Move User Guide

Move 5.1

March 4, 2024

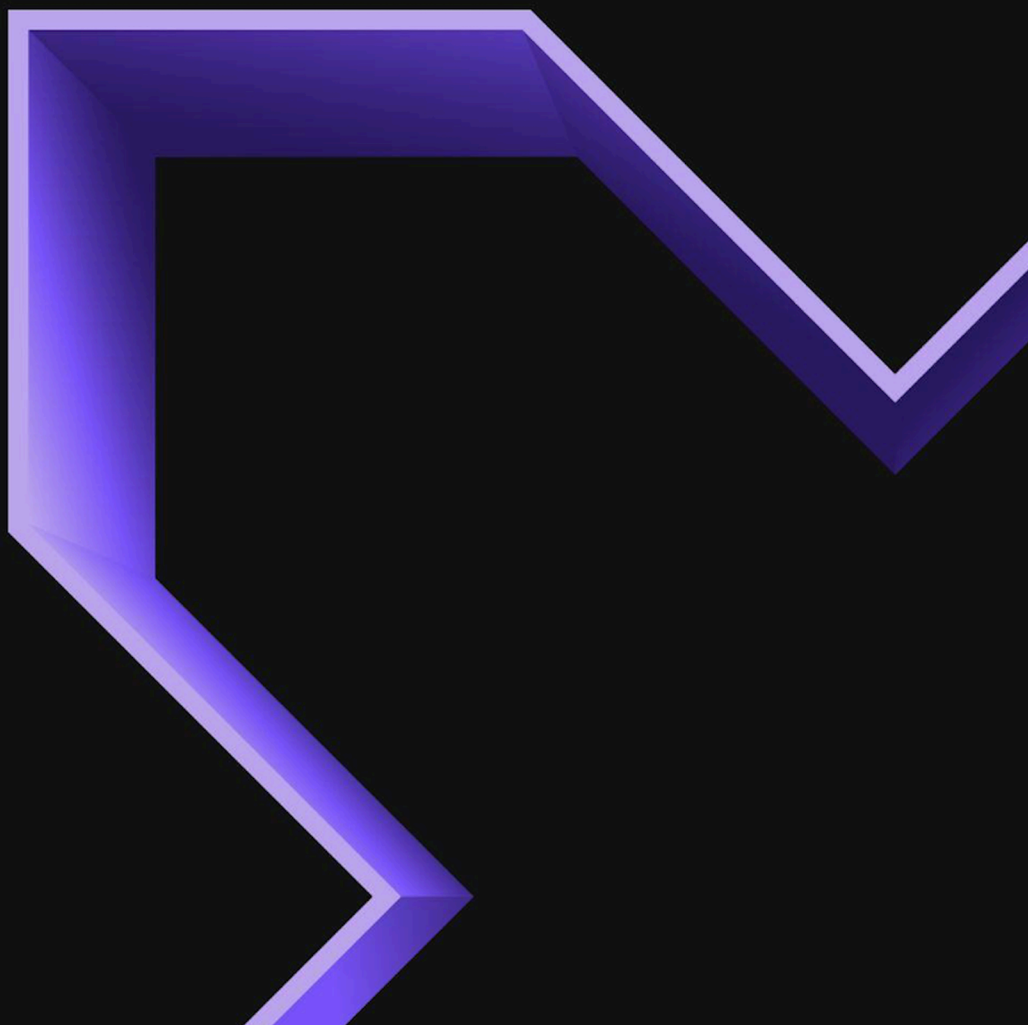

NUTANIX

Contents

MOVE OVERVIEW

Nutanix Move (Move) is a cross-hypervisor mobility solution to migrate virtual machines (VMs) and files with minimal downtime.

Move supports VM migration from the following sources to targets, where first platform is the source and second platform is the target.

- VMware ESXi (legacy infrastructure or Nutanix) to AHV
- VMware ESXi (legacy infrastructure or Nutanix) to VMware ESXi on Nutanix
- VMware ESXi to Nutanix Cloud Clusters (NC2) on AWS
- VMware ESXi to NC2 on Microsoft Azure
- Microsoft Hyper-V to AHV
- Microsoft Hyper-V to VMware ESXi on Nutanix
- Microsoft Hyper-V to NC2 on AWS
- AWS EC2 to AHV
- AWS EC2 to VMware ESXi on Nutanix
- AWS EC2 to NC2 on AWS
- Microsoft Azure Cloud to AHV
- Microsoft Azure Cloud to VMware ESXi on Nutanix
- Microsoft Azure Cloud to NC2 on Azure
- Nutanix AHV to Nutanix AHV
- Nutanix AHV to AWS EC2
- Nutanix AHV to Microsoft Azure Cloud
- Nutanix AHV to NC2 on AWS/Azure
- NC2 on AWS/Azure to Nutanix AHV
- NC2 on Azure to NC2 on Azure

Since the infrastructure underneath is different, a small downtime is incurred during cutover from any of the preceding sources to targets.

> **Note:**
>
> - As a part of this workflow, a service disruption is expected during cutover.
> - By default, Move uses the internal network of 172.17.0.0/16. The IP address assigned to Docker0 is 172.17.0.1. If you want to change the default Docker bridge IP address, refer to KB-7135 for more information.

Apart from VM migration, Move also supports the migration of files from external file servers to Nutanix file servers.

Move Architecture

Move is a distributed application which supports mobility from multiple sources such as ESXi, Hyper-V, AWS, Azure, and AHV.

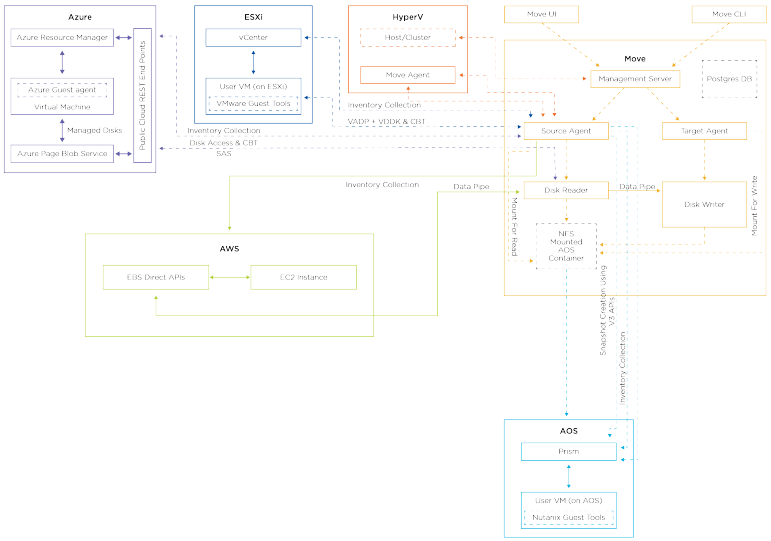

Figure 1: Distributed Architecture of Move

Components of Move

The preceding diagram depicts the distributed architecture of Move which has the following components.

- **Nutanix Move**: VM, which orchestrates the migration and can be accessed using the Move CLI or Move UI.
- **Move Agent** (only for Hyper-V migrations): A Move agent service is deployed on the Hyper-V host. This Move agent communicates with the Source Agent and interfaces with the source VMs on the Hyper-V host to facilitate the migration.

  Note: Move 3.7.1 and later versions support TLS 1.1 and TLS 1.2. If any operating system does not support TLS 1.1 and above, update the operating system with an appropriate patch before the migration or perform data only migrations for such VMs.

Move Components in various Migration Paths

The Move components are used in the following ways during various types of migration:

- **VMware ESXi to Nutanix AHV**

  When you are moving VMs from ESXi to AHV, Nutanix Move directly communicates with vCenter (vCenter is the management interface for the VMs running on the ESXi hypervisor) through the Management Server and the Source Agent. VMware Tools should be installed on the source VMs. The Source Agent collects information from the VMware library about the VMs being migrated.

- **VMware ESXi on legacy infrastructure to VMware ESXi on Nutanix**

  When you are moving VMs from VMware ESXi on legacy infrastructure to VMware ESXi on Nutanix, Nutanix Move directly communicates with Prism (management interface for the VMs running on AOS) through the Management Server and the Source Agent.

- **Microsoft Hyper-V to Nutanix AOS**

  When you are moving VMs from Hyper-V to AOS, a Move agent service is deployed on the Hyper-V host. Move agent communicates with the Source Agent and interfaces with the source VMs to facilitate migration, works with Hyper-V host to take snapshots and transfers data from source to target. Move agent is also used for collecting source inventory.

