- 8. Under VM details, select configuration details and provide service account credentials for each OS type:
 - Windows: Use a user with local administrator privileges (domain credentials can also be used, however ensure a user profile exists on the VM before conversion)
 - Linux: Use a user that can execute sudo commands without password prompt (user should be part of the sudoers list or added to /etc/sudoers.d/ folder)



- 9. Configure IP settings:
 - · Do not configure: Default option
 - Retain IP: Keep same IPs from source system

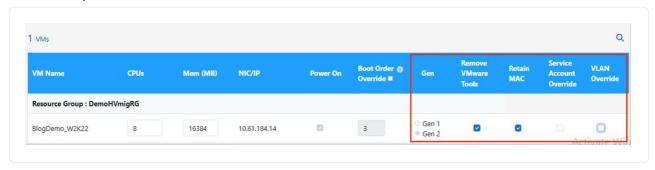
DHCP: Assign DHCP on target VMs

Ensure VMs are powered on during prepareVM phase, VMware Tools are installed, and preparation scripts run with proper privileges.

10. Configure VM settings:

- Resize CPU/RAM parameters (optional)
- Modify boot order and boot delay
- Power ON: Select to power on VMs after migration (default: ON)
- Remove VMware tools: Remove VMware Tools after conversion (default: selected)
- VM Firmware: Gen1 > BIOS and Gen2 > EFI (automatic)
- Retain MAC: Keep MAC addresses for licensing requirements
- Service Account override: Specify separate service account if needed
- VLAN override: Select correct tagged VLAN name when target hypervisor uses different vLAN name

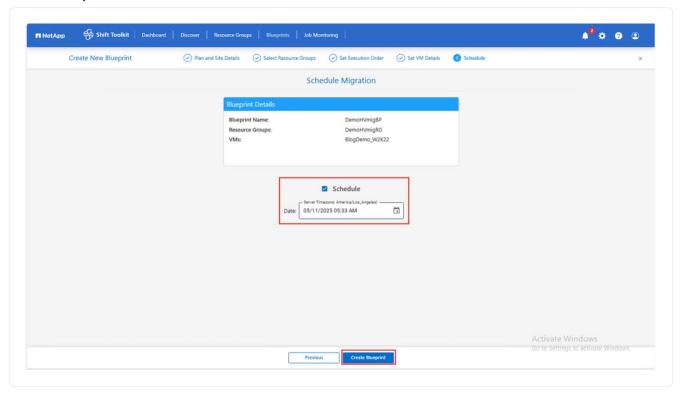
Show example



- 11. Click Continue.
- 12. Schedule the migration by selecting a date and time.



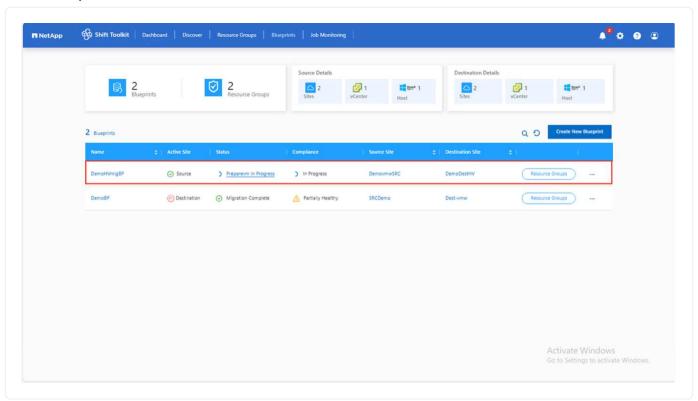
Schedule migrations at least 30 minutes ahead to allow time for VM preparation.



13. Click Create Blueprint.

Result

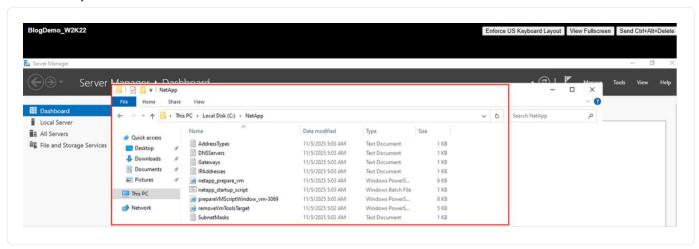
The Shift Toolkit initiates a prepareVM job that runs scripts on source VMs to prepare them for migration.



The preparation process:

- Injects scripts to add drivers (RHEL/CentOS, Alma Linux), remove VMware tools, and backup IP/route/DNS information
- Uses invoke-VMScript to connect to guest VMs and execute preparation tasks
- For Windows VMs: Stores scripts in C:\NetApp
- For Linux VMs: Stores scripts in /NetApp and /opt

Show example



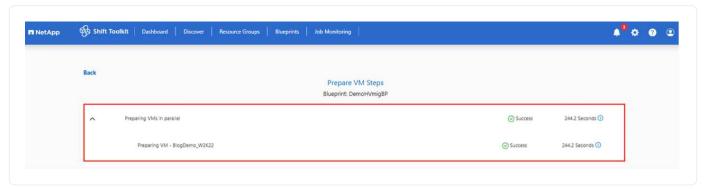
Show example



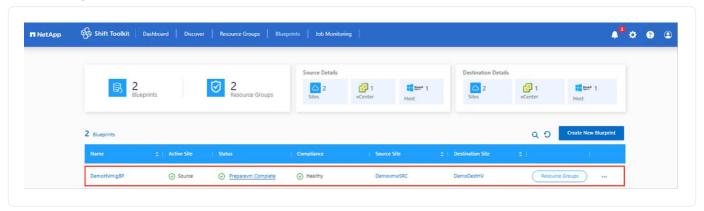


For Linux source VMs running CentOS or Red Hat, Shift Toolkit automatically installs necessary Hyper-V drivers before disk conversion to ensure successful boot after conversion. For detailed information, refer to System stuck in dracut after the migration of a RHEL VM to hyper-v.

When prepareVM completes successfully, the blueprint status updates to "Active." Migration will now happen at the scheduled time or can be started manually by clicking the **Migrate** option.



Show example



Step 4: Execute the migration

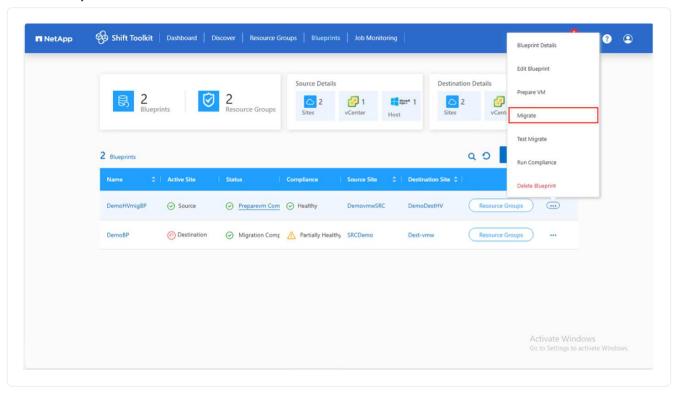
Trigger the migration workflow to convert VMs from VMware ESXi to Microsoft Hyper-V.

Before you begin

- All VMs are gracefully powered off according to the planned maintenance schedule
- Ensure the Shift VM is part of the domain
- Ensure CIFS share is configured with appropriate permissions
- The qtree used for migration or conversion has the right security style
- As a quick test, try creating a VM using Hyper-V Manager from any Hyper-V host within the cluster and place the VHDX on the CIFS share

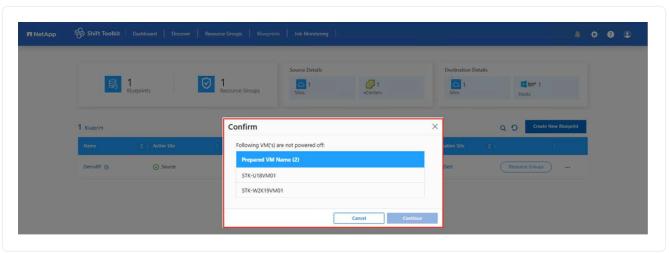
Steps

1. On the blueprint, click Migrate.



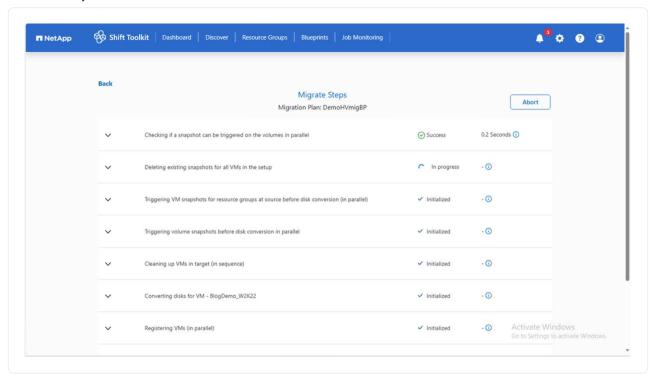
2. If VMs are not powered off, the Shift Toolkit will prompt for a graceful shutdown before proceeding.

Show example

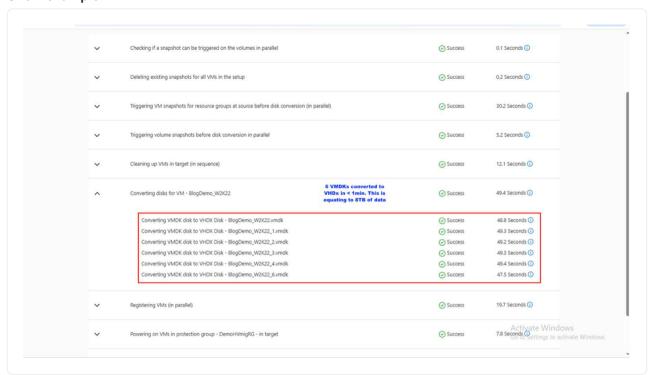


- 3. The Shift Toolkit takes the following actions:
 - Deletes existing snapshots for all VMs in the blueprint
 - Triggers VM snapshots at the source
 - Triggers volume snapshot before disk conversion
 - Converts VMDK to VHDx format for all VMs

The conversion happens in seconds, making this the fastest migration approach and reducing VM downtime.



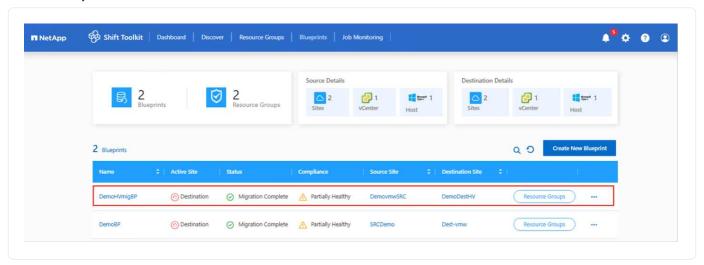
Show example

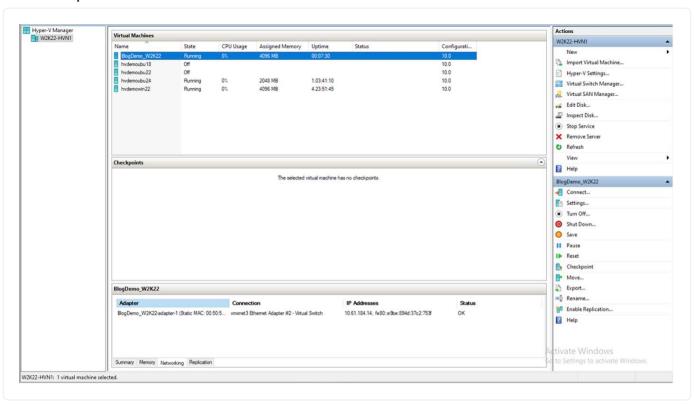


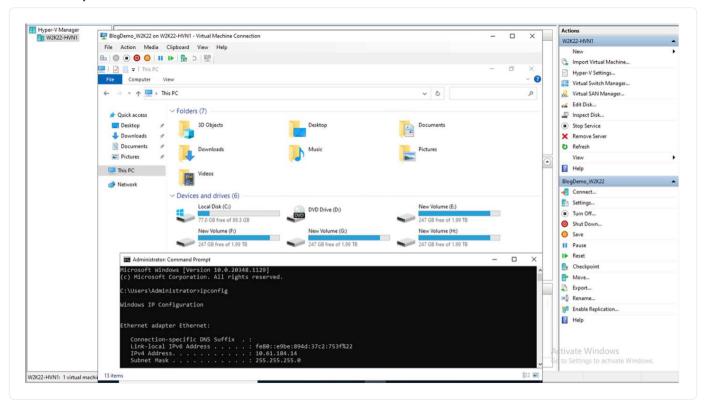
- $\,{}^{_{\odot}}$ Powers on VMs at the target
- Registers networks on each VM
- Removes VMware tools and assigns IP addresses using trigger scripts or cron jobs

Result

When the job completes, the blueprint status changes to "Migration Complete."







- (i)
- No more than ten conversions should be triggered in parallel from the same ESXi source to the same Hyper-V destination.
- (i)
- If there are failures, enable delegation using any authentication protocol.
- <u>(i)</u>

After migration when Windows VMs are powered on, Shift Toolkit uses PowerShell Direct to connect to Windows-based guest VMs regardless of network configuration or remote management settings.

(1)

After conversion, all VM disks on Windows OS except the OS disk will be offline because the NewDiskPolicy parameter is set to offlineALL on VMware VMs by default. Run this PowerShell command to fix: Set-StorageSetting -NewDiskPolicy OnlineAll

(i)

Shift Toolkit uses cron jobs that execute on boot for Linux-based distributions. No SSH connections are created for Linux-based VMs once they are brought on Hyper-V hosts.

Video demonstration

The following video demonstrates the process outlined in this solution.

Migrate VMs from ESXi to Hyper-V using Shift Toolkit

Migrate VMs from Microsoft Hyper-V to VMware ESXi using the Shift Toolkit

Migrate VMs from Microsoft Hyper-V to VMware ESXi using the Shift Toolkit by configuring source and destination sites, creating resource groups and blueprints, and

executing the migration workflow.

The Shift Toolkit enables direct VM conversion between hypervisors without creating additional disk copies, delivering copy-less migration with minimal downtime for both Windows and Linux virtual machines.

Before you begin

Verify that the following prerequisites are met before starting the migration.

VMware requirements

- vCenter and ESXi hosts are configured
- vCenter server account (RBAC user) with minimum required privileges
- vCenter and ESXi hosts are reachable from the Shift Toolkit and DNS entries are current
- Distributed port groups are configured with appropriate VLAN IDs (standard port groups are not supported)
- NFS share (for storing migrated VMs) and source share (for VMs to be migrated) reside on the same volume

Hyper-V requirements

- VM VHDx files are placed on an SMB share
 - If VMs are on a Cluster Shared Volume (CSV), perform a live migration to an SMB share
- Hyper-V integration services are enabled and running on guest VMs
- VMs to be migrated are in a RUNNING state for preparation
- · VMs must be powered off before triggering migration

Guest VM requirements

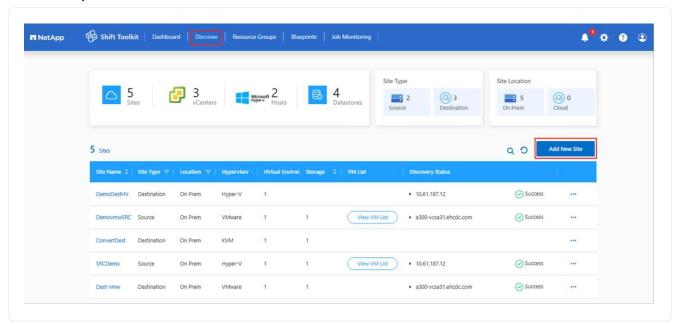
- For Windows VMs: Use local administrator credentials or domain credentials with an existing user profile on the VM
- For Linux VMs: Use a user with permissions to execute sudo commands without password prompt
- Shift Toolkit uses PowerShell Direct for Windows VMs and SSH for Linux VMs

Step 1: Add the source site (Hyper-V)

Add the source Hyper-V environment to the Shift Toolkit.

Steps

- 1. Open the Shift Toolkit in a supported browser and log in with the default credentials.
- 2. Navigate to **Discover > Add Sites**.