- **AWS EC2 to Nutanix AOS**

  When you are migrating EC2 instances from AWS to AOS (AHV and ESXi), the Amazon Elastic Block Store (EBS) direct APIs are used to establish a connection between AWS and Nutanix Move. For data transfer for the VM being migrated (EC2 Instance), Move takes snapshots of the EBS volumes of the VMs. The data path connection between the Amazon EBS direct APIs and Nutanix Move are used to transfer data from AWS to the target Nutanix cluster.

  > **Note:** Nutanix converts the source VM disks to AHV format.

  After the migration of the VM from the source to the target, Move deletes all EBS volume snapshots that were taken by it.

  > **Note:** No other copies of the data are stored by Move.

- **Nutanix AHV to AWS EC2**

  When you are moving VMs from AHV to AWS, Nutanix Guest Tools should be installed on the source VMs. EBS direct APIs will be used for writing data in the AWS. Prism V2 APIs are used for collecting the inventory, and Prism V3 APIs are used for managing the snapshot life cycle and querying the change block area.

- **Microsoft Azure to Nutanix AOS**

  When you are migrating a VM from Azure Public Cloud to Nutanix AHV, the Azure Page Blob APIs are used to establish a connection between Azure and Nutanix Move. Nutanix Move leverages Azure AD & Graph service (service principals) for authentication. For the actual transfer of data for the VM being migrated (guest VM), Move takes snapshots of the managed disks of the VMs. The data is securely transferred to the target Nutanix cluster using a SAS (Shared Access Signature) URI that is generated by Nutanix Move when taking snapshots. Nutanix Move groups the snapshots taken by it using a custom resource group created for each VM. These resource groups are cleaned up automatically after the completion of the migration. Automatic preparation of the VM is supported by Azure Guest Agent and uses artefacts hosted on a public Azure storage account by Nutanix. Access to these public storage account uses internal routing policy of Microsoft which does not require traffic over the internet for added security.

  > **Note:** Nutanix converts the source VM disks to AHV format

  After the migration of the VM from the source to the target, Move deletes all managed disk snapshots that were taken by it.

  > **Note:** No other copies of the data are stored by Move.

- **Nutanix AHV to Nutanix AHV**

  Move supports migrations from Nutanix AHV to Nutanix AHV.

- **Nutanix AHV to Microsoft Azure**

  When you are migrating a VM from Nutanix AHV to Microsoft Azure, Move provides the flexibility to migrate workloads between on-prem and public cloud.

- **Files Server (SMB/NFS) to Nutanix Files Server**

  When you are migrating files from external SMB/NFS files server to Nutanix files server, Move directly communicates with the target Nutanix files server through the management server and the files agent. Move uses the migration capabilities of the Nutanix files servers to perform the migration.

Move Operations

You can perform the following operations with Move.

- Migrate powered on or powered off virtual machines (VMs).

- Pause and resume migration.
- Schedule data seeding.
- Manage VM migrations between multiple clusters from a single management interface.
- Sort and group VMs for easy migration.
- Monitor progress of migrations for individual VMs as well as migration plans.
- Cancel in-progress migration for individual VMs.
- Migrate all AHV-certified guest operating systems.

  For more information about the supported guest operating systems on AHV, refer to Compatibility and Interoperability Matrix. It also indicates whether an operating system is community-supported, legacy, or deprecated on AHV.

Move Software Package

Move software package contains Move ZIP file for AHV, Offline upgrade package, and Move OVA file for ESXi.

You can download the Move software package from the Nutanix Support Portal.

> **Note:** *version-number* in the utility name is the version number of Move.

Table 1: Move ZIP file for AHV: move-<version-number>.zip

The Move software package is delivered as a zip file, which contains the following contents.

Utility name	Description
• Cli-darwin-amd64-*version-number* • Cli-linux-amd64-*version-number* • Cli-windows-amd64-*version-number*.exe	Deployer utility for the following operating systems: • macOS • Linux • Windows
move-*version-number*.qcow2	Base disk image
move-*version-number*.qcow2.md5	Base disk image checksum

Table 2: Move Offline Upgrade Package: move-offline-upgrade-<version-number>.zip

Utility Name	Description
• Cli-darwin-amd64-*version-number* • Cli-linux-amd64-*version-number* • Cli-windows-amd64-*version-number*.exe	Deployer utility for the following operating systems: • macOS • Linux • Windows

Utility Name	Description
move-offline-upgrade-*version-number*.tar.gz	This offline upgrade package is for upgrading Move VM to the latest available version without connecting to the Internet by uploading the binary.

Table 3: Move OVA file for ESXi: mmove-<version-number>-esxi.ova

Utility Name	Description
move-*version-number*-esxi.ova	The Move OVA file for ESXi is used to deploy Move on the ESXi target by uploading this file through the *ESXi host client*, *vCenter client*, or *Virtual Infrastructure Client* (VI Client).

Supported AOS, ESXi, and Hyper-V Version

For more information about the supported AOS, ESXi, and Hyper-V version with Move, refer to the *Supported AOS, ESXi, and Hyper-V Version* section in the latest version of the Move Release Notes in the Nutanix Support Portal.

Port Requirements

For information about the port requirements for Move, refer to Ports and Protocols.

Unsupported Features

This section lists the unsupported features of Move.

- IPV6
- VM names with non-English characters.
- VM names with single and double quotes.
- Windows VM installed with running antivirus software. Antivirus software prevents the installation of the VirtIO drivers.
- Guest operating systems with deduplication enabled.
- In case of Hyper-V source VMs, logical 4K aligned VHDX images are not supported.

Migration Limitations

This section lists the limitations for migration from ESXi and Hyper-V.

The following migrations are not recommended, but can be performed with some limitations.

- Exchange should be migrated by installing newer versions of Exchange Server in parallel with existing production environments, then move user mailboxes from the old system to the new one.
- Domain Controllers should be migrated by preparing a new domain controller, and then migrating the Flexible Single Master Operations (FSMO) roles.

- Move migrates database VMs similar to any other VMs; however, the migration happens at the VM level, not the application level. Move will not perform any application-level changes, disk layout updates, or best practices on the migrated VM.

  Nutanix recommends you to apply Nutanix best practices to create additional vDisks and rebalance data across those disks to take advantage of the parallel nature of Nutanix distributed storage. For Nutanix best practices, refer to Solutions Documentation.

- Time taken for migration depends on the size of the VM, data churn rate within the VM during migration, and network bandwidth/connectivity between source and on-prem data center.
- VMs managed by Citrix Hypervisor are known to have issues after migration.
- Migration to Metro clusters is not supported.
- For migrating Prism Central (PC) instances deployed on ESXi or AHV hypervisors on Nutanix, use the Prism Central-Disaster Recovery feature. Move must not be used for the migration of PC instances.

  Move can be used to migrate PC instances only if the following are true:

  - The PC instances are deployed on non-Nutanix ESXi environment.
  - PC version is 2022.6 and older, with Microservices Platform (MSP) disabled.

MOVE DEPLOYMENT

Nutanix recommends to deploy Move in the AHV or ESXi cluster by using Prism Element UI. Once the deployment is successful, you can log on to Move to perform the migration.

Note:

- As a best practice, it is recommended to deploy Move on the target cluster (AHV or ESXi on AOS).
- Deployment of Move in Prism Central is not supported using Move CLI. If a Prism Element is registered to Prism Central, IP address of the Prism Element can be used to deploy Move using Move CLI.
- For migration from any source (ESXi, Hyper-V, and AWS) to AHV target and from any source (ESXi, Hyper-V, and AWS) to NC2 on AWS target, Move should be deployed on the same destination target cluster where the VMs need to be migrated. However, for migrations from ESXi, we recommend to deploy Move on source (ESXi) if you want to migrate across geographical locations or if the latency between the source (ESXi) and the target (considering Move is to be deployed on AHV) is more than 200 ms.

Downloading Move and Invoking Move CLI

To get started with Move, you can first download and invoke Move CLI on the target cluster, and then deploy Move. If you are migrating to multiple AHV clusters, then you can deploy Move on any of the target clusters.

About this task

Note:

- Move CLI supports AHV Prism Element only.
- Use **Tab** to automatically complete the parameters. Press **Tab** to see the command completion, and press **Tab** again to enter the completion mode.
- Enter **help** along with any command, for information to display options for that command.

To download and invoke the Move CLI, do the following:

Procedure

1. Download the Move software package from Nutanix Support Portal.
2. Extract the zip file on your workstation.
3. Open the local CLI of your operating system.

4. Browse to the extracted folder location and run the following command.

```
$ binary_name -c cluster_virtual_ip_address
```

Replace the *binary_name* with the name of the binary for your operating system.

For more information, refer to Move Software Package on page 11.

> **Note:** The binaries for each operating system are similar to the following.
>
> - macOS: `./cli-darwin-amd64-version-number`
> - Windows: `cli-windows-amd64-version-number`
> - Linux: `./cli-linux-amd64-version-number`

Replace *cluster_virtual_ip_address* with the FQDN or the IP address of the cluster.

> **Note:** Use the `-u` parameter to log on. For more information, run the command `./binary_name --help`.

> **Note:** This address is either from Prism Element or the cluster VIP obtained from Prism. Prism Central is not supported in Move.

An example command looks like the following with the subsequent parameters.

- macOS deployer utility
- AHV cluster
- IP address: 10.1.1.100

```
$ ./cli-darwin-amd64-3.3.1 -c 10.1.1.100
```

The Move CLI appears.

What to do next

You can deploy Move VM once the CLI is invoked. For more information, refer to Deploying Move on AHV (CLI) on page 15

Move Deployment on AHV

Nutanix recommends to deploy the Move VM on the AHV target through CLI or Prism Element UI.

> **Note:** As a best practice, it is recommended to deploy Move on the target cluster (AHV or ESXi on AOS).

Deploying Move on AHV (CLI)

Nutanix recommends deploying the Move VM on the AHV cluster to which you want to migrate the VMs. By default, DHCP is supported. Deploying the Move VM creates, uploads, and starts the Move VM. When you invoke the `deploy-vm` command from the Move CLI with right set of options, the utility uploads the QCOW2 image in the target cluster, and deploys the Move application.

About this task

> **Note:**
>
> - Move deployment through CLI is not supported for ESXi on Nutanix.
> - While deploying Move VM by using the CLI, use the Prism Element admin user credentials.

- Network should support DHCP. If the network only supports static IP address, Move tries to fetch the IP address and if an IP address is not found, the deployment fails and you will be prompted for a clean up. If you select **Y**, complete cleanup is performed. If you select **N**, Move VM deployed will not be deleted.

  Following is an example of using a static IP address for deployment:

  ```
  » deploy-vm vm-container DM_Nutest_Ctr vm-network static
  Image is already present as a local file ./move-3.6.2.qcow2, nothing to
   download...
  Image Download complete... [OK]
  Creating image...
  Uploading file ./move-3.6.2.qcow2
  Image upload for Nutanix-Move completed [OK]
  VM Deployment completed [OK]
  VM Power on completed [OK]
  Timed out querying for 'Nutanix-Move' to get IP address
  Timed out querying for 'Nutanix-Move' to get IP address [ERROR]
  Do you want to perform cleanup ? (y/N):
  ```

  If DHCP is not supported, refer to Assigning a Static IP Address to Move on page 24 or Deploying Move on AHV (Prism Element UI) on page 17.

- Move VM CLI deployment can show that the IP address query failed, but Move VM is up and running. If Move VM is not allotted an IP address, contact Nutanix Support and refer to KB 6701.

- Use **Tab** to automatically complete parameters. Press **Tab** to see the command completion, and then press **Tab** again to enter the completion mode.

To deploy Move VM on AHV, do the following:

Procedure

1. Open the Move CLI.

   Refer to Downloading Move and Invoking Move CLI on page 14 for instructions about how to invoke the Move CLI.

2. To deploy the Move VM, run the following command, and press **Enter**.

   ```
   <cluster-name ip> >> deploy-vm vm-container storage_container vm-
   network virtual_machine_network
   ```

   `<cluster-name ip>` is the cluster name and IP address of the cluster. This prompt is displayed automatically.

   Replace *`storage_container`* with the name of the storage container.

   > **Note:** To view a list of available storage containers, press **Tab** after `deploy-vm vm-container`.

   Replace *`virtual_machine_network`* with the VM network name.

   Following is an example of deploying Move VM with DHCP:

   ```
   >> deploy-vm vm-container UserContainer-New vm-network Dynamic-Pool-10
   ```

   > **Note:** While deploying with DHCP if the IP address is not found, the deployment fails and you will be prompted for a clean up. If you select **Y**, complete cleanup is performed. If you select **N**, the deployed Move VM will not be deleted and you can later refresh to assign an IP address.

   Successful completion creates, uploads, and starts the Move VM.

What to do next

You can access the Move UI. For more information, refer to Logging in to Move UI on page 21. If DHCP is not enabled, refer to Assigning a Static IP Address to Move on page 24.

Deploying Move on AHV (Prism Element UI)

You can deploy the Move VM from the Prism Element UI. Deploying the Move VM creates, uploads, and starts the Move VM.

Before you begin

Download the Move software package from the Nutanix Support Portal.

For more information, refer to Move Software Package on page 11.

About this task

> **Note:** To deploy the Move VM, you must use the qCOW2 image.

To deploy Move VM on AHV from Prism Element UI, do the following:

Procedure

1. Log on to the cluster on which you want to deploy Move by using your Prism Element admin credentials.
2. Click the gear icon pull-down list of the main menu.
3. Click **Image Configuration**.
   The **Image Configuration** window is displayed.
4. To upload an image file to the cluster, click the **+ Upload Image** button.
   The **Create Image** window appears. Do the following in the indicated fields:
   a. **Name**: Enter a name for the image.
   b. **Annotation** (optional): Enter a description for the image.
   c. **Image Type** (optional): Select the **Disk** image type from the pull-down list.
   d. **Container**: Select the storage container to use from the pull-down list.
      The list includes all containers created for this cluster. If there are no containers currently, a **Create Container** link appears to create a container.
   e. **Image Source**: Do one of the following:
      - **From URL**: Click this option to import the qCOW2 image from the Internet. Enter the appropriate URL address in the field.
      - **Upload a file**: Click this option to upload a qCOW2 file from your workstation. Click the **Choose File** button, and then select the Move qCOW2 image to upload from the file search window.
   f. When all the fields are correct, click the **Save** button.
      The **Create Image** window closes and the **Image Configuration** window reappears with the new image appearing in the list.

      > **Note:** To verify, you can see the image upload progress in the **Recent Tasks** drop-down of the main menu.
5. Go to **Home** > **VM**.

6. Click **+ Create VM**
   The **Create VM** window appears.
7. Complete the indicated fields for creating a VM.
   - **Name**: Enter a name for the VM.
   - **Description** (optional): Enter a description for the VM.
   - **Timezone**: Select the timezone of the VM as UTC.
   - **Use this VM as an agent VM**: Select this option to make this VM as an agent VM.

     You can use this option for the VMs that must be powered on before the rest of the VMs (for example, to provide network functions before rest of VMs are powered on the host) and must be powered off and migrated after rest of the VMs (for example, during maintenance mode operations).
   - **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM.
   - **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU.
   - **Memory**: Enter the amount of memory (in GiB) to allocate to this VM.

     Note: The new VM must meet the following minimum configuration.

     - **vCPU for each Core**: 2
     - **Number of Cores**: 2
     - **Memory**: 8 GB
8. Select the appropriate option under **Boot Configuration**.
9. Remove the current CDROM disk by clicking the **X** next to the CDROM.
10. Click **+ Add New Disk** and select **Operation** > **Clone from Image Service**, and then in the **Image** drop-down, select the uploaded image. Click **Add**.
11. Click **+ Add New NIC** under **Network Adapters (NIC)**. Select the network name in the **Subnet Name** drop-down menu and then select the appropriate option in **Network Connection State** drop-down menu. Click **Add**.
    The NIC appears under **Network Adapters (NIC)** table in the **Create VM** window.

    Note:

    - NIC connected to the appropriate VM network - make sure that the Move VM can connect to both source vCenter Server as well as target AHV cluster over this network. After selecting the network on the **Create NIC** dialog box, you can provide a static IP address. You must provide a static IP address if the network has IP address management enabled and no DHCP pool is defined for it.
    - You can ping the Move VM to determine if the VM is present on the network, or for any other troubleshooting purposes.
12. Click **Save**.
13. Click the **Table** view and locate and select the new VM, and then click **Power On**.
14. Wait for the VM to detect an IP address.
    The new Move VM is powered ON.

What to do next

You can access the Move UI. For more information, refer to Logging in to Move UI on page 21. If DHCP is not enabled, refer to Assigning a Static IP Address to Move on page 24.

Move Deployment on ESXi

Nutanix recommends to deploy Move on the ESXi target by uploading the Move OVA file through the *ESXi host client*, *vCenter client*, or *Virtual Infrastructure Client* (VI Client).

> **Note:** As a best practice, it is recommended to deploy Move on the target cluster (AHV or ESXi on AOS).