- 3. Click Add New Site and select Source.
- 4. Enter the source site details:

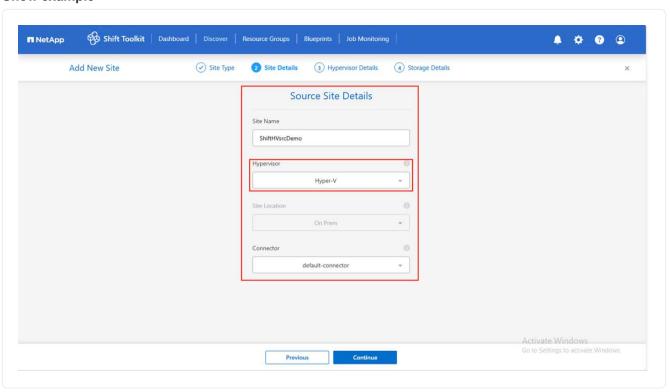
• Site Name: Provide a name for the site

• Hypervisor: Select Hyper-V

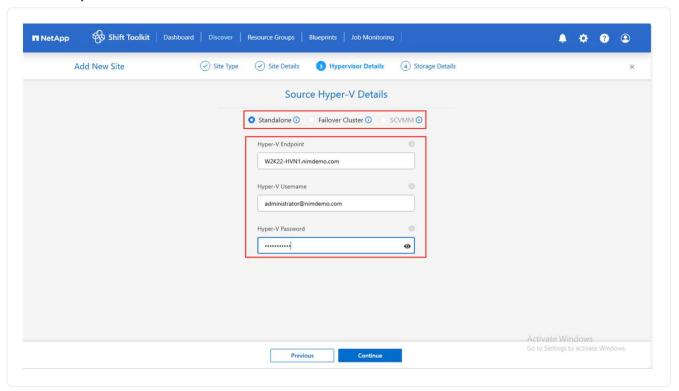
• Site Location: Select the default option

· Connector: Select the default selection

5. Click Continue.



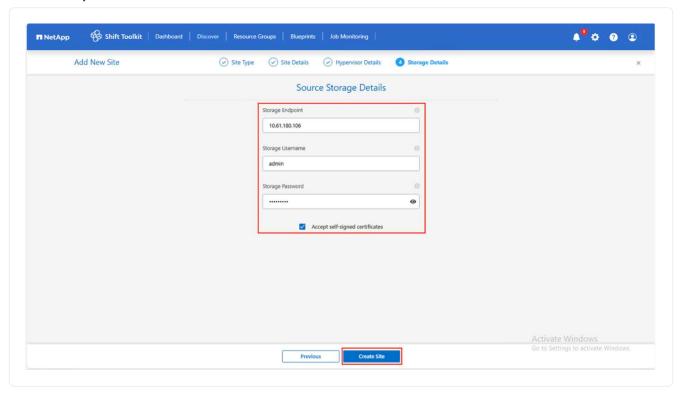
- 6. Enter the Hyper-V details:
 - Hyper-V standalone or failover cluster manager: IP address or FQDN
 - Username: Username in UPN format (username@domain.com or domain\administrator)
 - · Password: Password to access Hyper-V host or FCI instance
- 7. Click Continue.





The Hyper-V FCI and host discovery relies on DNS resolution. If resolution fails, update the host file (C:\Windows\System32\drivers\etc\hosts) and retry the discovery operation.

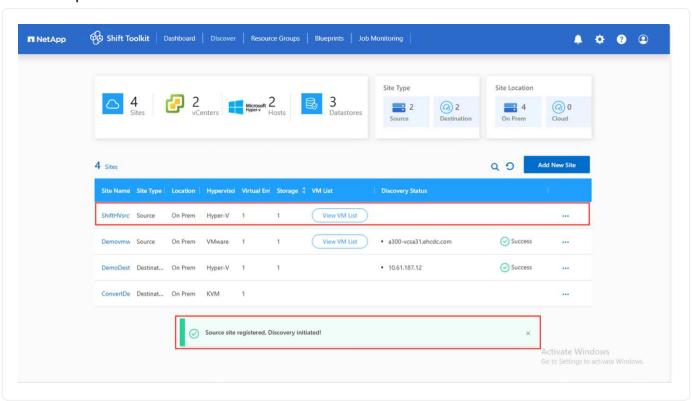
8. Enter the ONTAP storage system credentials.



9. Click Create Site.

Result

The Shift Toolkit performs automatic discovery and displays VMs with metadata information, including networks, virtual switches, and VLAN IDs.





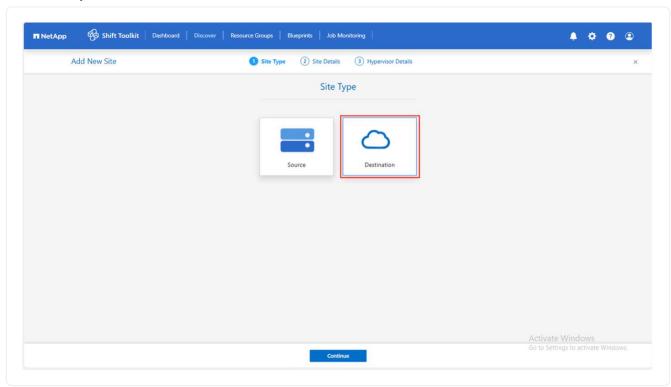
VM inventory auto-refreshes every 24 hours. To manually refresh after modifications, click the three dots next to the site name and select **Discover Site**.

Step 2: Add the destination site (VMware ESXi)

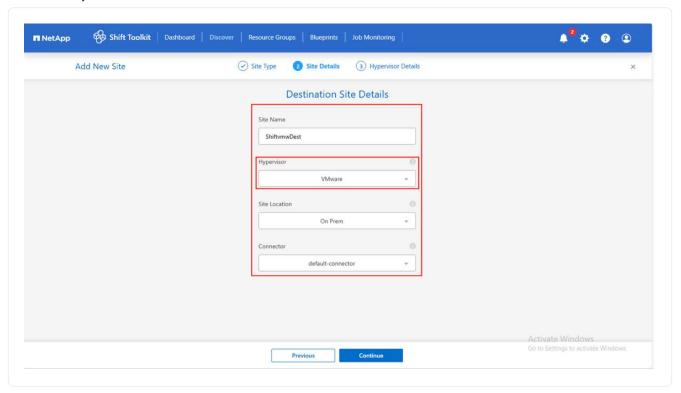
Add the destination VMware environment to the Shift Toolkit.

Steps

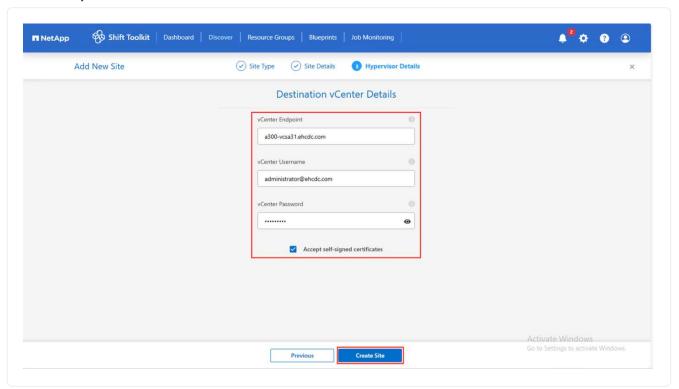
1. Click Add New Site and select Destination.



- 2. Enter the destination site details:
 - Site Name: Provide a name for the site
 - Hypervisor: Select VMware
 - · Site Location: Select the default option
 - Connector: Select the default selection
- 3. Click Continue.

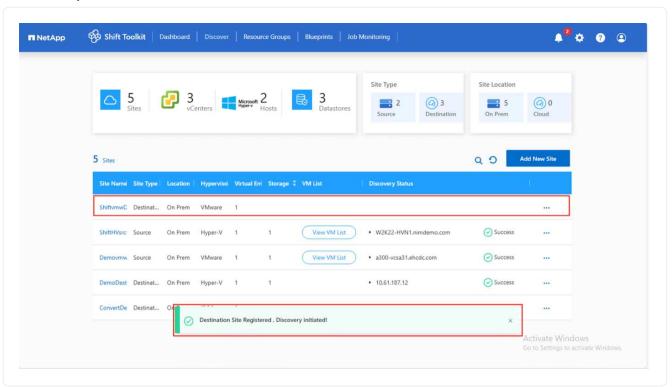


- 4. Enter the VMware vCenter details:
 - Endpoint: IP address or FQDN of the vCenter server
 - **Username**: Username in UPN format (username@domain.com)
 - vCenter Password: Password to access vCenter
 - vCenter SSL Thumbprint (optional)
- 5. Select Accept Self signed certificate and click Continue.



6. Click Create Site.

Show example



(i)

The source and destination storage system must be the same, as disk format conversion occurs at the volume level within the same volume.

Step 3: Create resource groups

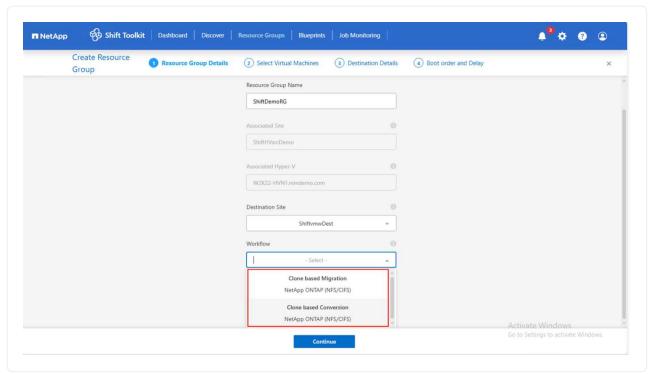
Organize VMs into resource groups to preserve boot order and boot delay configurations.

Before you begin

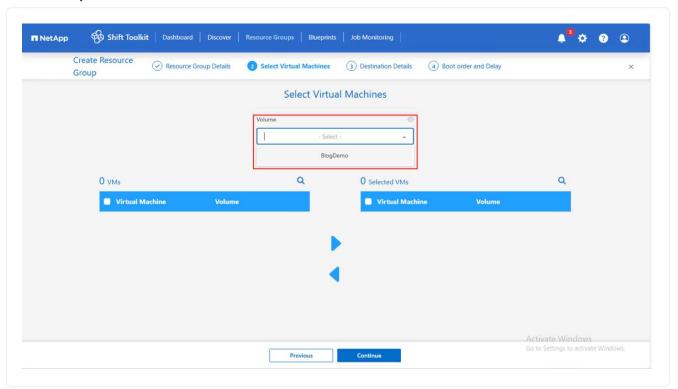
Ensure gtrees are provisioned as specified in the prerequisites.

Steps

- 1. Navigate to Resource Groups and click Create New Resource Group.
- 2. Select the source site from the dropdown and click Create.
- 3. Provide resource group details and select the workflow:
 - Clone based Migration: Performs end-to-end migration from source to destination hypervisor
 - · Clone based Conversion: Converts disk format to the selected hypervisor type



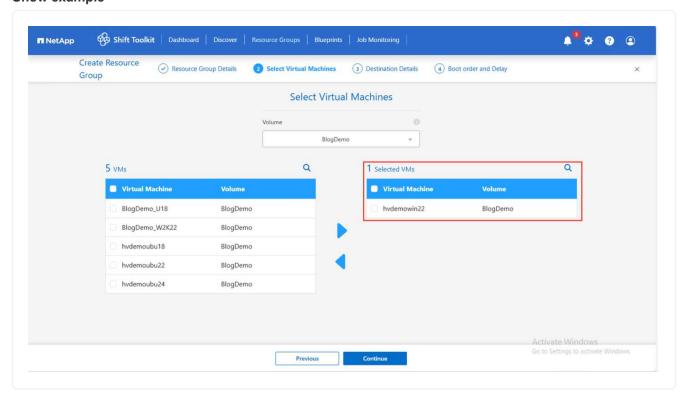
- 4. Click Continue.
- 5. Select VMs using the search option (default filter is "Datastore").





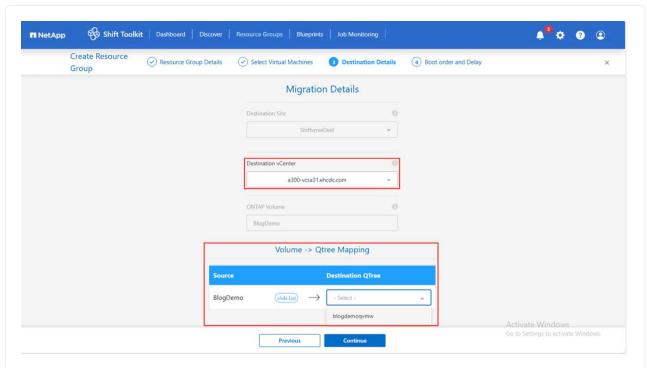
Move VMs to a designated SMB share on a newly created ONTAP SVM before conversion to isolate production shares from the staging area. The datastore dropdown only shows SMB shares; CSVs are not displayed.

Show example



6. Update migration details:

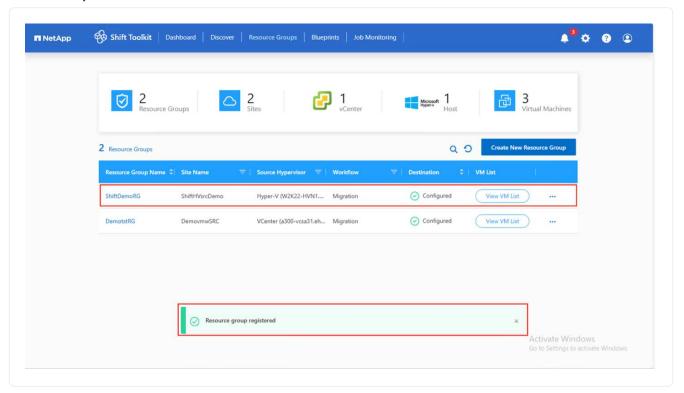
- Select Destination Site
- Select Destination VMware entry
- · Configure volume to qtree mapping





Set the destination path to the appropriate qtree when converting VMs from Hyper-V to ESXi.

- 7. Configure boot order and boot delay for all selected VMs:
 - 1: First VM to power on
 - 3: Default
 - **5**: Last VM to power on
- 8. Click Create Resource Group.



Result

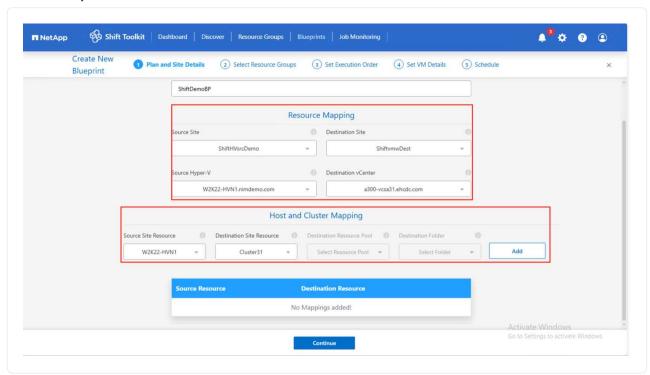
The resource group is created and ready for blueprint configuration.

Step 4: Create a migration blueprint

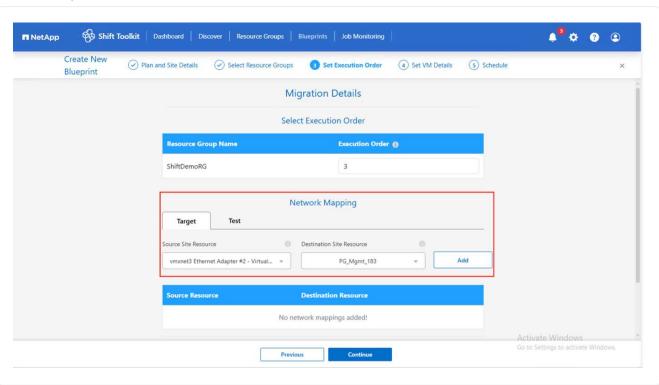
Create a blueprint to define the migration plan, including platform mappings, network configuration, and VM settings.

Steps

- 1. Navigate to **Blueprints** and click **Create New Blueprint**.
- 2. Provide a name for the blueprint and configure host mappings:
 - Select Source Site and associated Hyper-V hypervisor
 - Select **Destination Site** and associated vCenter
 - Configure host and cluster mapping



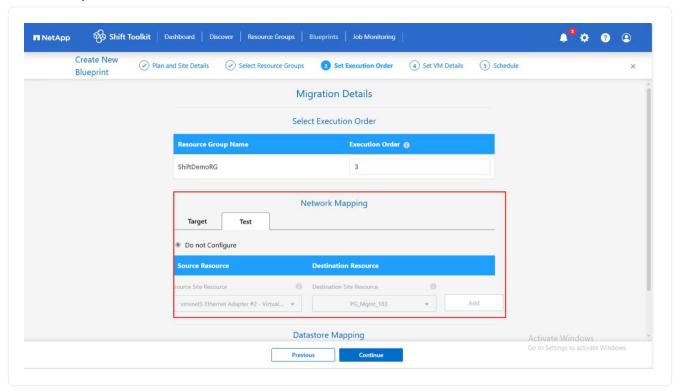
- 3. Select resource group details and click Continue.
- 4. Set execution order for resource groups if multiple groups exist.
- 5. Configure network mapping to appropriate port groups.



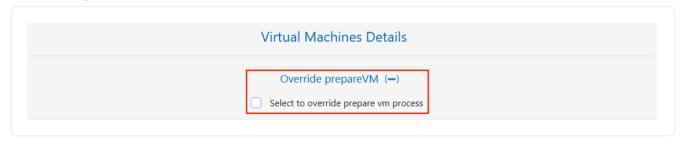


On VMware, Distributed Port Group is the only supported option. For test migration, select "Do not configure Network" to avoid production network conflicts; manually assign network settings after conversion.

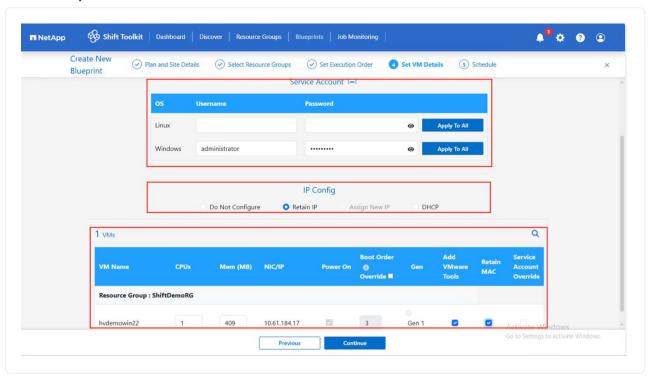
Show example



- 6. Review storage mappings (automatically selected based on VM selection).
 - Ensure qtrees are provisioned beforehand with necessary permissions.
- 7. Configure VM preparation override if needed for custom scripts or IP address customization.