Deploying Move OVA (ESXi Host Client)

You can deploy Move OVA through the ESXi Host Client for migrating VMs from ESXi to ESXi.

Before you begin

Download the Nutanix Move OVA file from Nutanix Support Portal.

About this task

To deploy Move OVA through the ESXi Host Client, do the following:

Procedure

1. Log on to ESXi Host Client.
2. Right-click **Host** in the VMware Host Client inventory, and then select **Create/Register VM**.
   The **New Virtual Machine** wizard opens.
3. On the **Select creation type** page of the wizard, select **Deploy a virtual machine from an OVF or OVA file**, and then click **Next**.
4. Click the blue pane, and select the Nutanix Move OVA file to deploy, and then click **Open**.
   The file you have selected is displayed in the blue pane.
5. Enter the VM name, NIC, and other required details, and then click **Next**.
6. Click **Finish**.

   > **Note:** Depending on your network speed, the deployment can take up to 10 minutes or more.

   The Move VM in now deployed.

Deploying Move OVA (vCenter Client)

You can deploy Move OVA through the vCenter client for migrating VMs from ESXi to ESXi.

Before you begin

Download the Nutanix Move OVA file from Nutanix Support Portal.

About this task

To deploy Move OVA through the vCenter client, do the following:

Procedure

1. Log on to vCenter.

2. Click **Deploy OVF Template**.
3. Browse the Nutanix Move OVA file, and then click **Next**.
4. Enter the name of the VM, NIC, and other required details, and click **Finish**.

   > **Note:** Depending on your network speed, the deployment can take up to 10 minutes or more.

   A success message appears when the Move VM is deployed.

Deploying Move OVA (VI Client)

You can deploy Move OVA through the VI client for migrating VMs from ESXi to ESXi.

Before you begin

Download the Nutanix Move OVA file from Nutanix Support Portal.

About this task

To deploy Move OVA through the VI client, do the following:

Procedure

1. Log on to vCenter or ESXi host.
2. Click **File** > **Deploy OVF Template**.
3. Browse the Nutanix Move OVA file.
4. Enter the VM name, Data-store, Disk Format, and Network-Mapping details.
5. Click **Finish**.

   > **Note:** Depending on your network speed, the deployment can take up to 10 minutes or more.

   The Move VM is now deployed.

LOGGING IN TO MOVE

You can login into Move VM by using both Move User Interface (UI) and Command Line Interface (CLI). This section provides detailed steps to login using both ways.

Logging in to Move UI

Once the Move VM is deployed successfully and the Move VM is started, you can login to the Move user interface (UI) using the Move VM IP address or the FQDN.

Before you begin

For static IP deployment, ensure to assign static IP address to Move.

> **Note:** For information on assigning static IP address to Move, refer to Assigning a Static IP Address to Move on page 24.

About this task

To log on to Move UI, do the following:

Procedure

1. Open a web browser, enter the FQDN or IP address of the VM.
2. (First-time log on only) If you are logging in for the first time, do the following:
   a. Read the Nutanix End User License Agreement (EULA) agreement, click the **I have read and agree to terms and conditions** option, and then click **Continue**.
   b. In the **Nutanix Customer Experience Program** screen, click **OK**.
      By participating in the Nutanix Customer Experience Program, Nutanix collects non-identifying information for product improvement. Information such as type of source and target, number of migrated VMs, Move version, operating system type of migrated VMs.
      You can opt out of Nutanix Customer Experience Program from the Move dashboard after logging in. Click the gear icon on the top-right corner, then click **Experience Improvement**. Clear the checkbox **Participate**.
   c. In the logon screen, set a password for the nutanix user in the **Enter new password** and **Re-enter new password** fields and click **Set Password**.
3. In the logon screen, type the password of the *nutanix* user and press **Enter**.

> **Note:** Default user of the Move UI is **nutanix**.

What to do next

Once logged on to the Move UI, you are now directed to the Move dashboard. For more information, refer to Move Dashboard on page 26.

Accessing Move VM with SSH

You can SSH to the Move VM.

About this task

To SSH to the Move VM, do the following:

Procedure

1. Open up an SSH terminal program.

   SSH terminal program such as *Terminal* on the Mac or *PuTTY* on Windows

   .

2. In Terminal, to SSH to the Move VM run the following command.

   ```
   $ ssh <move_vm_ipaddress>
   ```

3. Log on to the Move VM as the admin user.

   ```
   login as: admin
   ```

4. Enter the password as nutanix/4u.

   > **Note:** As a security measure, first-time SSH login into Move VM with user as admin and default password nutanix/4u requires a password change.

   ```
   admin@<<move_vm_ipaddress>'s password:
   ```

   You will be logged on to the Move VM and the command prompt will change to `admin@move on ~ $`.

FINDING MOVE VERSION

You can find the installed version of Move from the Move UI.

About this task

To find the installed version of Move, do the following:

Procedure

1. Log on to Move.
2. Click your user name in the top-right corner.
3. Click **About Move**.
   The **Above Move** window displays the version of Move.

INITIAL CONFIGURATIONS

You can do the initial configurations, such as changing the default password, assigning the static IP addresses and assigning the DHCP IP address from the static IP addresses.

Changing Admin User Password

Nutanix recommends that you secure the Move VM by changing the admin user password. After the initial log on and set up, change the password for the admin user before you SSH into the Move VM.

About this task

> **Note:** As a security measure, first-time SSH logon into Move VM with user as admin and default password nutanix/4u requires a password change.

> **Caution:** This password cannot be retrieved. Ensure that you keep this password in a secure location for your retrieval.

To change the admin user password, do the following:

Procedure

1. SSH into the Move VM with the admin credentials. For more information, refer to Accessing Move VM with SSH on page 21.
2. Change the admin password by entering `passwd`.

   ```
   admin@<<move_vm_ipaddress>'s password:
   ```

3. Complete the fields by entering the new password.

   ```
   admin@move on ~ $ passwd Changing password for user admin.
   New password:
   Retype new password:
   ```

   The window displays a confirmation, `passwd: all authentication tokens updated successfully.`

Assigning a Static IP Address to Move

If DHCP is not enabled, you can assign static IP addresses for the Move VM.

About this task

To assign a static IP address, do the following:

Procedure

1. Log on to the Prism Element UI of the cluster with the admin user credentials where Move VM is deployed.
2. Go to **Entity menu** > **Virtual Infrastructure** > **VMs**.
3. Select the VM named **Nutanix-Move**.
4. Open a remote console on the Move VM and log on with the Move admin user credentials.
   For more information about credential details, refer to Changing Admin User Password on page 24 topic.

5. Switch to the root user and enter the password of the Move VM.

```
admin@move on ~ $ rs
```

> **Note:** For the first time, the script is run automatically when the Move CLI is launched.

6. Configure the static IP address.

```
root@move on ~ $ configure-static-ip
```

7. Enter the required information as shown in the following example.

```
Do you want to configure static IPv4 address? (y/N)
y
Enter Static IPv4 Address (e.g. 192.168.1.3)
192.168.1.5
Enter Netmask (e.g. 255.255.255.0)
255.255.255.0
Enter Gateway IP Address (e.g. 192.168.1.254)
192.168.1.1
Enter DNS Server 1 IP Address (e.g. 128.91.2.13)
192.168.1.100
Enter DNS Server 2 IP Address (e.g. 128.91.2.14)
192.168.1.101
Enter Domain (e.g. my.dc.domain)
user.domain.com
```

The static IP address is assigned successfully.

Assigning DHCP IP Address from Static IP Address

If your Move VM is assigned a static IP address, you can assign the IP address from static to DHCP.

About this task

To assign an IP address from static to DHCP, do the following:

Procedure

1. Deploy Move VM using the static network pool.
2. Launch the Move VM from the Prism Element UI, and then check the IP address.
   The IP address is not assigned.
3. Delete the static NIC, and then add the DHCP NIC.
   The IP address is still not assigned.
4. To configure DHCP, run the following command:

```
root@move on ~ $ configure-dhcp
```

The DHCP IP address is now assigned.

MOVE DASHBOARD

The Move dashboard displays dynamically updated information about the migration plans for the VMs between source and target clusters.

Dashboard Options

The dashboard includes the following options:

- *Environment*. Displays all the added environments. Environments can be VMware ESXi, Nutanix AOS, Microsoft Hyper-V and Amazon Web Services. Click **+ Add Environment** to add locations to add locations for migrations between them.

  For more information, refer to Adding vCenter Server or Standalone ESXi Host Environment on page 38, Adding a Nutanix AOS Cluster Environment on page 39, Adding an AWS Environment on page 123, Adding Hyper-V Environment on page 79, and Adding an Azure Environment on page 167.

- *Create a Migration Plan*. Set up the migration plan for one or more VMs you want to migrate to the target environment. A migration plan includes scheduling options but does not start the cutover process.

  For more information, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 or Creating a Migration Plan (AHV to AWS) on page 214 or Creating a Migration Plan (Azure to AHV) on page 171 or Creating a Migration Plan (Azure to ESXi) on page 184.

- *Manage an Environment or Migration Plan*. Manage the existing environment or migration plans on the Move dashboard. You can **Refresh**, **Edit**, or **Remove** an environment, and **Start**, **Pause**, **Resume**, **Cancel**, **Edit**, or **Delete** the migration plans.

  For more information, refer to Environments and Migration Plan Management on page 259.

- *View Metrics*. View the metrics of Move appliance health.

  For more information, refer to View Metrics on page 284.

- *Bandwidth Throttling*. Create bandwidth throttling policies to limit the bandwidth available for data transfer from a specific source provider.

  For more information, refer to Move Bandwidth Throttling on page 286.

- *Upgrade Software*. Upgrade to a new version of Move to use the latest available features.

  For more information, refer to Move Upgrade Management on page 271.

- *Download Support Bundle*. Generate and download a support bundle that you can send to Nutanix Support for assistance.

  For more information, refer to Move Support Bundle Collection on page 292.

- *Experience Improvement*. Participate to improve Nutanix Customer Experience.

- *Realtime Logs*. Check the real time logs for all the Move components.

- *Appliance Settings*. Configure settings (such as snapshot intervals) for the Move appliance.

  For more information about configuring snapshot intervals, refer to Snapshot Configuration on page 290.

- *Events*. Displays the list of events happening in Move.

- *Help*. Opens the latest version of the Move documentation.

- *Rest API Docs*. Opens the latest version of the Move API documentation
- *About Move*. Check the latest version of the Move.
- *Log Out*. Sign out of Move VM.

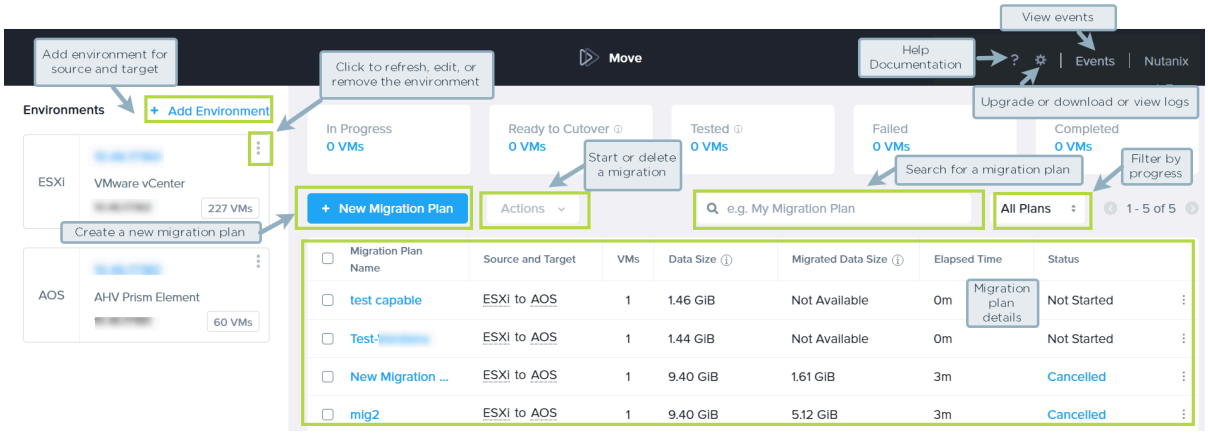

Figure 2: Move Dashboard

MOVE MIGRATION WORKFLOW

You can prepare and migrate VMs by using Move once the Move VM is deployed successfully. The steps involved during the entire migration process are as follows:

1. Log on to the Move UI.
2. Add environments for migration.

   For more information, refer to *Adding Environment* section for specific environments.
3. Create a migration plan.

   For more information, refer to *Creating Migration Plan* section for the specific source and target.
4. Prepare VMs automatically or manually.

   For more information, refer to *Preparing VMs Manually* and *Preparing VMs Automatically* section for the specific source and target.
5. Perform test migration or cutover for migration in target environment.

   For more information, refer to Creating a Test Capable VM Migration Plan on page 248 or *Performing Cutover* section for the specific source and target.
6. Manage migration plans.

   For more information, refer to Environments and Migration Plan Management on page 259.

Note: Before you migrate VMs, Move does not detect whether the source VMs have security constructs or policies (such as vTPM, BitLocker, or encrypted vDisks) enabled on those VMs. During the creation of a VM on the target hypervisor (AHV), Move does not apply any security constructs or policies. Therefore, Move does not warrant the integrity of such VMs if you decide to migrate them.

ESXI TO AHV AND ESXI TO NUTANIX CLOUD CLUSTERS (NC2)

You can prepare and migrate VMs running on ESXi hypervisor to AHV and ESXi to Nutanix Cloud Clusters (NC2) by using Move. For VM migration from ESXi to NC2, the target can be one of the following:

- NC2 on AWS
- NC2 on Azure

Move supports the following:

- Migration of vDisk partitions encrypted through cryptsetup in the Linux Unified Key Setup-on-disk-format (LUKS) format.
- (For ESXi to AHV only) Migration of vTPM-enabled VM from ESXi to AHV with Prism Central (PC) as the target. The minimum versions of software required are as follows:
  - AOS: 6.5.2
  - PC: pc.2022.9

  After the migration, the configuration of the vTPM-enabled VM at the source is retained at the target. However, the data that was stored in the vTPM VM at the source is not retained. For example, if the vTPM is enabled at the source, it will remain enabled at the target after migration.
- (For ESXi to AHV only) Migration of virtualization-based security (VBS) from ESXi to AHV. The minimum version of AOS must be 6.5.2. VBS is known as *Windows Defender Credential Guard* in AHV.

Note:

- For migration from any source (ESXi, Hyper-V, and AWS) to AHV target and from any source (ESXi, Hyper-V, and AWS) to NC2 on AWS target, Move should be deployed on the same destination target cluster where the VMs need to be migrated. Move appliance is recommended to be deployed on the target cluster (AHV). However, Move can be deployed on the source (ESXi) side for either of the following:
  - Your source (ESXi) and target (AHV) are across geographical locations.
  - The latency between your source (ESXi) and target (AHV) is more than 200ms.
- For NC2 on AWS and on Azure, static IP retention is not enabled.

Migration Considerations

You must consider the supported guest operating systems, requirements, recommendations, unsupported features, and limitations provided in this section before starting the migration process.

Supported Guest Operating Systems

Move provides full migration support for some common operating systems, and data-only support for other operating systems. Unless otherwise specified, Nutanix has qualified the following 64-bit guest operating system versions.

Full migration support migrates the data, prepares the operating system with the required device drivers and scripts for retaining the IP addresses, and recreates the VM on a target cluster. For full migration support in Windows, the UAC must be disabled.

> **Caution:** During full migration, UAC enabled on a Windows guest breaks the workflow of Move if a built-in local administrator user is not used for migration.
>
> If UAC is enabled or automatic VM preparation fails for certain VMs, you can choose to use manual preparation to prepare such VMs.

Data-Only support migrates the data and recreates the VM on the target. Data-only support requires the user to install the appropriate VirtIO drivers to each of these VMs. The following lists show the fully supported and data-only supported guest operating systems.

For more information about the supported operating systems for the VMs created using UEFI firmware, refer to Compatibility and Interoperability Matrix. It also indicates whether an operating system is community-supported, legacy, or deprecated on AHV.

> **Note:** Either Windows 7 or Windows Server 2008 R2 and earlier versions are not supported with UEFI on AHV.

Fully Supported

- Windows 7, 8, 8.1, 10, 11
- Windows Server 2008 R2, 2012, 2012 R2, 2016, 2019, 2022
- RHEL 6.3 (32-bit and 64-bit supported) to 6.10, 7.0–8.7, 9.0–9.2

  > **Note:** RHEL 6.3 is supported only with IDE as the disk controller.

- CentOS 6.3 (32-bit and 64-bit supported) to 6.10, 7.0–8.3

  > **Note:** CentOS 6.3 is supported only with IDE as the disk controller.