- 8. Under VM details, provide service account and credentials for each OS type:
 - · Windows: Local administrator or domain credentials (ensure user profile exists on VM)
 - Linux: User with sudo privileges without password prompt



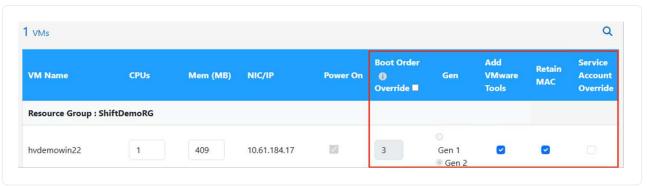
9. Configure IP settings:

- · Do not configure: Default option
- · Retain IP: Keep same IPs from source system
- DHCP: Assign DHCP on target VMs

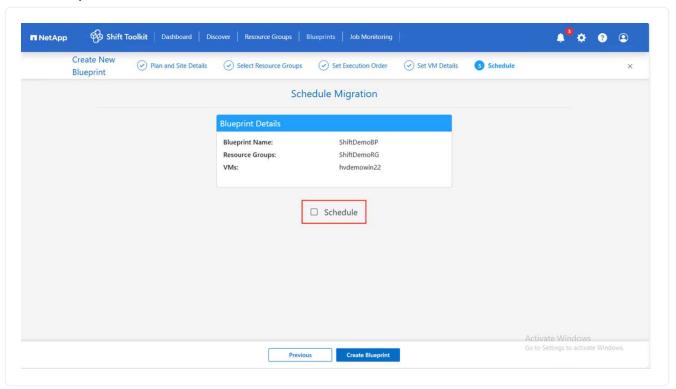
Ensure VMs are powered on during prepareVM phase and integration services are enabled.

10. Configure VM settings:

- Resize CPU/RAM parameters (optional)
- Modify boot order and boot delay
- Power ON: Select to power on VMs after migration (default: ON)
- Add VMware tools: Install VMware Tools after conversion (default: selected)
- Retain MAC: Keep MAC addresses for licensing requirements
- Service Account override: Specify separate service account if needed



- 11. Click Continue.
- 12. (Optional) Schedule the migration by selecting a date and time.

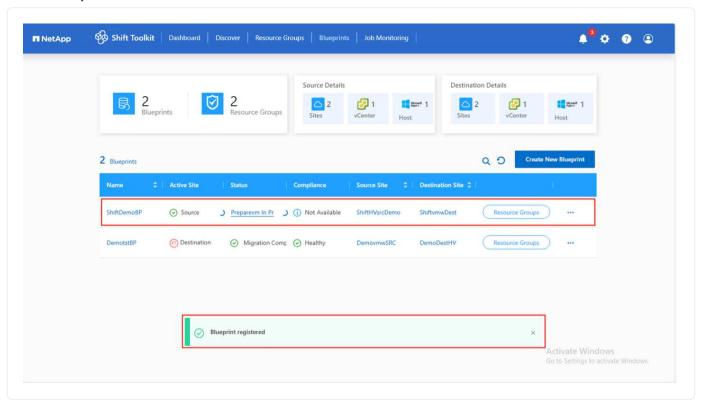


Schedule migrations at least 30 minutes ahead to allow time for VM preparation.

13. Click Create Blueprint.

Result

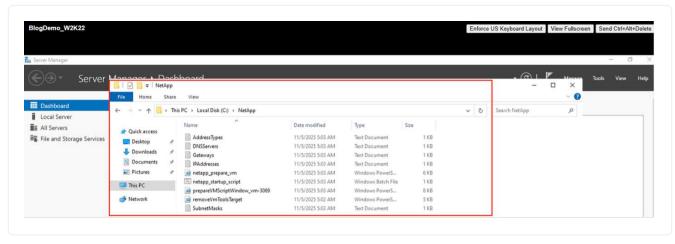
The Shift Toolkit initiates a prepareVM job that runs scripts on source VMs to prepare them for migration.



The preparation process:

• For Windows VMs: Stores scripts in $C: \$

Show example



• For Linux VMs: Stores scripts in /NetApp and /opt

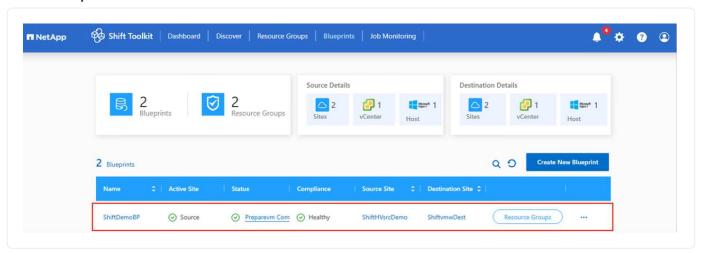




For CentOS or Red Hat VMs, the Shift Toolkit automatically installs necessary drivers before disk conversion to ensure successful boot after conversion.

When prepareVM completes successfully, the blueprint status updates to "Active."

Show example



Step 5: Execute the migration

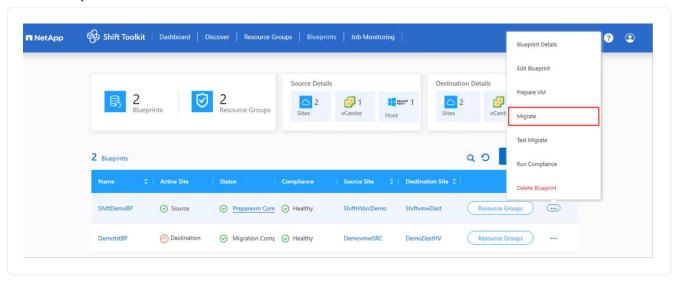
Trigger the migration workflow to convert VMs from Hyper-V to VMware ESXi.

Before you begin

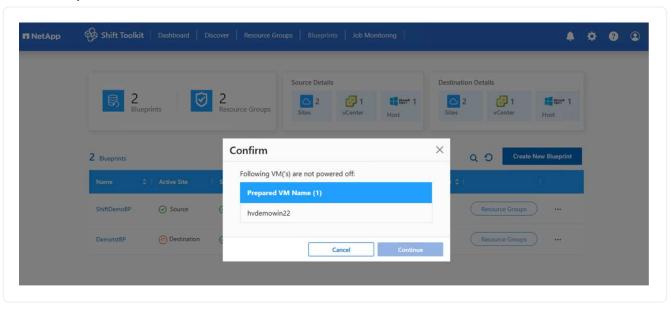
- · All VMs are gracefully powered off according to the planned maintenance schedule
- · The Shift Toolkit VM is part of the domain
- · CIFS share is configured with appropriate permissions
- · Qtrees have the correct security style
- · Integration Services are enabled on all guest VMs
- SSH is enabled on Linux-based guest VMs

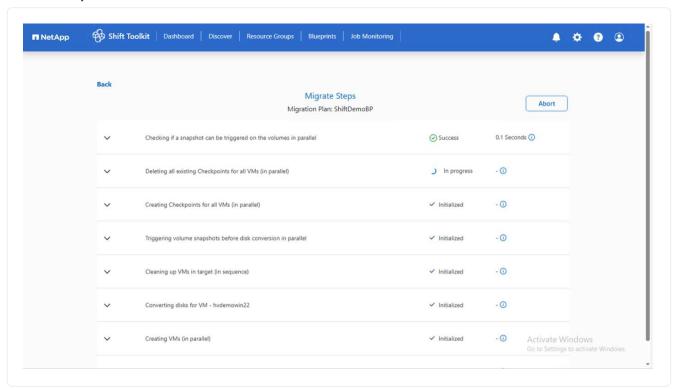
Steps

1. On the blueprint, click Migrate.



2. If VMs remain powered on, respond to the graceful shutdown prompt.



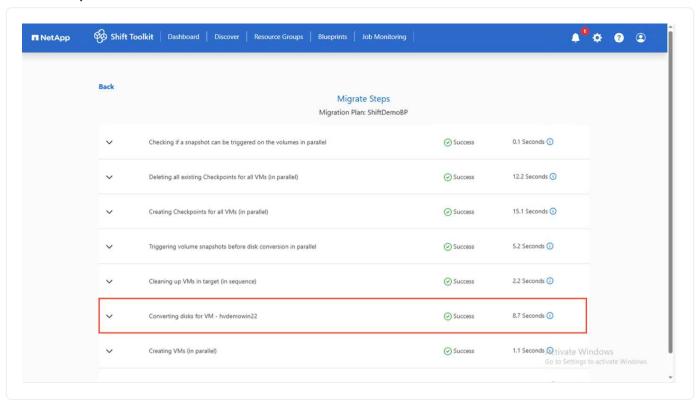


Result

The Shift Toolkit performs the following steps:

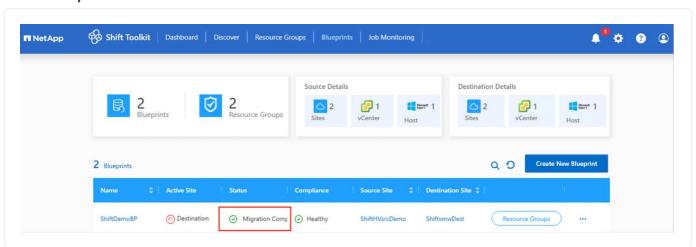
- 1. Powers off source VMs
- 2. Deletes existing checkpoints
- 3. Triggers VM checkpoints at the source
- 4. Triggers volume snapshot before disk conversion
- 5. Clones and converts VHDx files to VMDK format
- 6. Powers on VMs at the target site
- 7. Registers network settings
- 8. Adds VMware Tools and assigns IP addresses

The conversion completes in seconds, minimizing VM downtime.



When migration completes, the blueprint status changes to "Migration Complete."

Show example

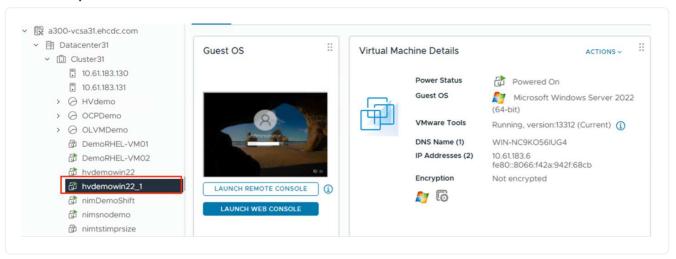


Step 6: Validate the migration

Verify that VMs are running successfully on the VMware ESXi host.

Steps

- 1. Log in to vCenter or the ESXi host.
- 2. Verify that VMs are running on the specified ESXi host.



- 3. Verify VM connectivity and application functionality.
- 4. (Windows VMs only) Bring offline disks online if needed:

Set-StorageSetting -NewDiskPolicy OnlineAll



After conversion, all VM disks on Windows OS except the OS disk will be offline due to the default Microsoft Windows SAN policy (offlineALL). This prevents data corruption when LUNs are accessed by multiple servers.

Result

The migration from Hyper-V to VMware ESXi is complete.



The Shift Toolkit uses cron jobs (Linux) and scheduled tasks (Windows) for post-migration operations. No SSH connections or equivalent are created after VMs are running on ESXi hosts.

Migrate VMs from VMware ESXi to Red Hat OpenShift Virtualization

Migrate VMs from VMware ESXi to Red Hat OpenShift Virtualization using the Shift Toolkit by preparing VMs, converting disk formats, and configuring the target environment.

The Shift Toolkit enables VM migration between virtualization platforms through disk format conversion and network reconfiguration on the destination environment.

Before you begin

Verify that the following prerequisites are met before starting the migration.

Red Hat OpenShift Virtualization requirements

- OpenShift Cluster endpoint with the following operators installed:
 - OpenShift Virtualization operator
 - NetApp Trident CSI driver

- NMstate
- NetApp Trident CSI configured with appropriate backends and storage classes
- NodeNetworkConfigurationPolicy and NetworkAttachmentDefinitions (NAD) configured with proper VLANs
- · OpenShift cluster is network reachable with current host file entries
- · Administrator level privileges on the cluster
- · Kubeconfig file downloaded

VMware requirements

• VMDKs are placed on individual volumes (mimicking VMDK to a PVC/PV construct) using symotion



This limitation will be removed in the next release where NAS-economy driver can be used for PVC provisioning.

- VMware tools are running on guest VMs
- VMs to be migrated are in a RUNNING state for preparation
- VMs must be powered off before triggering migration
- VMware tools removal happens on the destination hypervisor once VMs are powered on

Guest VM requirements

- For Windows VMs: Use local administrator credentials
- · For Linux VMs: Use a user with permissions to execute sudo commands without password prompt
- For Windows VMs: Mount the VirtlO ISO to the VM (download from here)



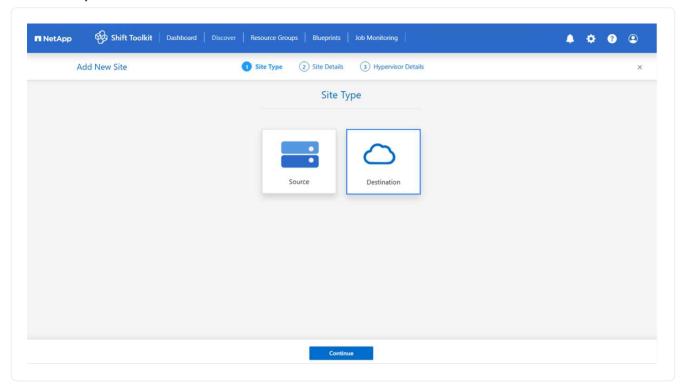
The preparation script uses the .msi package to install drivers and qemu-guest-agents.

Step 1: Add the destination site (OpenShift)

Add the destination OpenShift Virtualization environment to the Shift Toolkit.

Steps

1. Click Add New Site and select Destination.



2. Enter the destination site details:

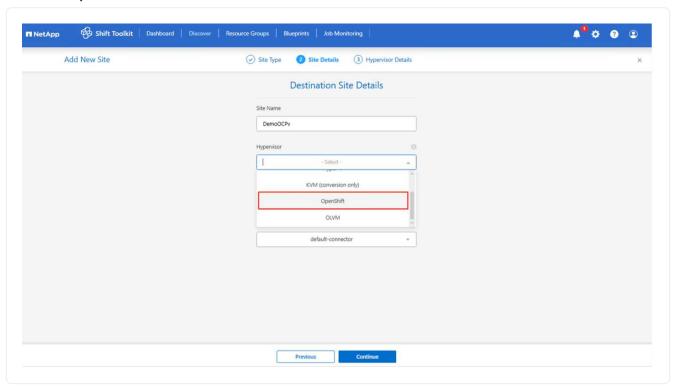
• Site Name: Provide a name for the site

• **Hypervisor**: Select OpenShift

• Site Location: Select the default option

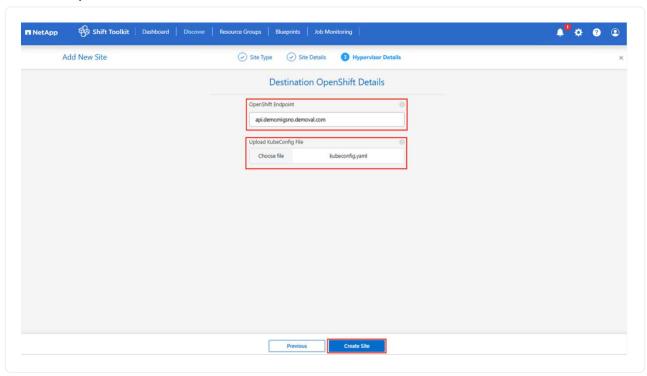
· Connector: Select the default selection

3. Click Continue.



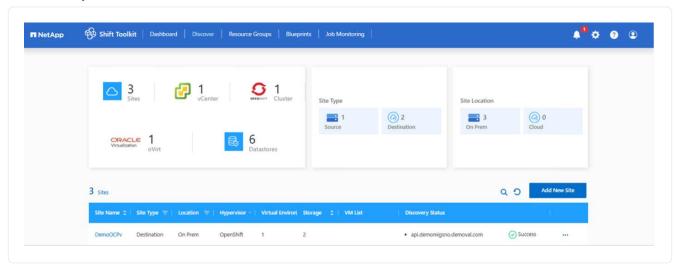
- 4. Enter the OpenShift details:
 - Endpoint: FQDN of OpenShift Cluster endpoint (for example, api.demomigsno.demoval.com)
 - Upload Kubeconfig file: Use the kubeconfig file with minimal permissions
 - (i)

The file extension must be yaml.



5. Click Create Site.

Show example





The source and destination volume will be the same as the disk format conversion happens at the volume level within the same volume.

Step 2: Create resource groups

Organize VMs into resource groups to preserve boot order and boot delay configurations.

Before you begin

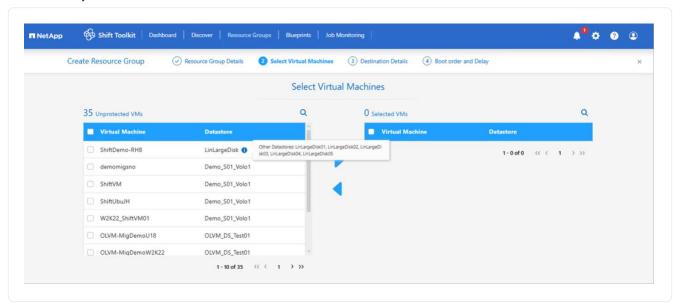
Ensure VM VMDKs are moved to individual datastore volumes on a newly created ONTAP SVM.