- Ubuntu Server and Desktop 12.04.5, 14.04.x, 16.04.x, 16.10 (32-bit and 64-bit supported)
- Ubuntu Server 12.0.4, 18.04, 19.04, 20.04
- FreeBSD 9.3 and 11.0
- SUSE 11 SP3–SP4, SUSE 12, SUSE 12 SP1–SP5, SUSE 15, SUSE 15 SP1– SP2
- Oracle Linux 6.4 and later, 7.x, 8, 8.2, 8.3, 8.4, 9.1

  > **Note:** If you face kernel panic issue on Oracle Linux versions after migration to AHV, then refer and apply the KB article 000004604 for these Oracle Linux VMs.

- Rocky Linux 9.3
- Debian 8.7, 9.4, 10.9, 10.10, 11

Data-Only Support

- Windows with UAC enabled
- RHEL 5.11 with SATA disk controller (32-bit and 64-bit supported), 6.3
- CentOS 5.11 (32-bit and 64-bit supported), 6.3
- VMs requiring PCI or IDE bus

Supported Operating Systems for UEFI Enabled VMs

Move supports the following operating systems for UEFI enabled VMs.

Table 4: Supported Operating Systems

Operating systems
Windows 7, 10
Windows Server 2012, 2016, 2019, 2022
Ubuntu 18.04, 20.04, 22.04
RHEL 7, 7.1, 7.5-8, 8.5, 8.6, 9.0, 9.1
CentOS 7, 7.1, 7.3, 7.6, 8
SUSE 12 SP3, SP5
Oracle Linux 8.4 UEFI

Support for UEFI with Secure Boot Enabled VMs

Move supports UEFI with secure boot enabled VMs.

Table 5: Supported Guest Operating Systems

Operating systems
Windows 10
Windows Server 2016, 2019, 2022
RHEL 7.7, 8.5, 8.6, 9.0-9.2
CentOS 7.3, 8.4
Ubuntu 20.0.4
Oracle Linux 8.4, 8.6, 9.1

Requirements

Before attempting to migrate VMs running on ESXi by using Move, make sure to conform to the requirements listed here.

General Requirements

Ensure to conform to the following requirements for ESXi to AHV and ESXi to NC2 on AWS migration.

- Supported browser: Google Chrome
- Ensure you have PowerShell version 4.0 or later.
- VMware Tools must be installed and up-to-date on the guest VMs for migration.
- The VMs hardware version should be 7 or above to support the Changed Block Tracking (CBT) feature.
- Source VMs must support *Changed Block Tracking* (CBT). For more information, refer to *VMware KB 1020128, Changed Block Tracking (CBT) on Virtual Machines*.
- Disks must be either sparse or flat format and must have a minimum version of 2.
- ESXi version must be minimum 5.1.
- Hosts must not be in maintenance mode.

- vCenter reachable from Move on port TCP 443.

  Note: If vCenter is running on different port, make sure https://`ip-address`:port/sdk is accessible from the Move VM.

- ESXi hosts should be reachable from Move on ports TCP 443 and TCP 902.
- Every VM must have a UUID.
- The configuration file (.vmx) of the VMs to be migrated should be present in the ESXi host.
- The VMs must be compatible with the multiple (more than one) snapshots taken by Move.
- Allow ports (TCP and UDP) 2049 and 111 between the Move network and the AHV CVM network.
- Accounts used for performing in-guest operations require **Login as Batch Job** rights in the local security policy on Windows or within the group policy. Administrator users do not have sufficient rights.

  This requirement is only applicable if the VM preparation mode is automatic.

- Ensure that you are an administrator for Windows source VMs or a root for Linux source VMs to run the source VM preparation scripts.
- If local built-in administrator user performs guest preparation, admin approval mode should be disabled. By default, admin approval mode is in disabled state. If the admin approval mode is in enabled state, refer to KB 7672 in the Nutanix Support Portal.
- Ensure that the Move user must belong in a group with **Restore files and directories** security policy.
- Ensure that the boot mode is configured as `Legacy` on the migrated target VMs if the source VMs are configured with `EFI in legacy compatibility mode`.
- Network Security Services (NSS) 3.44 is required for Linux VMs.
- Before you initiate a migration, ensure to disable backups.

Prerequisites for Linux guest VMs:

- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

Service Accounts

Move requires the following service accounts with admin privileges.

- vCenter Server
- Prism Element UI for the AHV cluster

Privileges Required on vCenter User for VM Migration

Following are the privileges that you must grant to the vCenter user for VM migration. For more information about how to grant privileges, refer to the following topics:

- Granting Privileges to a User for a vCenter Server Inventory on page 35
- Granting Privileges to a User for Specific vCenter Server Inventory Objects on page 35

  Note: Update the privileges on the user level, not on the group level.

Cryptographic Operations

- Direct Access

Global

- Disable methods
- Enable methods

Sessions

- Validate session
- View and stop sessions

Virtual machine

- Change Configuration
  - Add existing disk
  - Advanced configuration
  - Change Settings
  - Configure Raw device
  - Modify device settings
  - Remove disk
  - Set annotation
  - Toggle disk change tracking
- Guest operations
  - Guest operation modifications
  - Guest operation program execution
  - Guest operation queries
- Interaction
  - Connect devices
  - Power off
  - Power on
- Provisioning
  - Allow read-only disk access
  - Allow virtual machine download
- Snapshot management
  - Create snapshot
  - Remove snapshot

Steps to Prevent IDE to SCSI conversion During VM Migration

To prevent VMs vdisks from being converted from IDE to SCSI during VM migration, perform the following steps:

1. SSH into the Move VM with the admin credentials. For more information, refer to Accessing Move VM with SSH on page 21.
2. Switch to the root user by entering the password for the admin user.

```
admin@move on ~ $ rs
[sudo] password for admin:
```

3. Browse to the directory /opt/xtract-vm/conf.

```
cd /opt/xtract-vm/conf
```

4. Create or open the file tgtagent.json.

```
vi tgtagent.json
```

5. Create the JSON content.

   Following is an example of tgtagent.json. You can use this example JSON file, and change or remove the values of the flags as necessary.

```
{
  "AHVTargetBootConfig": {
    "IsGlobalBootConfig": true,
    "GlobalBootDeviceIndex": 1,
    "GlobalBootDeviceType": "SCSI",
    "MigrationPlanBootConfig": [
      {
        "MigrationPlanName": "MP1",
        "BootDeviceIndex": 1,
        "BootDeviceType": "IDE"
      },
      {
        "MigrationPlanName": "MP2",
        "BootDeviceIndex": 2,
        "BootDeviceType": "SCSI"
      }
    ]
  }
}
```

   Note:

   - When IsGlobalBootConfig is set to true or false, it turns custom configuration on or off respectively. By default, this value is false if JSON is not present. You can set it to true and define GlobalBootDeviceIndex and GlobalBootDeviceType based on the source VMs.
   - MigrationPlanBootConfig is optional. This is only required if you have different VMs with different boot device index order.

     Plans mentioned in MigrationPlanBootConfig overrides the above global variables on the migration plan level for a particular plan.

6. Restart the tgtagent container.

```
restart-tgtagent
```

   Note: This JSON file can be created while the migration is in progress or before the migration is initiated.

   Once IsGlobalBootConfig is set to true, disks boot order for VMs will be decided based on the values given in the JSON such as global or migration plan level as applicable.

Granting Privileges to a User for a vCenter Server Inventory

You need to grant privileges to the source and target vCenter users for successful VM migration using Move. This topic details the procedure to grant the necessary privileges to a user for a vCenter Server Inventory.

Before you begin

Ensure that you have added a vCenter single sign-on (SSO) user.

For information on adding a vCenter SSO user, refer to the *Add vCenter Single Sign-On Users* section in *VMware vSphere product documentation*.

About this task

To grant the privileges (for VM migration) to a user for a vCenter Server inventory, do the following:

Procedure

1. Login to the vCenter Server using the vSphere Client.
2. Create a role that has all the privileges required for VM migration using Move.

   Note:

   - For information on creating a role in vCenter, refer to *Create a vCenter Server Custom Role* section in *VMware vSphere product documentation*.
   - To know about the privileges required for a vCenter user for VM migration in Move, refer to Requirements (ESXi to ESXi) on page 55 or Requirements on page 31.

3. From the menu bar, select **Menu** > **Hosts and Clusters**.
4. In the vSphere Client object navigator, select the vCenter Server object for which you want to modify the permissions.
5. Go to the **Permissions** tab and click the **Plus (+)** icon.
   Add Permission window appears.
6. In the **User** field, select the user to whom you want to grant the privileges for VM migration.
7. From the **Role** dropdown menu, select the role that you created in 2 on page 35.
8. Select the **Propagate to children** checkbox and click **OK**.
   The user is assigned the selected role for the selected vCenter Server object and the data centers under it. All the privileges associated with that role are granted to the user.