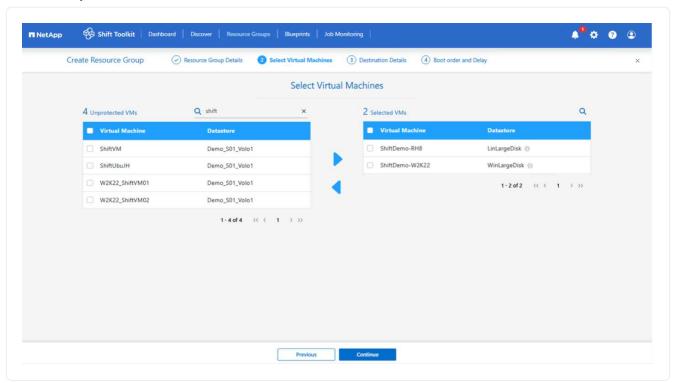
Steps

- 1. Navigate to Resource Groups and click Create New Resource Group.
- 2. Select the source site from the dropdown and click Create.
- 3. Provide resource group details and select the workflow:
 - Clone based Migration: Performs end-to-end migration from source to destination hypervisor
 - · Clone based Conversion: Converts disk format to the selected hypervisor type
- 4. Click Continue.
- 5. Select VMs using the search option.

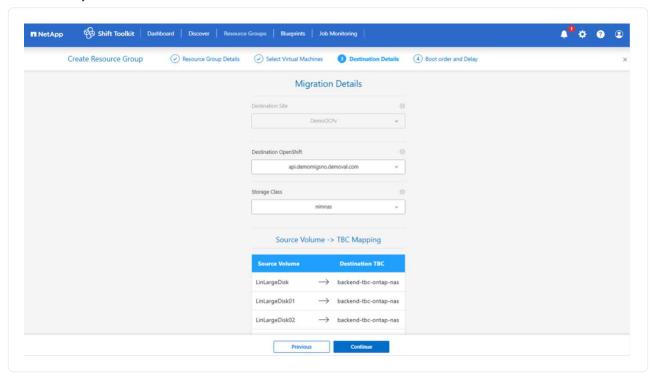


VM selection for resource groups is based on virtual machine and not at the datastore level.





- 6. Update migration details:
 - Select Destination Site
 - Select Destination OpenShift entry
 - Select the storage class

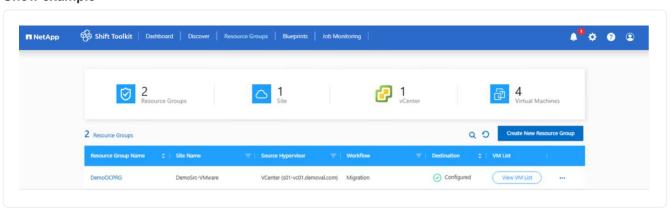




The Trident backend will be mapped to the source volume automatically if there is only one TBC; however, if there are multiple TBCs, the backend can be selected.

- 7. Configure boot order and boot delay for all selected VMs:
 - 1: First VM to power on
 - 3: Default
 - ∘ **5**: Last VM to power on
- 8. Click Create Resource Group.

Show example



Result

The resource group is created and ready for blueprint configuration.

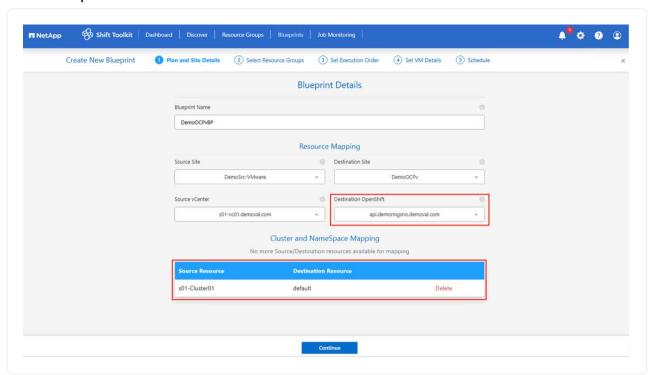
Step 3: Create a migration blueprint

Create a blueprint to define the migration plan, including platform mappings, network configuration, and VM settings.

Steps

- 1. Navigate to **Blueprints** and click **Create New Blueprint**.
- 2. Provide a name for the blueprint and configure host mappings:
 - Select Source Site and associated vCenter
 - Select **Destination Site** and associated OpenShift target
 - Configure cluster and host mapping

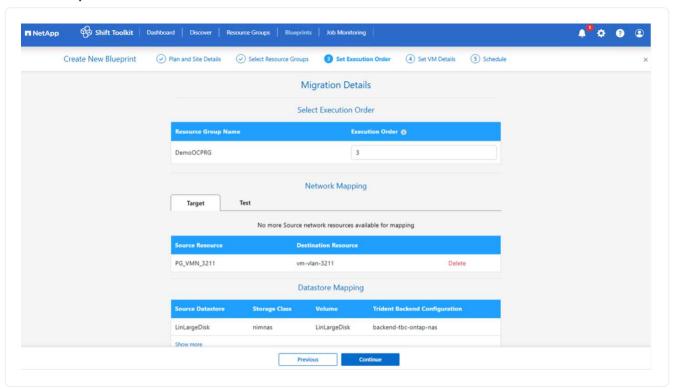
Show example



- 3. Select resource group details and click Continue.
- 4. Set execution order for resource groups if multiple groups exist.
- 5. Configure network mapping to appropriate logical networks.



Network attachment definitions should already be provisioned within the OpenShift cluster with the appropriate VLAN and trunk options. For test migration, select "Do not configure Network" to avoid production network conflicts; manually assign network settings after conversion.

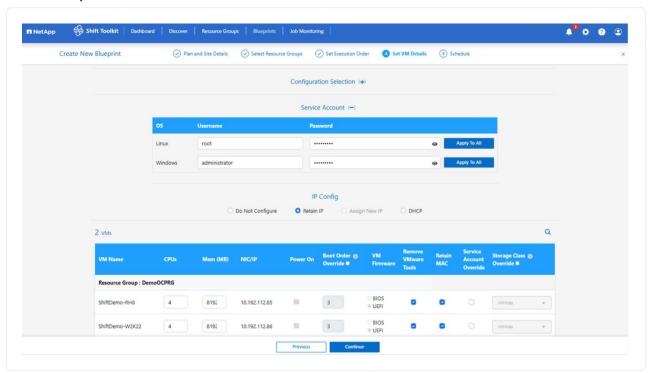


6. Review storage class and backend mappings (automatically selected based on VM selection).



Ensure VMDKs are symotioned to individual volumes beforehand so the virtual machine can be created and powered on from the PVC.

- 7. Under VM details, select configuration details and provide service account credentials for each OS type:
 - Windows: Use a user with local administrator privileges (domain credentials can also be used)
 - Linux: Use a user that can execute sudo commands without password prompt





The configuration selection allows you to select the disk image format, skip override prepareVM, and choose whether to split the volume from the parent. By default, split clone is disabled and the workflow defaults to RAW format.

8. Configure IP settings:

- · Do not configure: Default option
- Retain IP: Keep same IPs from source system
- DHCP: Assign DHCP on target VMs

Ensure VMs are powered on during prepareVM phase and VMware Tools are installed.

9. Configure VM settings:

- Resize CPU/RAM parameters (optional)
- Modify boot order and boot delay
- Power ON: Select to power on VMs after migration (default: ON)
- Remove VMware tools: Remove VMware Tools after conversion (default: selected)
- VM Firmware: BIOS > BIOS and EFI > EFI (automatic)
- Retain MAC: Keep MAC addresses for licensing requirements



If interface name needs to be retained while retaining the MAC address, ensure appropriate udev rules are created on the source VM.

- · Service Account override: Specify separate service account if needed
- 10. Click Continue.
- 11. (Optional) Schedule the migration by selecting a date and time.



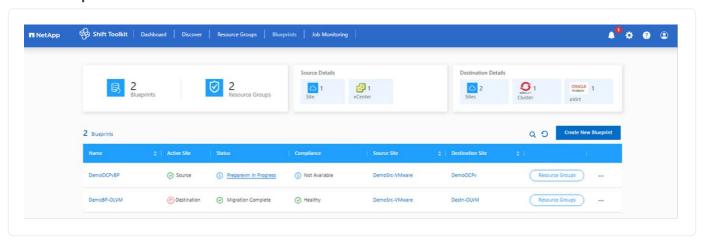
Schedule migrations at least 30 minutes ahead to allow time for VM preparation.

12. Click Create Blueprint.

Result

The Shift Toolkit initiates a prepareVM job that runs scripts on source VMs to prepare them for migration.

Show example



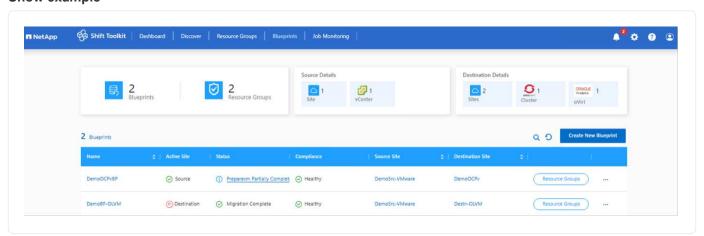
The preparation process:

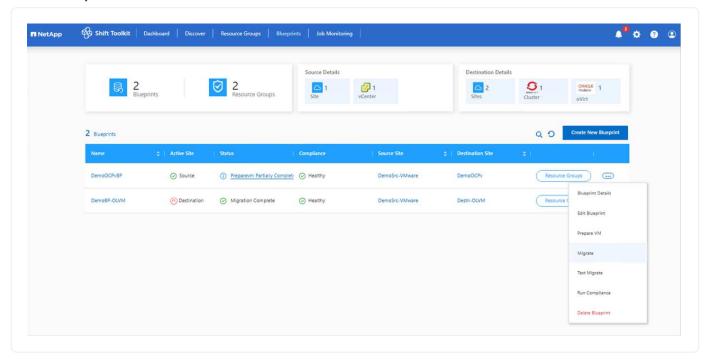
- Injects scripts to update VirtlO drivers, install qemu-agent, remove VMware tools, backup IP details, and update fstab
- · Uses PowerCLI to connect to guest VMs (Linux or Windows) and update VirtIO drivers
- For Windows VMs: Stores scripts in C:\NetApp
- For Linux VMs: Stores scripts in /NetApp and /opt



For any supported VM OSes, the Shift Toolkit automatically installs necessary VirtlO drivers before disk conversion to ensure successful boot after conversion.

When prepareVM completes successfully, the blueprint status updates to "PrepareVM Complete." Migration will now happen at the scheduled time or can be started manually by clicking the **Migrate** option.





Step 4: Execute the migration

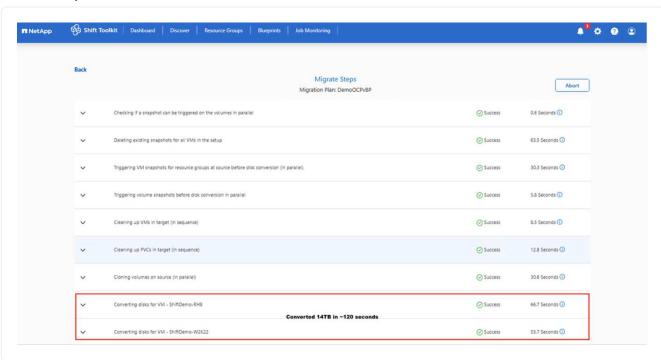
Trigger the migration workflow to convert VMs from VMware ESXi to OpenShift Virtualization.

Before you begin

All VMs are gracefully powered off according to the planned maintenance schedule.

Steps

1. On the blueprint, click **Migrate**.



- The Shift Toolkit performs the following steps:
 - Deletes existing snapshots for all VMs in the blueprint
 - Triggers VM snapshots at the source
 - Triggers volume snapshot before disk conversion
 - Clones the individual volumes
 - Converts VMDK to RAW format for each VMDK

The Shift Toolkit automatically finds all VMDKs associated with each VM, including the primary boot disk.



If there are multiple VMDK files, each VMDK will be converted. In this release (v4.0), each VMDK should be placed on an individual volume/datastore.

· Cleans up the volumes to have just the disk.img file

With the virtual machine disk image converted to RAW format, the Shift Toolkit cleans up the volumes, renames the raw file to disk.img, and assigns necessary permissions.

Imports the volumes as PVCs using Trident import

The volumes are then imported as PVCs using NetApp Trident APIs.

· Creates VMs using VM-specific yaml files

Once the PVCs are imported and PVs are in place, the Shift Toolkit uses OC CLI to create each VM depending on the OS using yaml files.



VMs are created under the "Default" namespace.

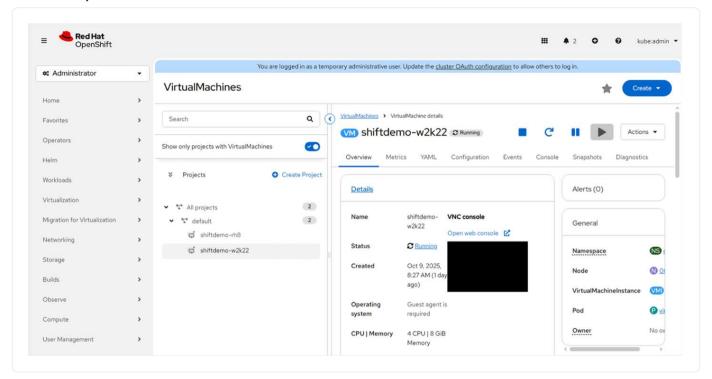
Powers on VMs at the target

Depending on the VM OS, the Shift Toolkit auto-assigns the VM boot option along with storage controller interfaces. For Linux distributions, VirtlO or VirtlO SCSI is used. For Windows, the VM powers on with SATA interface, then the scheduled script auto-installs VirtlO drivers and changes the interface to VirtlO.

· Registers networks on each VM

Networks are assigned based on the blueprint selection.

Removes VMware tools and assigns IP addresses using cron jobs



Use Migration Toolkit for Virtualization with Shift Toolkit

This section describes how to use Migration Toolkit for Virtualization (MTV) with NetApp Shift Toolkit for seamless migration to Red Hat OpenShift Virtualization.

Before you begin

Ensure the following prerequisites are met:

- OpenShift cluster with OpenShift Virtualization operator and NetApp Trident CSI driver installed
- MTV 2.9.4 (which includes conversion mode)
- · Shift Toolkit installed



Since only Shift Toolkit API is used, there is no need to configure Shift Toolkit resource groups or blueprints.

- · Administrator level privileges on the OpenShift cluster
- · A Linux instance with tridentctl and OC command line tool installed
 - Kubeconfig exported or OC login executed to connect to the cluster
 - Download the script named "OpenShift-MTV" from Shift Toolkit UI (Settings > Developer Access > Script Blocker)
 - Unzip the file: unzip openshift-mtv.zip
 - Ensure Python3 is installed: dnf install python3
 - o Install OpenJDK 8 or later: yum install java-1.8.0-openjdk
 - o Install requirements: pip install -r requirements.txt
- Virtual Machine requirements for MTV: VMDKs for a VM must be placed on individual volumes. For a

VM with 3 disks, each disk should be on its individual volume (mapping datastore to PVC construct). This must be done manually using storage vmotion.

Steps

1. Create migration plans using MTV.

To leverage fast VMDK conversion, create a migration plan for the VMs and ensure the following parameters are in the YAML:

```
o targetNamespace: default
o type: conversion
o storage: {}
```



The plan should be created beforehand to ensure preserve IP settings are configured by MTV.

2. Map VMs from vCenter and volumes on ONTAP storage.

Use the script to create necessary PVCs and import them to the OpenShift cluster. The PVCs must have the following labels and annotations:

Labels:

vmID and vmUUID in the PVC (Forklift looks for these values)

Annotation:

The vmdk disk name for forklift.konveyor.io/disk-source

The script ensures these attributes are set for every PVC and updates disk.img permissions:

```
o "owner": { "id": 107 }
o "group": { "id": 107 }
o "mode": "0655"
```

- 3. Update the JSON file with the following details:
 - ONTAP Cluster: Can be a SVM; vsadmin can be used. Set splitclone to "False" if clone volume doesn't need immediate detachment
 - vCenter: Minimum RBAC rights to discover VMs and associated VMDK files
 - Trident storage class: Should be NFS backend with correct version in yaml
 - OpenShift: Specify the project name (default is used as an example)



Keep the rest of the values as default.

- 4. Once prerequisites are met, execute python3 main.py to create PVCs and import them to the OpenShift Cluster.
- 5. Once PVCs are imported, trigger migration using MTV to create the VM with the appropriate specification.

```
rocckin-Num-V201/Nome/wasedmin/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv/openshift-mtv
```

Show example



6. Convert VMDK with MTV.

The script automatically finds all VMDKs associated with each VM, including the primary boot disk.



If there are multiple VMDK files, each VMDK will be converted.

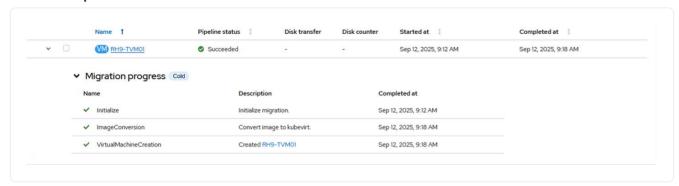
7. Upload RAW image to OpenShift Virtualization.