Granting Privileges to a User for Specific vCenter Server Inventory Objects

You need to grant privileges to the source and target vCenter users for successful VM migration using Move. This topic details the procedure to grant the necessary privileges to a user for specific vCenter Server inventory objects.

Before you begin

Ensure that you have added a vCenter single sign-on (SSO) user.

For information on adding a vCenter SSO user, refer to the *Add vCenter Single Sign-On Users* section in *VMware vSphere product documentation*.

About this task

To grant the privileges (for VM migration) to a user for specific vCenter Server inventory objects, do the following:

Procedure

1. Login to the vCenter Server using the vSphere Client.
2. Create a role that has only the following privileges.
   - Global: Disable methods
   - Global: Enable methods
   - Sessions: Validate session
   - Sessions: View and stop sessions

   Note: For information on creating a role in vCenter, refer to *Create a vCenter Server Custom Role* section in *VMware vSphere product documentation*.
3. Create another role that has all the privileges required for VM migration using Move.

   Note: To know about the privileges required for a vCenter user for VM migration in Move, refer to Requirements (ESXi to ESXi) on page 55 or Requirements on page 31.
4. From the menu bar, select **Menu** > **Hosts and Clusters**.
5. In the vSphere Client object navigator, select the vCenter Server object for which you want to modify the permissions.
6. Go to the **Permissions** tab and click the **Plus (+)** icon.
   Add Permission window appears.
7. In the **User** field, select the user to whom you want to grant the privileges for VM migration.
8. From the **Role** dropdown menu, select the role that you created in 2 on page 36.
9. Do not select the **Propagate to children** checkbox and click **OK**.
   The user is assigned the selected role for the selected vCenter Server object only. The privileges associated with that role are granted to the user.
10. Under the selected vCenter Server object in the vSphere Client object navigator, select the data center object for which you want to modify the permissions.
11. In the **Permissions** tab, click the **Plus (+)** icon.
    Add Permission window appears.
12. Select the same user that was selected in 7 on page 36.
13. From the **Role** dropdown menu, select the role that you created in 3 on page 36.
14. Select the **Propagate to children** checkbox and click **OK**.
    The user is assigned the selected role for the selected data center object only. The privileges associated with that role are granted to the user.

What to do next

Under the same vCenter Server, if you have multiple data centers for which you want to modify the permissions, then perform 10 on page 36 to 14 on page 36 for each of those data centers.

Recommendations for Migration

Nutanix recommends the following for optimal VM migration from ESXi.

Recommendations

- Clear all the VM alerts in vCenter, if any.
- Refresh the inventory.
- Convert the templates to VMs.
- Enable access to VMs from vCenter.
- Ensure that all VMs are connected through vCenter.
- Ensure that all VMs are valid in vCenter.
- Recover VMs fully, if any are orphaned.
- Ensure the disk compatibility with CBT.
- Disable fault tolerance for VMs, if any in a fault tolerance pair.
- Install the latest version of VMware Tools on the VMs to be migrated.
- Ensure that the CBT-enabled VMs have fewer than 30 snapshots in the inventory.

> **Note:**
>
> Move migrates maximum of 8 disks from single ESXi hosts in parallel. The other VMs for migration from the same ESXi hosts are queued and only progress as and when the earlier disk data seeding completes. The limit is 32 disks in parallel at the appliance level. Refer to KB-9460 for further information.

Unsupported Features

This topic lists the unsupported features for migration from ESXi to AHV and ESXi to NC2 on AWS.

- Guest operating systems not supported by AHV.

  For more information about supported guest operating systems on AHV, refer to Compatibility and Interoperability Matrix.
- PCIE pass-through
- Independent disks
- VMs with multi-writer disks attached
- VMs with 2 GB sparse disk attached
- VMs with SCSI controllers with a SCSI bus sharing attached

  > **Note:** Change SCSI bus controller to `None`.
- Migration of VMs from ESXi standalone hosts with free license

Limitations

This section lists the limitations for migration from ESXi to AHV and ESXi to NC2 on AWS, and DNS configuration for the Move.

ESXi Migration Limitations

In addition to Migration Limitations on page 12, following are the limitations for ESXi migration.

- Migration of 100 VMs in a plan is qualified and supported for ESXi migration.
- VMs migrated with more than one network interfaces might not retain all the IP addresses.

  Workaround: Manually assign IP addresses after migration.
- After migration, Windows VMs with connected NICs only will retain their IP address. Those with disconnected NICs will not retain their IP address.
- (For ESXi to NC2) Retention of static IP address is not supported for VM migrations.

DNS Configuration for Migration from ESXi

DNS is configured during deployment and resolves the following.

- FQDN of the source ESXi host if the host is added using its FQDN in vCenter
- FQDN of the source vCenter if it is added using its FQDN
- FQDN of the target cluster

Warnings and Cautions

- Source VM disks attached to mix of PVSCSI and LSI adapters might get different device names (`sda`, `sdb`, and so on).

  Note: For Linux VMs, manually edit `fstab`. Then, arrange the correct order for the UUIDs.

Adding vCenter Server or Standalone ESXi Host Environment

While creating a migration plan for migration from ESXi to any target, ensure to add at least one vCenter Server or standalone ESXi host environment for migration.

About this task

Note: This procedure is only applicable for migration from ESXi.

To add a vCenter Server or standalone ESXi host environment, do the following:

Procedure

1. Log on to the Move UI.

2. Click **+ Add Environment** under **Environments**.
   The **Add Source Environment** window appears.

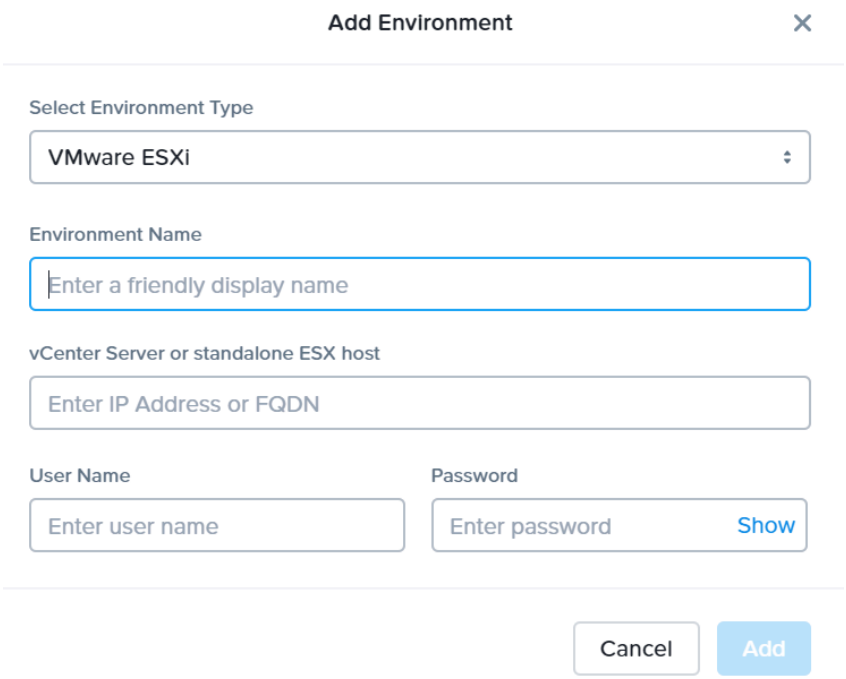

Figure 3: Add VMware ESXi Environment Dialog Box

3. Select **VMware ESXi** as the environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the ESXi environment.
   b. **vCenter Server or standalone ESX host**: Enter the IP address or the FQDN of the vCenter Server, or the IP address of the ESXi host.
      If you do not have a standard port, use the custom port for vCenter in the `vCenter IP address:custom port number` format. For example, if you are using IP address, use the format `10.136.72.150:8443` or if you are using FQDN, use the format `vcenter.nutanix.com:8443`.
   c. **User Name**: Enter the username or User Principal Name (UPN) for logging on to the vCenter Server.
   d. **Password**: Enter the password for logging on to vCenter Server.

   The VMware ESXi environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

You can now add Nutanix AHV cluster environment. For more information, refer to Adding a Nutanix AOS Cluster Environment on page 39

Adding a Nutanix AOS Cluster Environment

While creating a migration plan, AHV AOS cluster can be added as both source or target. If you want to use ESXi on Nutanix cluster as target, then add the corresponding AOS cluster environment for ESXi.

About this task

> **Note:** When you add a AOS cluster, Move VM IP address with the subnet 255.255.255.255 is added to the global NFS allowlist.
>
> **Caution:** Modifying the NFS allowlist of the destination container disables inheritance of the NFS allowlist at the global level and lead to issues with Move operations.