The script uses Trident CSI to import volumes as PVCs to the cluster. The PVC yaml is populated with labels and annotations.

Create virtual machine with MTV.

After import, call the MTV plan to start migration. The UI shows as "Cold," but based on the yaml specification of conversion, MTV checks for each PVC and the vmID/vmUUID, maps them, and initializes the migration.

Show example



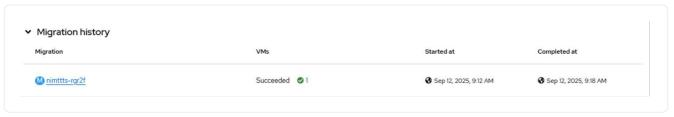


VMs are created under "Default" project for virtual machines, but this can be modified within the MTV migration plan YAML.

9. Boot VM for the first time with MTV.

Depending on the VM OS, MTV auto-assigns the VM boot option along with storage controller interfaces.

Show example



Migration completed in 6 minutes for a VM with 1.5TB data disk (spread across 3 PVCs). This showcases a streamlined, low-impact approach to re-homing VMs using ONTAP storage.



Before getting started with this specific integration, contact your Red Hat account team.

Video demonstration

The following video demonstrates the process outlined in this solution.

Zero touch migration from ESX to Oracle Linux Virtualization Manager (OLVM)

Migrate VMs from VMware ESXi to Oracle Linux Virtualization Manager

Migrate VMs from VMware ESXi to Oracle Linux Virtualization Manager (OLVM) using the Shift Toolkit by preparing VMs, converting disk formats, and configuring the target environment.

The Shift Toolkit enables VM migration between virtualization platforms through disk format conversion and network reconfiguration on the destination environment.

Before you begin

Verify that the following prerequisites are met before starting the migration.

Oracle Linux Virtualization Manager requirements

- · Oracle Linux Virtualization Manager with Oracle Linux KVM hosts added to the datacenter
- ONTAP NFS storage added as storage domain
- · Administrator level privileges on the cluster
- Oracle Linux Virtualization Manager and VDSM releases are >= 4.5
- Oracle Linux Virtualization Manager (destination) hosts are network reachable
- NFSv3 Storage domain configured with the appropriate volume and qtree
 - Ensure read-write access to the vdsm user (UID 36) and kvm group (GID 36) is allowed
- · Networks configured with the appropriate VLANs

VMware requirements

- VM VMDKs are placed on NFSv3 volume (all VMDKs for a given VM should be part of the same volume)
- · VMware tools are running on guest VMs
- VMs to be migrated are in a RUNNING state for preparation
- VMs must be powered off before triggering migration
- VMware tools removal happens on the destination hypervisor once VMs are powered on

Guest VM requirements

- For Windows VMs: Use local administrator credentials
- For Linux VMs: Use a user with permissions to execute sudo commands without password prompt
- For Windows VMs: Mount the VirtlO ISO to the VM (download from here)



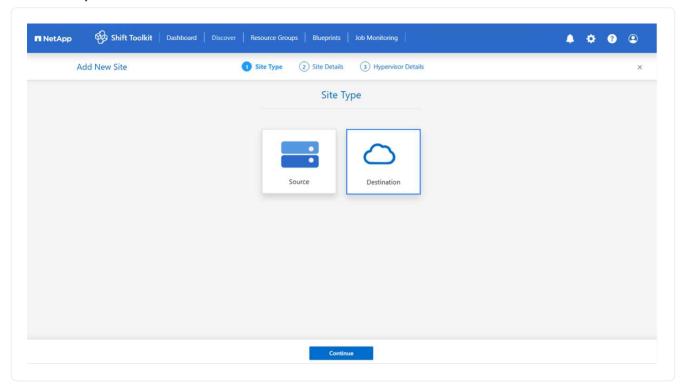
The preparation script uses the .msi package to install drivers and qemu-guest-agents.

Step 1: Add the destination site (OLVM)

Add the destination Oracle Linux Virtualization Manager environment to the Shift Toolkit.

Steps

1. Click Add New Site and select Destination.



2. Enter the destination site details:

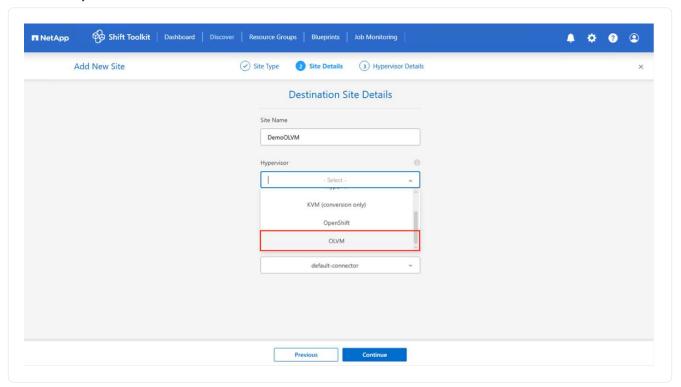
• Site Name: Provide a name for the site

• **Hypervisor**: Select OLVM

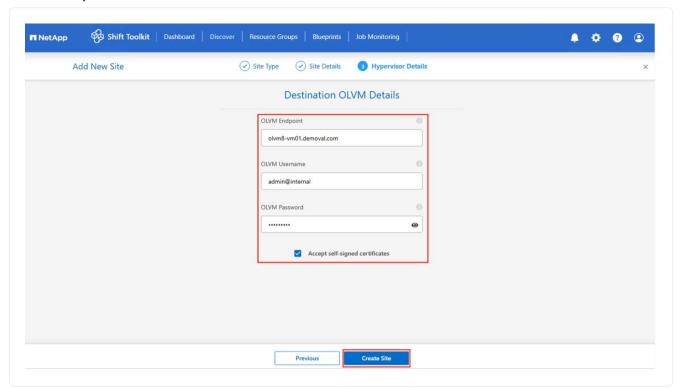
• Site Location: Select the default option

· Connector: Select the default selection

3. Click Continue.

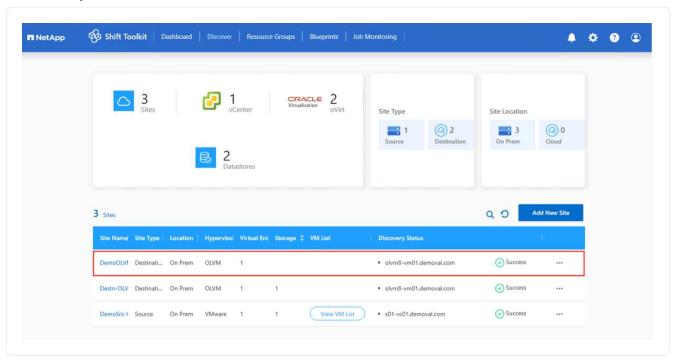


- 4. Enter the OLVM details:
 - Endpoint: IP address or FQDN of Virtualization Manager
 - **Username**: Username in format username@profile (for example, admin@internal)
 - · Password: Password to access Virtualization Manager
- 5. Select Accept Self signed certificate and click Continue.



6. Click Create Site.

Show example





The source and destination volume will be the same as the disk format conversion happens at the volume level within the same volume.

Step 2: Create resource groups

Organize VMs into resource groups to preserve boot order and boot delay configurations.

Before you begin

- Ensure gtrees are provisioned as specified in the prerequisites
- Move VMs to a designated datastore on a newly created ONTAP SVM before conversion to isolate production NFS datastores from the staging area

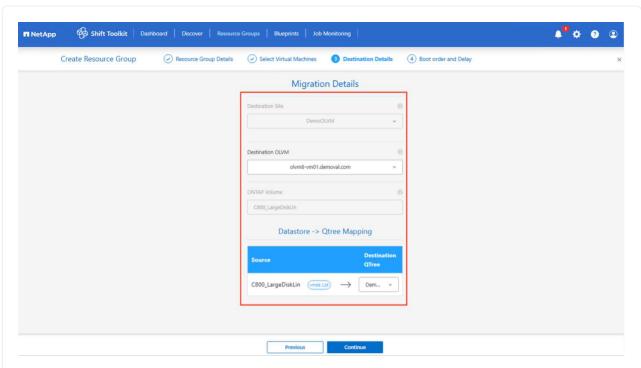
Steps

- 1. Navigate to Resource Groups and click Create New Resource Group.
- 2. Select the source site from the dropdown and click Create.
- 3. Provide resource group details and select the workflow:
 - Clone based Migration: Performs end-to-end migration from source to destination hypervisor
 - · Clone based Conversion: Converts disk format to the selected hypervisor type
- Click Continue.
- 5. Select VMs using the search option (default filter is "Datastore").



The datastore dropdown only shows NFSv3 datastores. NFSv4 datastores are not displayed.

- 6. Update migration details:
 - Select Destination Site
 - Select Destination OLVM entry
 - Configure Datastore to Qtree mapping

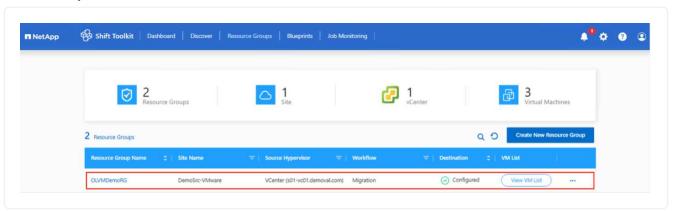




Ensure the destination path (where the converted VMs are stored) is set to a qtree when converting VMs from ESXi to OLVM. Also ensure this qtree is added to the storage domain. Multiple qtrees can be created and used for storing converted VM disks.

- 7. Configure boot order and boot delay for all selected VMs:
 - 1: First VM to power on
 - 3: Default
 - ∘ **5**: Last VM to power on
- 8. Click Create Resource Group.

Show example



Result

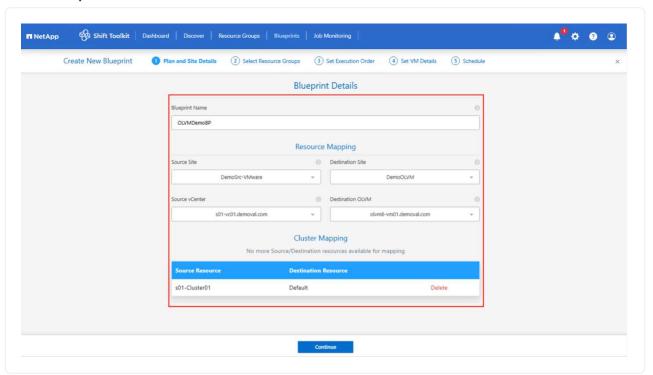
The resource group is created and ready for blueprint configuration.

Step 3: Create a migration blueprint

Create a blueprint to define the migration plan, including platform mappings, network configuration, and VM settings.

Steps

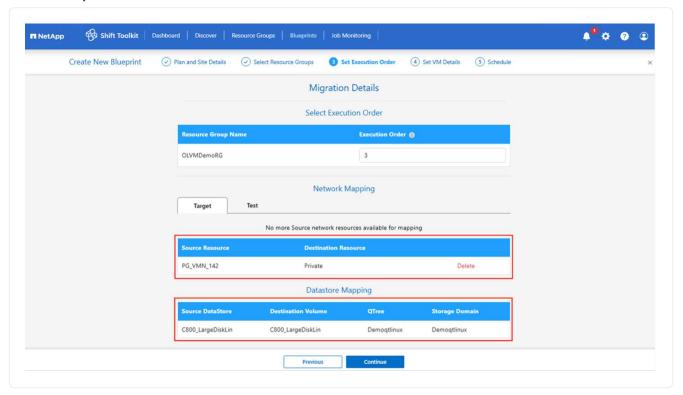
- 1. Navigate to **Blueprints** and click **Create New Blueprint**.
- 2. Provide a name for the blueprint and configure host mappings:
 - Select Source Site and associated vCenter
 - Select **Destination Site** and associated OLVM target
 - Configure cluster and host mapping



- 3. Select resource group details and click Continue.
- 4. Set execution order for resource groups if multiple groups exist.
- 5. Configure network mapping to appropriate logical networks.



Networks should already be provisioned within OLVM with the appropriate VLAN tagging. For test migration, select "Do not configure Network" to avoid production network conflicts; manually assign network settings after conversion.

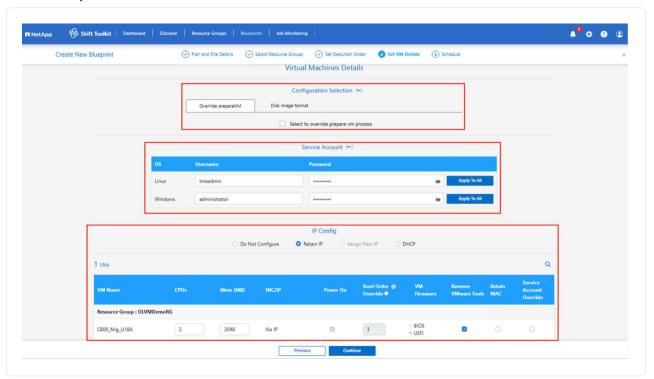


6. Review storage mappings (automatically selected based on VM selection).



Ensure the qtree is provisioned beforehand and necessary permissions are assigned so the virtual machine can be created and powered on from NFS volume.

- 7. Under VM details, select configuration details and provide service account credentials for each OS type:
 - Windows: Use a user with local administrator privileges (domain credentials can also be used)
 - · Linux: Use a user that can execute sudo commands without password prompt





The configuration selection allows you to select the disk image format and skip override prepareVM. The workflow defaults to QCOW2 format, but RAW format can be selected if required. The override prepareVM option allows administrators to skip VM preparation and run custom scripts.

8. Configure IP settings:

- Do not configure: Default option
- Retain IP: Keep same IPs from source system
- DHCP: Assign DHCP on target VMs

Ensure VMs are powered on during prepareVM phase and VMware Tools are installed.

- 9. Configure VM settings:
 - Resize CPU/RAM parameters (optional)
 - Modify boot order and boot delay
 - Power ON: Select to power on VMs after migration (default: ON)
 - Remove VMware tools: Remove VMware Tools after conversion (default: selected)
 - VM Firmware: BIOS > BIOS and EFI > EFI (automatic)
 - Retain MAC: Keep MAC addresses for licensing requirements
 - · Service Account override: Specify separate service account if needed
- 10. Click Continue.
- 11. Schedule the migration by selecting a date and time.



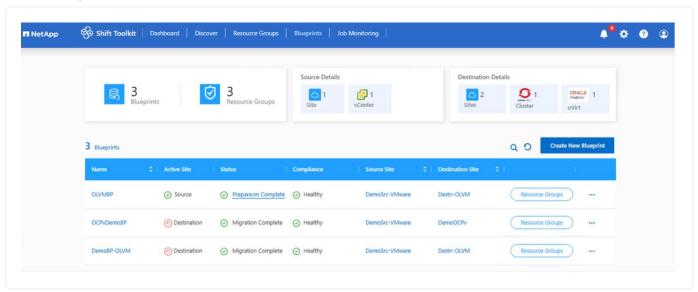
Schedule migrations at least 30 minutes ahead to allow time for VM preparation.

12. Click Create Blueprint.

Result

The Shift Toolkit initiates a prepareVM job that runs scripts on source VMs to prepare them for migration.

Show example



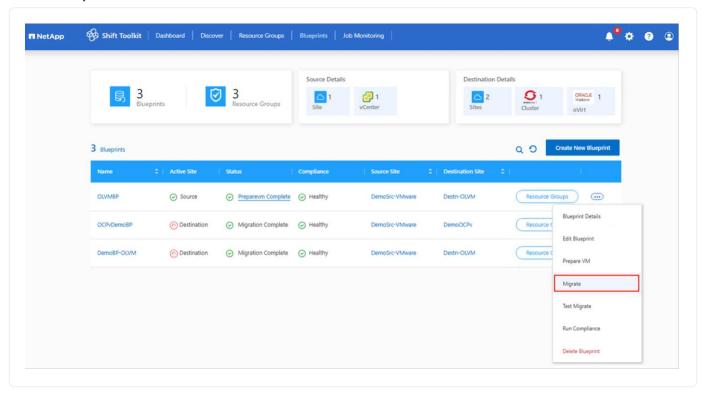
The preparation process:

- Injects scripts to update VirtlO drivers, install qemu-agent, remove VMware tools, backup IP details, and update fstab
- Uses PowerCLI to connect to guest VMs (Linux or Windows) and update VirtIO drivers
- For Windows VMs: Stores scripts in C:\NetApp
- For Linux VMs: Stores scripts in /NetApp and /opt



For any supported VM OSes, the Shift Toolkit automatically installs necessary VirtlO drivers before disk conversion to ensure successful boot after conversion.

When prepareVM completes successfully, the blueprint status updates to "PrepareVM Complete." Migration will now happen at the scheduled time or can be started manually by clicking the **Migrate** option.



Step 4: Execute the migration

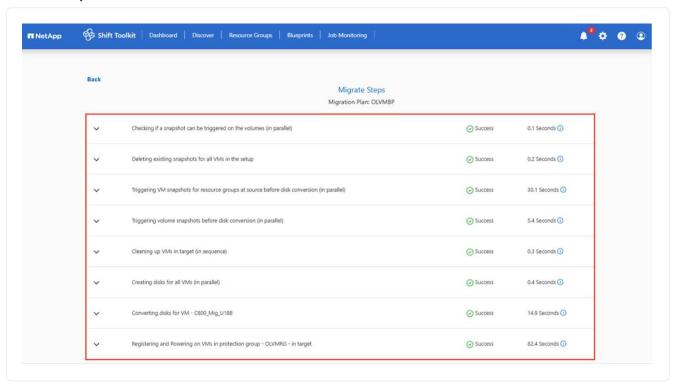
Trigger the migration workflow to convert VMs from VMware ESXi to Oracle Linux Virtualization Manager.

Before you begin

All VMs are gracefully powered off according to the planned maintenance schedule.

Steps

1. On the blueprint, click **Migrate**.



- 2. The Shift Toolkit takes the following actions:
 - Deletes existing snapshots for all VMs in the blueprint
 - Triggers VM snapshots at the source
 - Triggers volume snapshot before disk conversion
 - Converts VMDK to QCOW2 or RAW format for all VMs

The Shift Toolkit automatically finds all VMDKs associated with each VM, including the primary boot disk.



If there are multiple VMDK files, each VMDK will be converted.

Uploads the QCOW2 or RAW image to OLVM storage domain

With the virtual machine disk image converted to QCOW2 or RAW format, the Shift Toolkit uploads the file to the appropriate storage domain and adds each disk.

Creates virtual machines

The Shift Toolkit makes REST API calls to create each VM depending on the OS.



VMs are created under the "Default" cluster.

Powers on VMs at the target

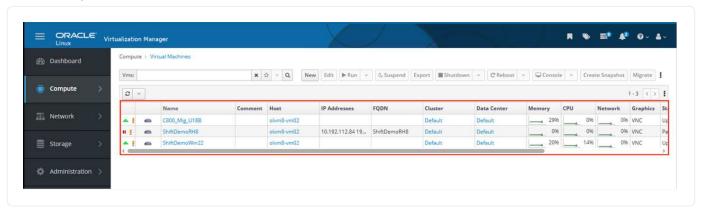
Depending on the VM OS, the Shift Toolkit auto-assigns the VM boot option along with storage controller interfaces. For Linux distributions, VirtlO or VirtlO SCSI is used. For Windows, the VM powers on with SATA interface, then the scheduled script auto-installs VirtlO drivers and changes the interface to VirtlO.

Registers networks on each VM

Networks are assigned based on the blueprint selection.

Removes VMware tools and assigns IP addresses using trigger scripts or cron jobs

Show example



Video demonstration

The following video demonstrates the process outlined in this solution.

Zero touch migration from ESX to Oracle Linux Virtualization Manager (OLVM)

Convert VMs using the Shift Toolkit

Use the Shift Toolkit to convert VMware ESX virtual machine disks (VMDK) to Microsoft Hyper-V (VHDX) disk format or Red Hat KVM (QCOW2) disk format. This process includes setting up resource groups, creating conversion blueprints, and scheduling conversions.

Overview

Shift toolkit supports disk-level conversions of virtual disks between hypervisors for the following disk formats:

- VMware ESX to Microsoft Hyper-V (VMDK to VHDX)
- Microsoft Hyper-V to VMware ESX (VHDX to VMDK)
- VMware ESX to Red Hat KVM (VMDK to QCOW2)
- VMware ESX to Red Hat KVM (VMDK to RAW)

The converted qcow2 files are compatible with any KVM hypervisors. For example, a qcow2 file can be utilized with RHEL-based KVM using virt-manager to create a VM, as well as with ubuntu KVM Rocky Linux based KVM and others. The same can be used with Oracle Linux virtualization manager with a tweak and with OpenShift virtualization after importing using NetApp Trident. The goal is to provide the disk (converted in secs to mins) which can then be integrated into existing automation scripts used by organizations to provision the VM and assign the network. This approach helps reduce overall migration times, with disk conversion handled by Shift toolkit APIs and the remaining script bringing up the VMs.

Shift toolkit supports end-to-end migration from VMware to other compatible KVM hypervisors. However, convert option provides ability to migration administrators to run these conversion and migration APIs.

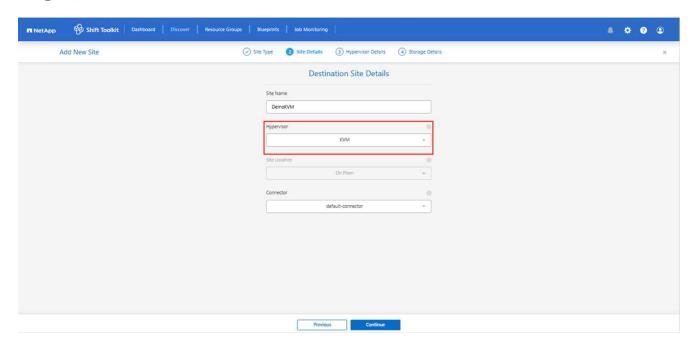
Convert to QCOW2 format

To convert the virtual disks to QCOW2 format with NetApp Shift toolkit, follow these high-level steps:

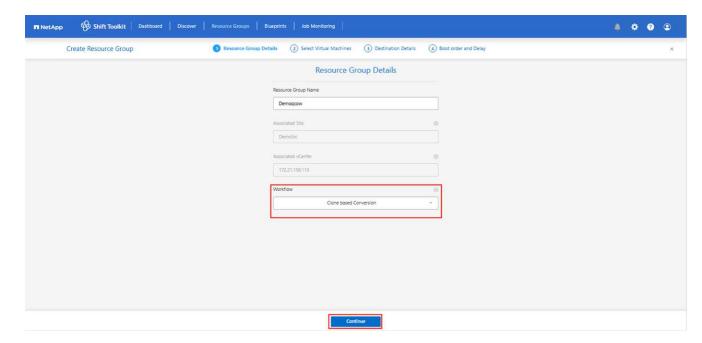
• Create a destination site type specifying KVM (conversion only) as the hypervisor.

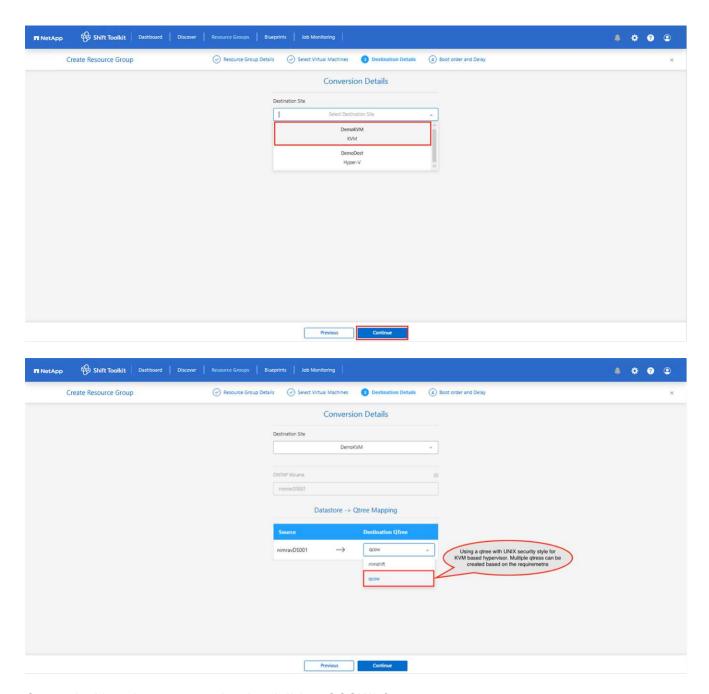


Hypervisor details are not required for KVM.

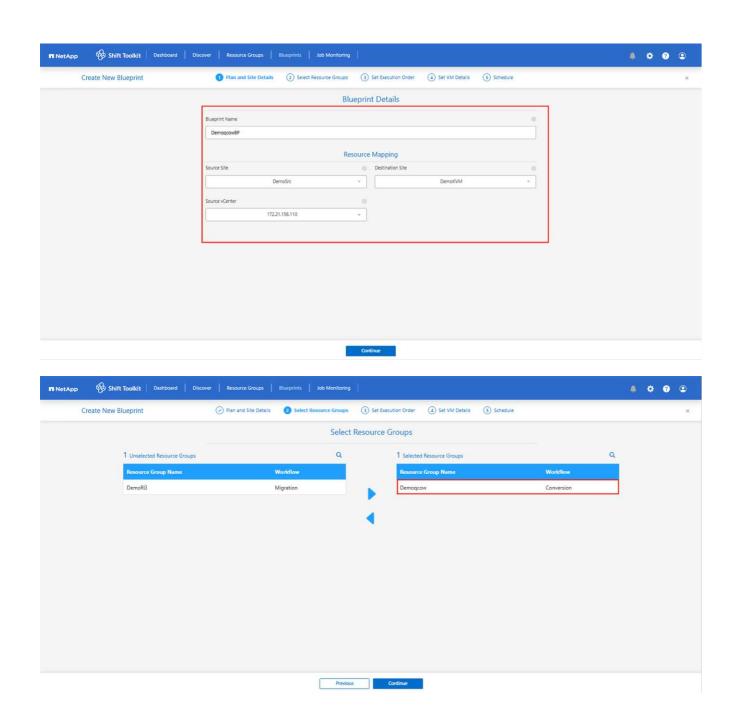


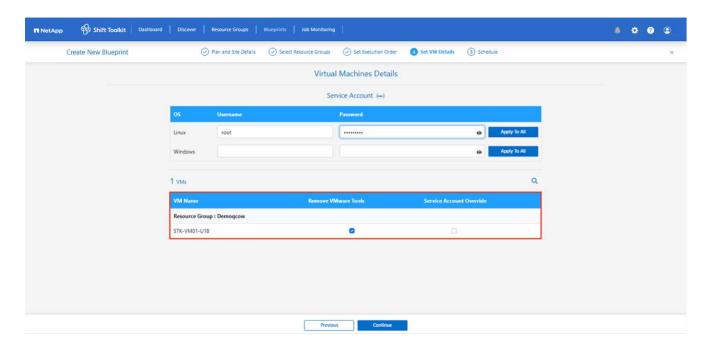
• Create a resource group with the VMs for which the disk conversion is required



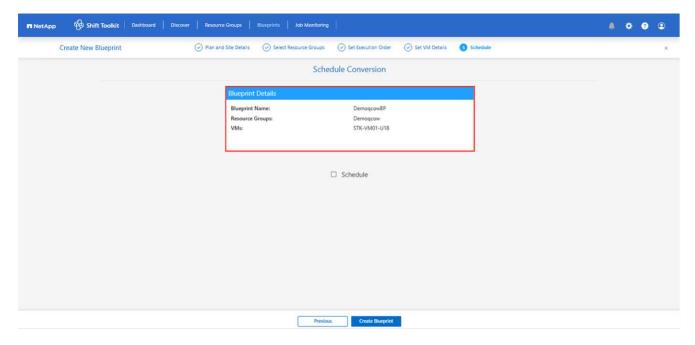


• Create the blueprint to convert the virtual disk to QCOW2 format.

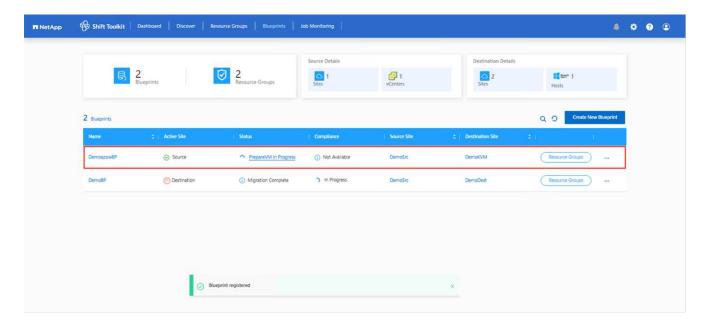




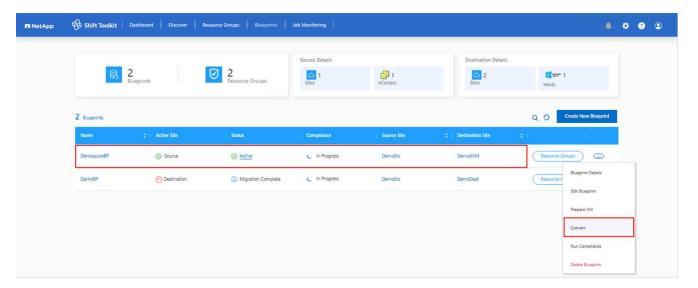
• Designate a slot using the scheduling option. If the conversion is to be performed on an ad-hoc basis, leave the scheduling option unchecked.



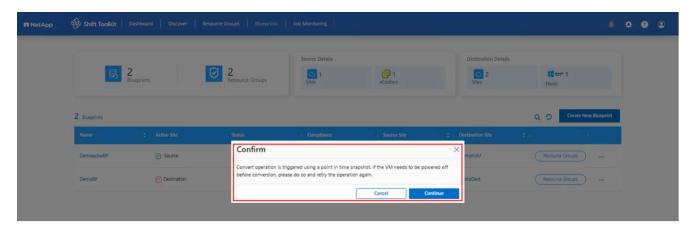
 After the blueprint is created, a prepareVM job is triggered. This job automatically executes scripts on the source VMs to ready them for conversion. These scripts remove VMware Tools and update drivers to match the target hypervisor requirements.



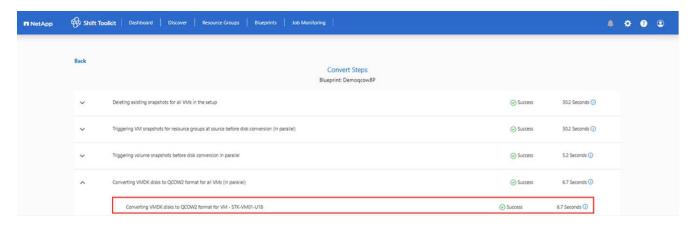
- Once the prepareVM job completes successfully (as shown in the screenshot below), the VM disks associated with the VMs are ready for conversion, and the blueprint status will update to "Active."
- Click "Convert" after scheduling the required downtime for the VMs.



• The convert operation uses a point-in-time snapshot. Power off the VM if needed and then retrigger the operation.



• The convert operation executes each operation against the VM and respective disk to generate the appropriate format.



• Use the converted disk by manually creating the VM and attaching the disk to it.



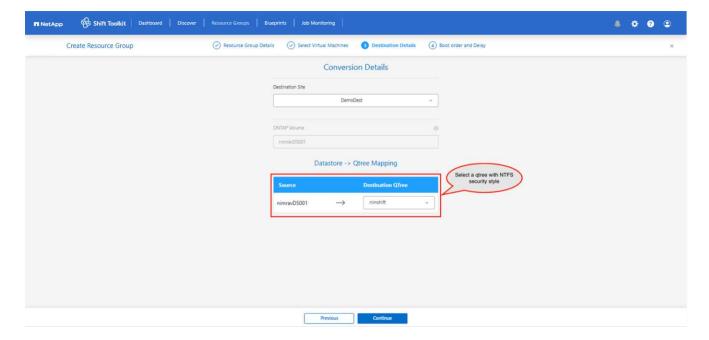


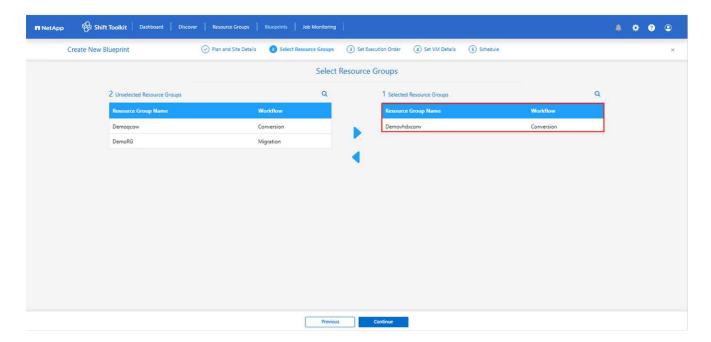
The Shift toolkit supports disk conversions only for the qcow2 format. It doesn't support VM creation or registration. To use the converted disk, manually create the VM and attach the disk.

Convert to VHDX format

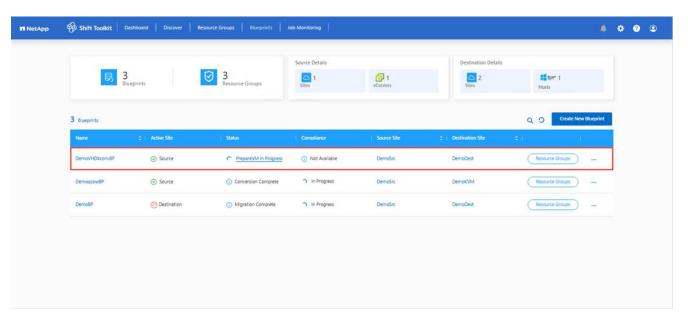
To convert the virtual disks to VHDX format with NetApp Shift toolkit, follow these high-level steps:

- Create a destination site type specifying Hyper-V as the hypervisor.
- · Create a resource group with the VMs for which the disk conversion is required

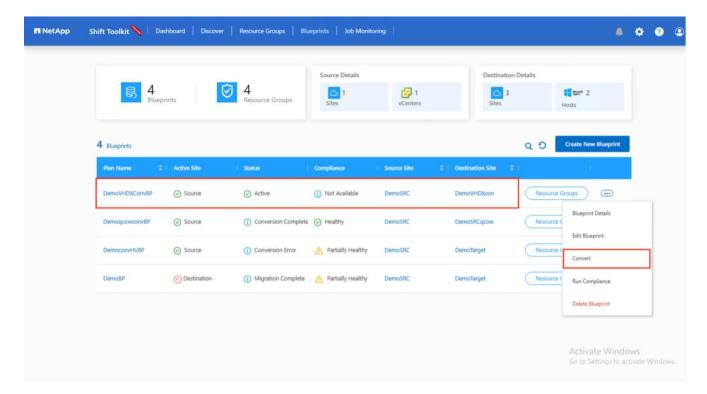




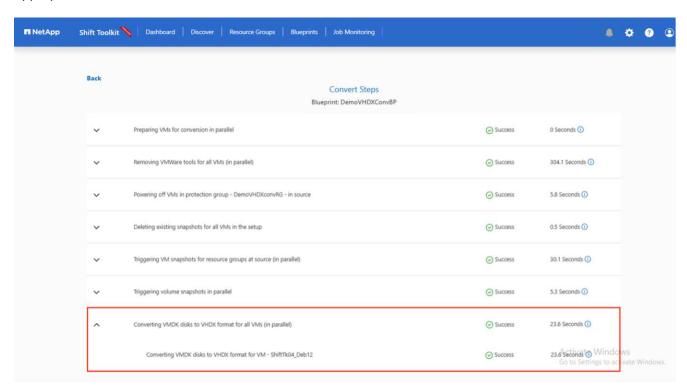
• Create the blueprint to convert the virtual disk to VHDX format. Once the blueprint is created, the preparation jobs will be automatically initiated.



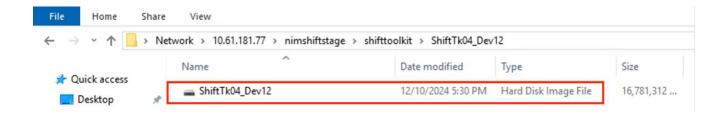
• Choose "Convert" once the required downtime for the VMs has been scheduled.



 The convert operation executes each operation against the VM and respective disk to generate the appropriate VHDX format.



• Use the converted disk by manually creating the VM and attaching the disk to it.



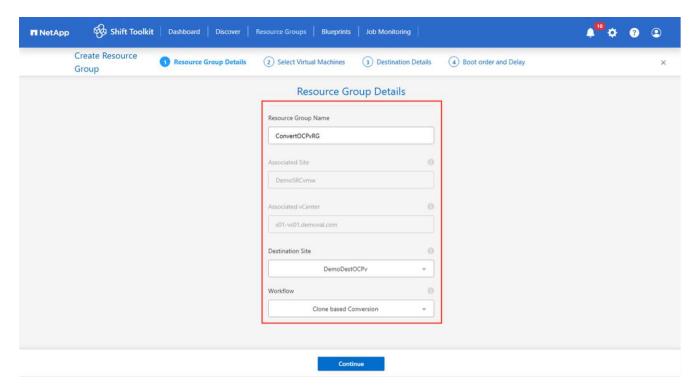


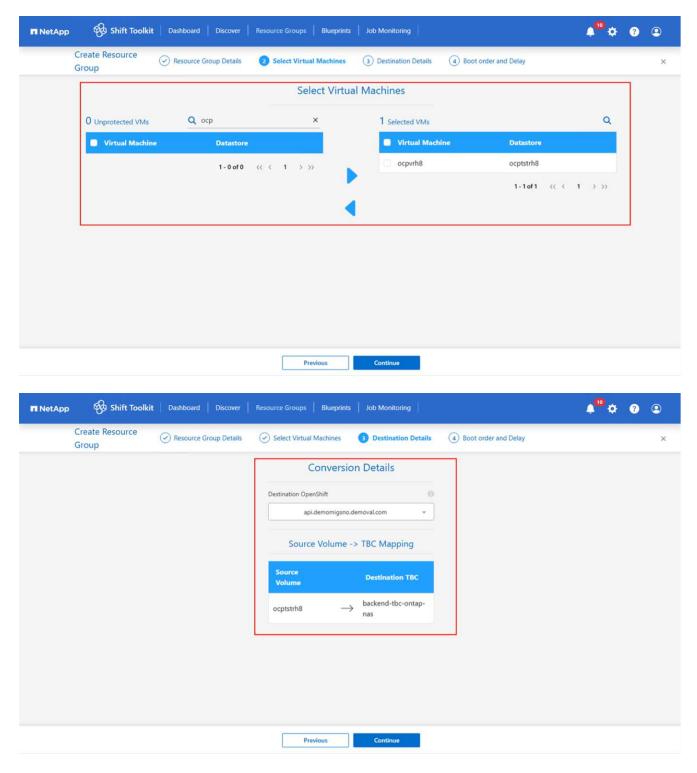
To use the converted VHDX disk in a VM, the VM must be created manually via Hyper-V manager or PowerShell commands, and the disk must be attached to it. Along with this, network should also be mapped manually.

Convert to RAW format

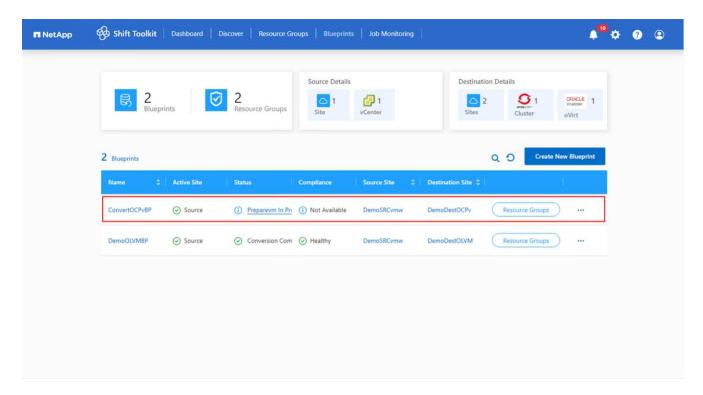
To convert the virtual disks to RAW format with NetApp Shift toolkit, follow these high-level steps:

- Create a destination site type specifying OpenShift or OLVM as the hypervisor.
- Create a resource group with the VMs for which the disk conversion is required

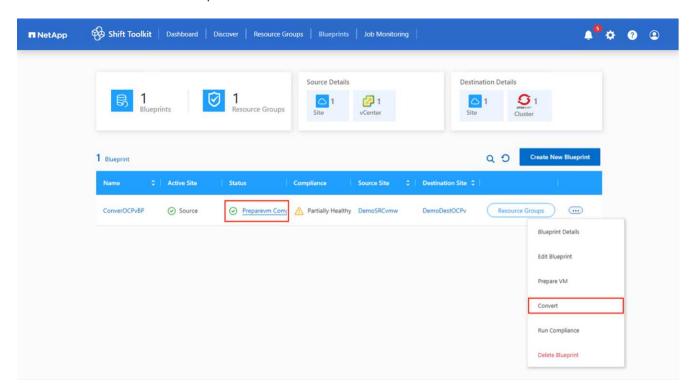




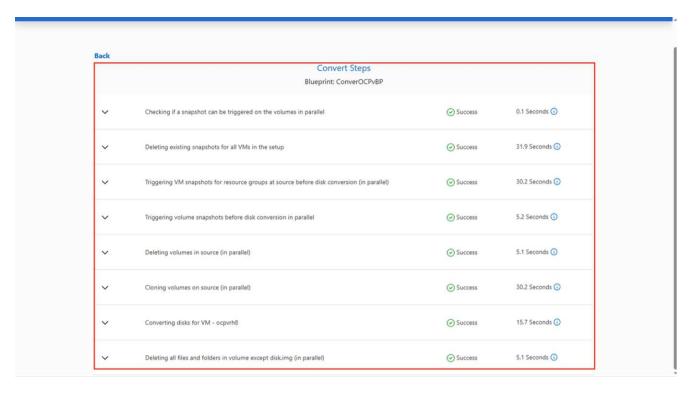
• Create the blueprint to convert the virtual disk to RAW format. Once the blueprint is created, the preparation jobs will be automatically initiated.



• Choose "Convert" once the required downtime for the VMs has been scheduled.



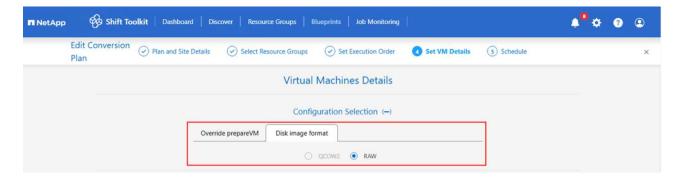
• The convert operation executes each operation against the VM and respective disk to generate the appropriate RAW format.



- · Use the converted disk by manually creating a VM.
 - For OpenShift, import the volume as PVC using the tridentctl and then create the VM using the imported disks.
 - For OLVM, head to ovirt engine URL and create a new VM by attaching the Shift toolkit converted RAW file as the OS disk. Ensure to select the appropriate interface.



For OLVM, qcow2 file format can also be used. This selection can be made during the creation of blueprint.





To use the converted RAW disk image in a VM, the VM must be created manually via OpenShift console or OC commands via YAML for OpenShift or using OLVM REST APIs/UI for OLVM, and the disk must be attached to it. Along with this, network should also be mapped manually.



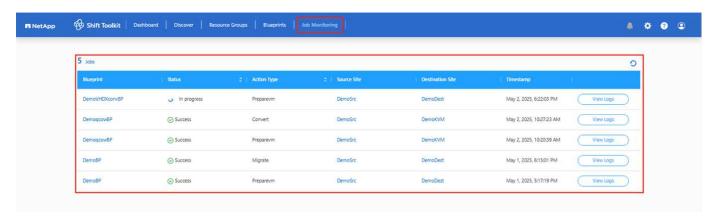
Ensure the appropriate Boot type (EFI or BIOS) is specified when manually creating the VM.

Monitor migration jobs with the Shift Toolkit dashboard

Use the Shift Toolkit Job Monitoring dashboard to track migration, conversion, and blueprint operations in real time, allowing you to quickly identify job status and resolve issues.

Job monitoring dashboard

The Job Monitoring dashboard provides a centralized view of all active and completed operations within the Shift Toolkit. Use this dashboard to monitor the progress of your migration, conversion, and blueprint jobs.

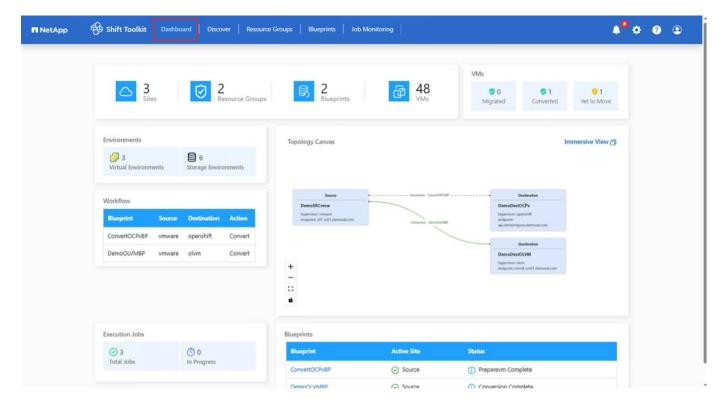


The dashboard displays key information for each job:

- Job type (migration, conversion, or blueprint)
- Current status (running, successful, failed, or partially failed)
- · Progress indicators and completion percentage
- · Number of VMs processed
- · Start and end times

Understanding job status

The intuitive interface enables you to quickly evaluate the status of all operations and identify jobs that require attention.



Job status indicators help you understand the outcome of each operation:

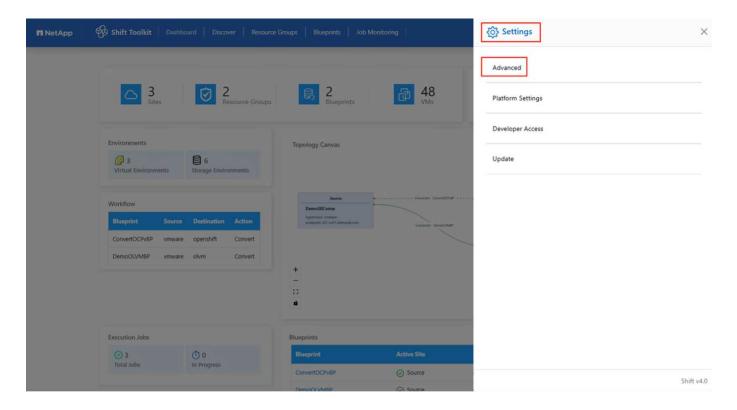
- Successful: All VMs in the job completed without errors
- Failed: The job encountered errors and could not complete
- Partially failed: Some VMs completed successfully while others encountered errors
- Running: The job is currently in progress

Use the status information to prioritize troubleshooting efforts and ensure smooth migration workflows.

Configure advanced settings in the Shift Toolkit

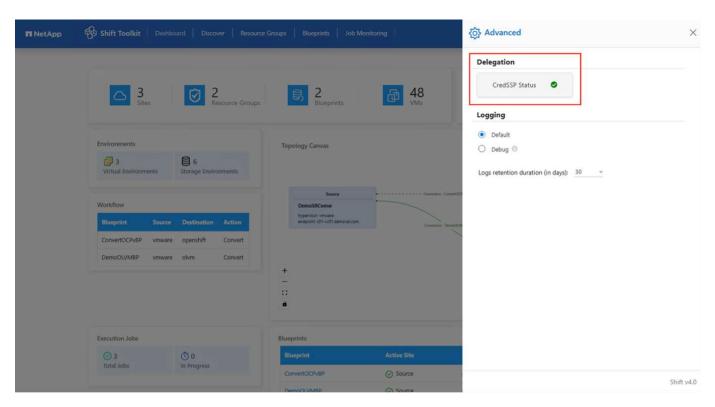
Configure advanced settings in Shift Toolkit to manage CredSSP authentication, enable logging and debugging, access REST APIs, and set up email notifications for migration jobs.

Access the advanced settings by clicking the **Settings** icon in the top toolbar.



Credential Security Service Provider (CredSSP)

The Shift Toolkit uses Credential Security Service Provider (CredSSP) to manage credential transfers during the conversion process. The Shift server runs scripts on the guest operating system of the VM being converted, passing credentials through a "double-hop" from the Shift server to the guest OS through the Hyper-V server.



Configure the Shift server as a CredSSP client

The Advanced Settings wizard automatically configures the Shift server as a CredSSP client, enabling it to delegate credentials to the Hyper-V servers.

Behind the scenes

The Shift Toolkit executes the following commands and policy configurations to set itself up as a client:

Commands executed:

- Set-Item WSMan:\localhost\Client\TrustedHosts -Value "fqdn-of-hyper-v-host"
- * Enable-WSManCredSSP -Role client -DelegateComputer "fqdn-of-hyper-v-host"

Group policy configured:

 Computer Configuration > Administrative Templates > System > Credentials Delegation > Allow delegating fresh credentials with NTLM-only server authentication

Enable this policy and add wsman/fqdn-of-hyper-v-host.

Configure the Hyper-V server as a CredSSP server

Use the Enable-WSManCredSSP cmdlet on the Hyper-V server to configure it as a CredSSP server, enabling it to receive credentials from the Shift server.

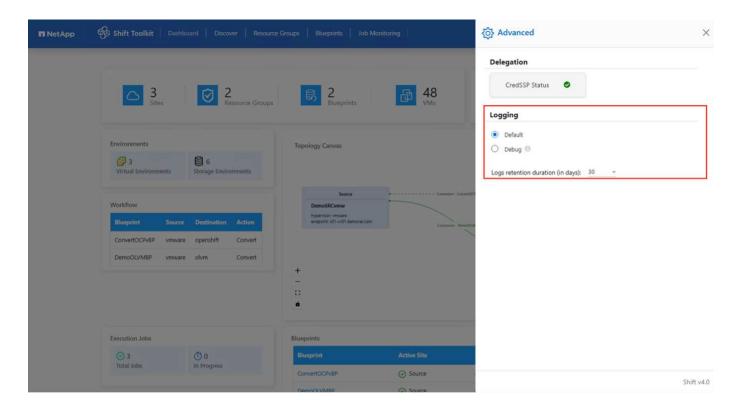
Steps

- 1. On the Hyper-V host where VMs will be provisioned by the Shift Toolkit server, open a Windows PowerShell session as Administrator.
- 2. Run the following commands:

```
Enable-PSRemoting
Enable-WSManCredSSP -Role server
```

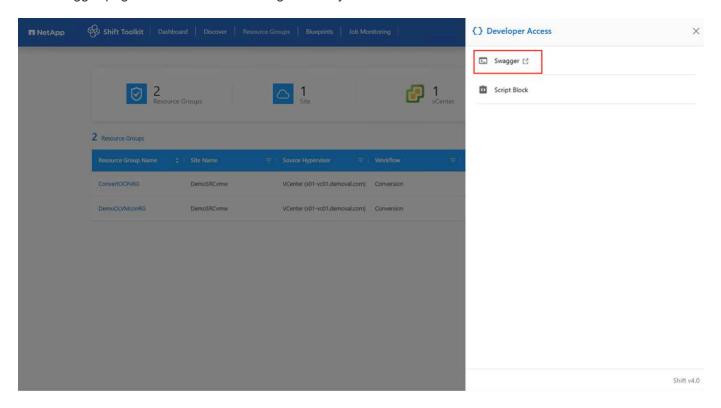
Logging and debugging

The Shift Toolkit includes default logging with a 30-day retention period. Logging can be switched to debug mode upon request from support for troubleshooting purposes.

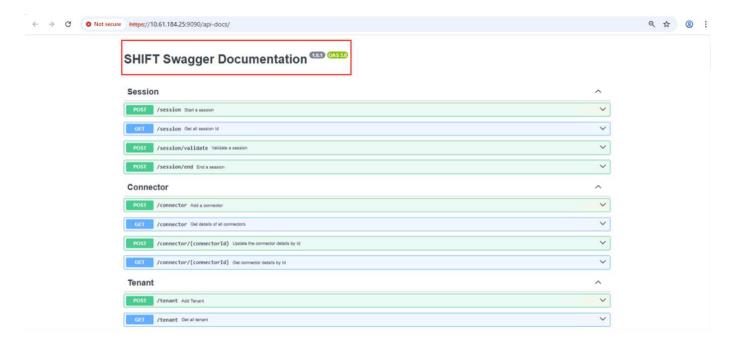


Swagger

The Swagger page in the advanced settings allows you to interact with the Shift Toolkit REST APIs.



The Shift Toolkit REST API provides programmatic access to migration, conversion, and automation capabilities. The APIs are organized by functional workflow to help you quickly find the resources you need for specific tasks.



Authentication and setup APIs

Use these APIs to establish connections, manage users, and configure authentication for the Shift Toolkit server.

Session

Manage user authentication and obtain authorization tokens for API requests:

- · Start a session
- · Validate a session
- · Get all session IDs
- End a session

User

Manage user accounts and permissions:

- · Add a user
- · Get all users
- · Change user password
- Accept EULA

CredSSP

Configure Credential Security Service Provider for credential delegation:

- Enable CredSSP
- · Get CredSSP status

Connector

Manage connections to infrastructure components:

- · Add a connector
- · Get details of all connectors
- · Update connector details by ID
- · Get connector details by ID

Tenant

Manage multi-tenant configurations:

- · Add a tenant
- · Get all tenants

Infrastructure management APIs

Use these APIs to configure and discover your source and target environments.

Site

Manage migration sites and their associated virtual and storage environments:

- · Get count of sites
- · Get all site details
- Add a site
- · Get site details by ID
- · Delete a site by ID
- · Add virtual environment to a site
- · Add storage environment to a site
- · Get virtual environment details for a site
- · Update virtual environment details for a site
- · Delete virtual environment details for a site
- · Get storage environment details for a site
- · Update storage environment details for a site
- · Delete storage environment details for a site

Discovery

Discover and inventory VMs and resources in source and target sites:

- · Discover source site
- · Get all discovery requests for source site
- · Discover target site
- · Get all discovery requests for target site
- · Get discovery steps for source site by ID

Get discovery steps for target site by ID

VM and resource management APIs

Use these APIs to inventory, organize, and manage VMs and resources for migration.

VM

Query and manage virtual machines:

- · Get VMs for a site and virtual environment in source
- · Get unprotected VMs for a site and virtual environment
- · Get VM count
- · Get protected VM count

Resource

View resource utilization and availability:

- · Get resource details for a site and virtual environment
- · Get source site resource count

Resource Group

Organize VMs into protection groups for migration:

- · Get protection group count
- · Get all protection group details
- · Add a protection group
- · Get protection group details by ID
- · Delete a protection group by ID
- · Update protection group details by ID
- · Get VMs of a protection group by ID
- Get blueprints containing the protection group

Migration and recovery APIs

Use these APIs to execute migrations, monitor compliance, and manage recovery operations.

Blueprint

Define and manage migration blueprints:

- · Get blueprint count
- · Get all blueprint details
- · Add a blueprint
- · Get blueprint details by ID
- · Delete blueprint by ID

- · Update blueprint details for ID
- · Get VMs of a blueprint
- · Get power status of VMs present in the blueprint

Compliance

Verify readiness and compatibility before migration:

- · Get compliance check result for a blueprint
- · Get compliance check final status for a blueprint
- Add on-demand compliance check for a blueprint

Execution

Monitor migration and conversion job execution:

- · Get all execution details
- · Get details of execution in progress
- · Get execution count
- · Get count of executions in progress
- · Get steps for execution ID

Recovery

Execute and manage migration and recovery operations:

- · Add new execution request for a blueprint
- · Add retry request of execution for a blueprint
- · Get execution statuses of all blueprints
- · Get execution status for blueprint ID

Automation APIs

Use these APIs to extend and automate Shift Toolkit functionality.

Script Block

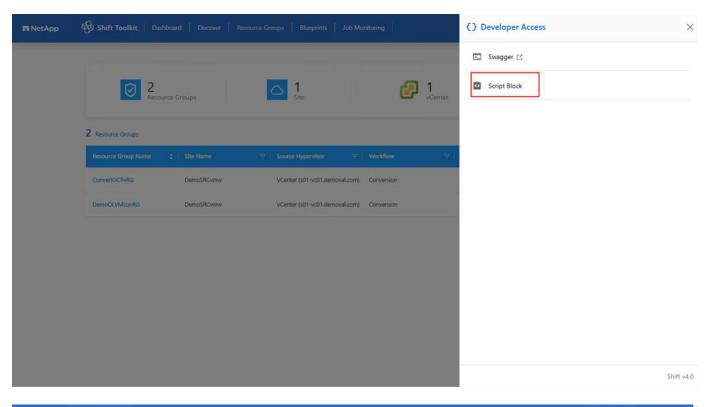
Access and execute automation scripts:

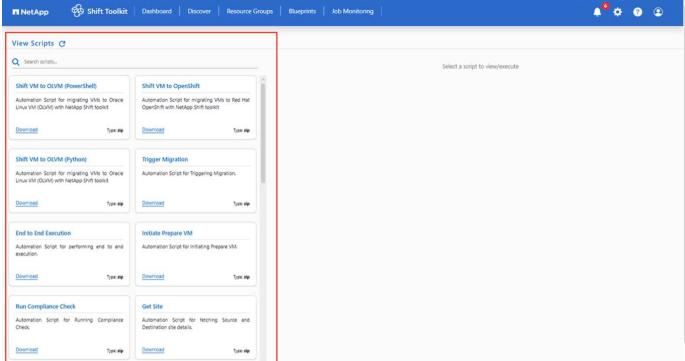
- · Get all scripts metadata
- · Get script metadata by ID
- · Get all refresh metadata
- · Execute script

Script block and automation

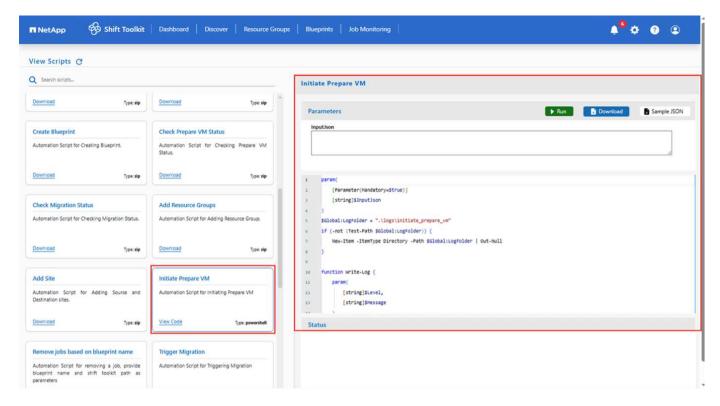
The script block within the Shift Toolkit provides sample code to help you automate, integrate, and develop features using internal and external APIs. Browse and download samples in the Code Samples section, written

by the Shift Toolkit automation team and community members. Use these samples to get started with automation, management, or integration tasks.





The following example shows a PowerShell script that deletes a specific job within the Shift Toolkit UI. While this capability is not exposed through the standard workflow, it can be accomplished using the script block. The script is also available as a batch script that can be easily downloaded and executed.



The objective of the script block is to provide sample scripts for day 0 and ongoing operations on specific hypervisors using the Shift Toolkit APIs and the respective hypervisor published APIs.

The objective of the script block is to provide sample scripts for day 0 and ongoing operations on specific hypervisors using the Shift Toolkit APIs and the respective hypervisor published APIs.

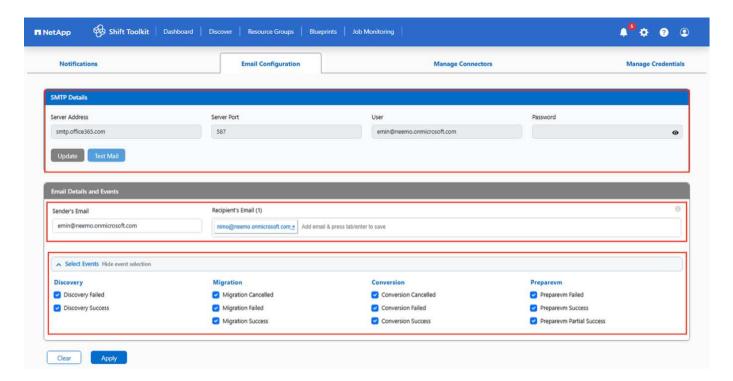
Email notifications and alerts

Configure email notifications to send alerts about discovery, conversion, or migration jobs to specified recipients. UI notifications (alerts within the interface) are also available and stored for 7 days.

Access the email notification settings from Settings > Platform Settings > Email configuration.

Steps

- 1. Log in to the Shift Toolkit UI.
- 2. Navigate to Settings > Platform Settings.
- 3. Select Email notifications and update SMTP details:
 - SMTP server address
 - Port
 - Username
 - Password
- 4. Update the recipient field and select the events from the available categories.
- 5. Click Apply.



The screenshot shows the breakdown for each notification category and event.



The email notification uses basic SMTP authentication and SendGrid in this release. A future release will support modern authentication.



The email notification uses basic SMTP authentication and SendGrid in this release. A future release will support modern authentication.

Abort and revert capabilities

The Shift Toolkit provides the option to cancel a running job at any step within the workflow. When a job is cancelled, all orphaned components are automatically cleaned up, including:

- Powering off VMs on the hypervisor if they were powered on
- · Removing disk entries from the appropriate qtree
- Deleting persistent volume claims (PVCs)

Since the Shift Toolkit does not modify the source VM in any way, rollback is straightforward, simply power on the source VM. No other rollback actions are required.

Migrate VMs from SAN environments for conversion with Shift Toolkit

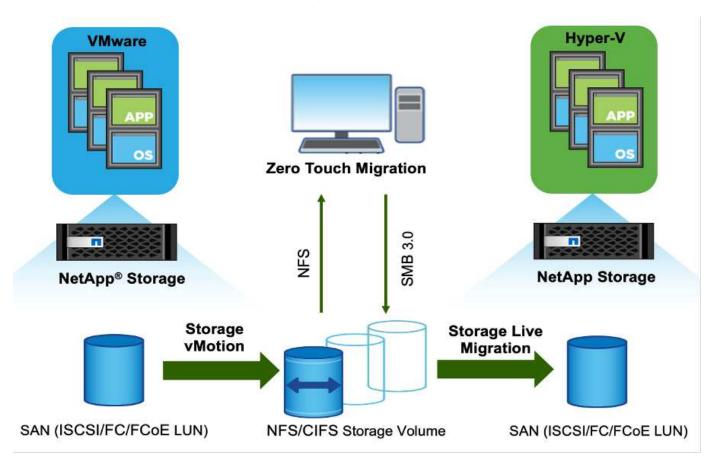
Migrate VMs from SAN-based datastores to NAS before converting them with the Shift Toolkit, using VMware Storage vMotion and Storage Live Migration to maintain business continuity.

Requirements for SAN-based VMs

The Shift Toolkit requires VMs to reside in a NAS environment (NFS for VMware ESXi) before conversion. If your VMs are currently stored on SAN-based datastores using iSCSI, Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), or NVMe over Fibre Channel (NVMe/FC), you must first migrate them to an NFS datastore.

Migration workflow for SAN environments

The following diagram illustrates the complete migration workflow for VMs stored in a SAN environment.



The migration process consists of three main phases:

Migrate from SAN to NAS (VMware environment)

Use VMware vSphere Storage vMotion to migrate VMs and their disks from the SAN datastore to an NFS datastore. This operation can be performed without VM downtime.

Convert VMs with the Shift Toolkit

After the VMs reside on the NFS datastore, the Shift Toolkit uses NetApp FlexClone technology to convert the VMs from VMware ESXi to any hypervisor. The converted VMs and their disks are placed on a qtree accessible by the respective hypervisor host.

Migrate back to SAN

After conversion, use storage migration to move the converted VMs and their disks from the qtree to a SAN-enabled volume. This allows you to maintain your SAN infrastructure in the respective hypervisor environment.

Handling processor compatibility issues

When performing live VM migration between nodes with different processor capabilities, the migration may fail due to processor compatibility checks.

To resolve this issue:

- 1. Enable the "Migrate to a physical computer with a different processor" option in Hyper-V.
- 2. Use the processor compatibility script available in the Shift Toolkit script block to configure VMs for cross-processor migration.

This setting allows VMs to migrate between hosts with different processor feature sets while maintaining compatibility.

Next steps after converting or migrating VMs using the Shift Toolkit

After you convert or migrate VMs using the Shift Toolkit, review key post-migration tasks to validate the new environment. You can verify system health, perform cleanup steps, and troubleshoot common issues using detailed examples.

Conclusion

NetApp Shift toolkit helps an administrator to rapidly and seamlessly convert VMs from VMware to Hyper-V. It can also convert just the virtual disks between the different hypervisors. Therefore, Shift toolkit saves you several hours of effort each time that you want to move workloads from one hypervisor to the other. Organizations can now host multi-hypervisor environments without having to worry about whether workloads are tied down to a single hypervisor. This capability increases flexibility and reduces licensing costs, lock-in, and commitments to a single vendor.

Next Steps

Unlock the potential with Data ONTAP by downloading Shift toolkit package and start migrating or converting the virtual machines or the disk files to simplify and streamline migrations.

To learn more about this process, feel free to follow the detailed walkthrough:

Shift Toolkit walkthrough

Troubleshooting and Known Issues

1. Trigger script for setting IP address and removal VMware tools fails for Windows VM with the following error: The credential is invalid

Error message:

Enter-PSSession: The credential is invalid.

Potential causes:

The guest credentials couldn't be validated

a. The supplied credentials were incorrect

2. Windows virtual machine encounters BSOD errors

NOTE: This is not a Shift toolkit problem, however environment related.

b. There are no user accounts in the guest

Error message:

Bluescreen error during initial boot after migration.

Potential cause:

Local group policy setup to block the installation of applications including new drivers for Microsoft Hyper-V.

a. Update the policy to allow installation of drivers.

3. No datastores listed while trying to create a resource group

Error message:

Mount paths are empty while getting volumes for mountpaths for site.

Potential causes:

The NFS volume used as a datastore is using v4.1

a. Shift toolkit filters out NFS v3 datastores during the resource group creation. NFS 4.1 or 4.2 is not supported in the current release.

4. Unable to access Shift toolkit UI after enabling SSL.

Error message:

Login failed, Network error

Potential causes:

MongoDB service not running
Using Firefox browser to access Shift UI

- a. Ensure Mongo service is running
- b. Use Google Chrome or IE to access Shift UI.
- 5. Unable to migrate VMs with encryption enabled.

Error message:

Boot failure on Hyper-V side

Potential causes:

VMDK encrytped using vSphere encryption

a. Decrypt the VMDK inside VMware and retry the operation.

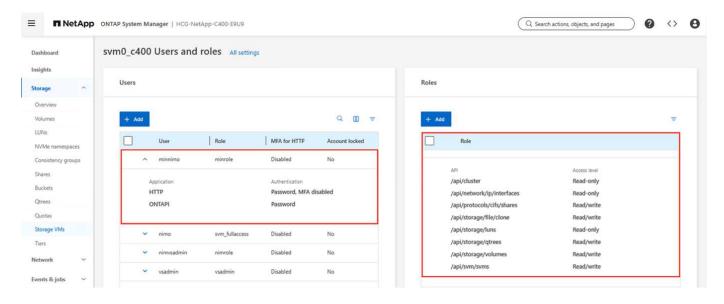
Appendix

Custom ONTAP role for Shift toolkit

Create an ONTAP role with minimum privileges so that there is no need to use the ONTAP admin role to perform operations in Shift toolkit. These minimum roles are required at the SVM level on the ONTAP Storage Side.



vsadmin can also be used.



Use ONTAP System Manager to create the role.

Perform the following steps in ONTAP System Manager:

Create a custom role:

- To create a custom role at the SVM level, select Storage > Storage VMs > required SVM> Settings > Users and Roles
- Select the arrow icon (→) next to Users and Roles.
- · Select +Add under Roles.
- · Define the rules for the role and click Save.

Map the role to the Shift toolkit user:

Perform the following steps on the Users and Roles page:

- · Select Add icon + under Users.
- Select the required username and select the role created in the previous step in the drop-down menu for Role.
- · Click Save.

Once done, use the above created user while configuring the source and destination sites within Shift toolkit UI.

Minimum permissions role required on VMware

To migrate virtual machines from VMware vSphere using Shift toolkit, create a RBAC user with the below mentioned privileges using Administration > Access Control > Roles.

DESCRIPTION

USAGE

PRIVILEGES

Datastore

- Browse datastore
- Update virtual machine files

Virtual machine

- Edit Inventory
 - Register
 - Unregister
- Interaction
 - Answer question
 - Console interaction
 - · Power off
 - Power on
- Snapshot management
 - Create snapshot
 - Remove snapshot
 - Rename snapshot

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.