To add a Nutanix AOS cluster environment, do the following:

Procedure

1. Log on to the Move UI.
2. Click **+ Add Environment** under **Environments**.

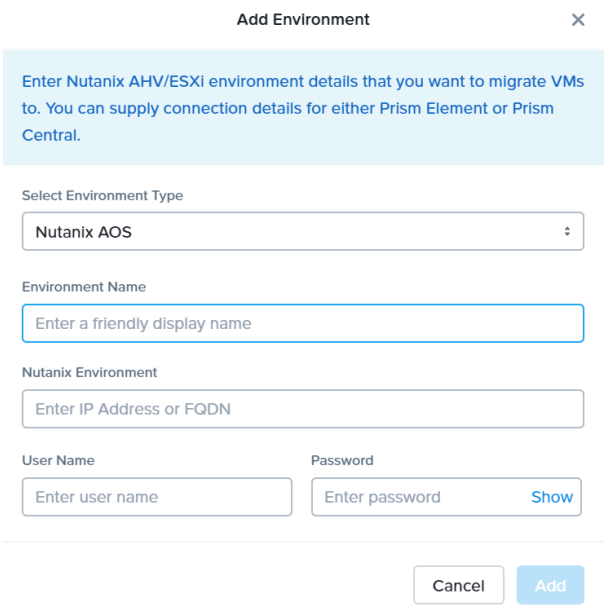

Figure 4: Add AOS Environment Dialog Box

The **Add Environment** window appears.

3. Select **Nutanix AOS** as the target environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the NC2 on AWS and AHV or VMware ESXi on Nutanix environment.
   b. **Nutanix Environment**: Enter the IP address or the FQDN for the target Prism Central or Prism Element.
   c. **User Name**: Enter the username for logging on to the target Nutanix environment.
   d. **Password**: Enter the password for logging on to the target Nutanix environment.

5. (Only for ESXi to ESXi migration) Enter credentials for registered vCenter.

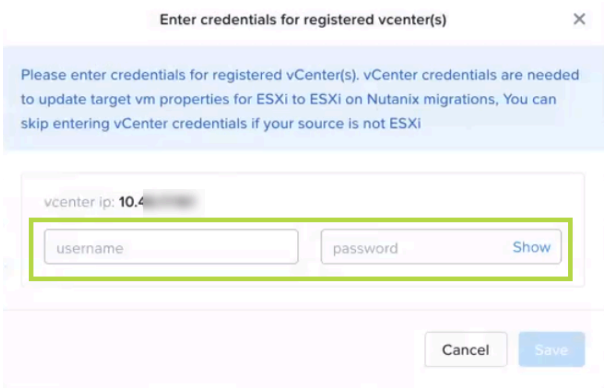

Figure 5: vCenter Credentials Dialog box

vCenter credentials are required to update the target VM properties for ESXi to ESXi on Nutanix migration.

Note:

- You can skip adding vCenter credentials if your source is not ESXi.
- To retain the VM properties on the target VM after migration, be sure to provide the target vCenter credentials. The following properties will be retained:
  - SCSI controller types
  - Network adaptor type
  - MAC address
  - Video card
  - Memory overcommit variables
  - Tool upgrade flag
  - Sync time with host flag
  - Disable acceleration flag
  - Enable logging flag

The AOS environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

Once you have added the environments, you can proceed to create the migration plan. For more information, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan (Hyper-V to ESXi) on page 103 or Creating a Migration Plan (Azure to AHV) on page 171 or Creating a Migration Plan (Azure to ESXi) on page 184

Creating a Migration Plan

You can create a migration plan to seed the data, cutover, and monitor the VMs. You can create the migration plan in Move UI without initiating the cutover process.

Before you begin

Ensure that you have added ESXi and AOS environments.

For more information, refer to Adding vCenter Server or Standalone ESXi Host Environment on page 38 and Adding a Nutanix AOS Cluster Environment on page 39

About this task

Note:

- This procedure is only applicable for migration from ESXi to AHV and ESXi to NC2 on AWS.
- If you are logging in for the first time, log on to the Move UI with your default credentials.
- You must have admin user credentials to complete the migration process.
- If you restart the management server, scheduled VM migration does not begin automatically.
- If the source boot is set to UEFI, set up the boot device manually in the VM post migration for the following operating systems.
  - CentOS 7.4 2, 6.8 3, 8, 8.1, 8.2
  - Ubuntu 12.04 4, 19.04
  - OEL 7 5
  - RHEL 6.8, 8.1

  For more information about setting up the boot device, refer to AHV Administration Guide.
- While migrating Prism Central to AHV, the DHCP IP address of the Prism Central is not retained post migration, and you have to reconfigure the IP address. IP address must be same before and after the migration for proper connectivity between the Prism Central and the Prism Element.
- When Move encounters a large source VM disk of size greater than 2 TB that belongs to one ESXi host, Move prioritize the migration of this disk and does not migrate any other disks from the same host in parallel. Only after the large VM disk migration gets completed, Move migrates the other disks from the same host. Meanwhile, Move migrates the VM disks belongs to other source ESXi hosts in parallel.
- If you cancel or discard the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.

  For information on canceling an ongoing migration, see Pausing or Canceling a VM Migration on page 257.

  For information on performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

To create a migration plan, do the following:

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.

2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.

3. On the Move dashboard, click **Create a Migration Plan**.

   > **Note:** If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.

4. Enter the new migration plan name, and then click **OK**.

   > **Note:** You can edit the migration plan name by clicking the pencil icon next to the migration plan name.

   The **New Migration Plan** window appears.

5. Complete the following fields, and then click **Next**.

   a. **Select a Source**: Select any vCenter Server or standalone ESXi host as source for migration. Once you select the source, an appropriate target appears.

      > **Note:** You might at times see a message relating to inventory collection as shown below. During this time, the environments undergoing Refresh will not be available for selection. Such environments do not automatically show up for selection once the inventory collection is over. In case you wish to select one of the

environments undergoing Refresh (not available for selection), you will need to select **Cancel** and wait for inventory collection to get over and then create the migration plan again.

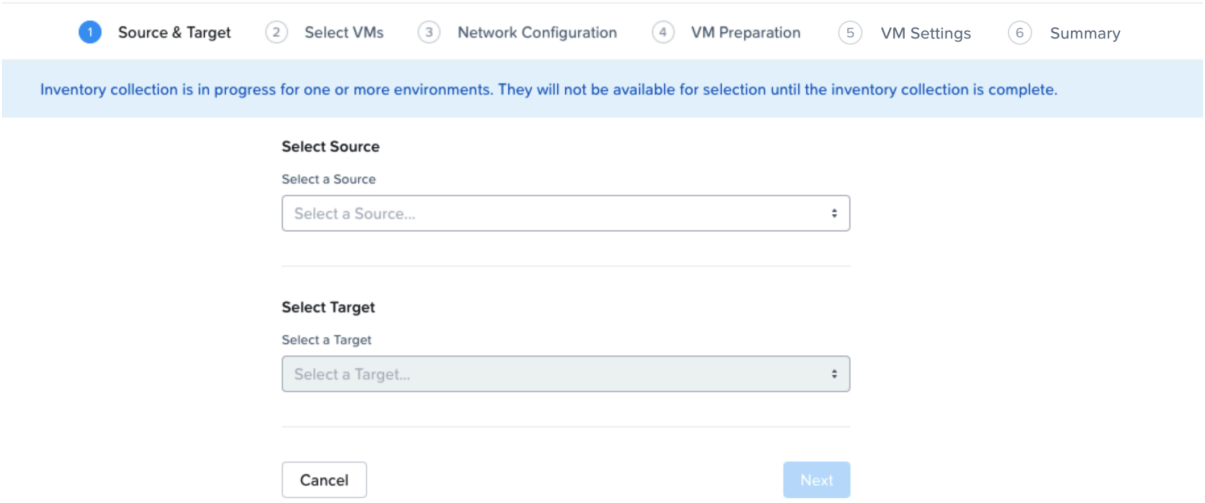

Figure 6: Inventory Collection Message

b. **Select a Target**: Select any AHV or NC2 on AWS as target for the migrating VMs.

c. **Target Project** (optional): Select the project you want as the target.
This field is available only with Prism Central and if AHV is the target.

d. **Target Owners**: Select the owners for the selected target project.
This field is available only if a target project is selected.

e. **Target Cluster**: Select the cluster on which you will migrate the VMs.

f. **Target Containers**: Select the container on which you will migrate the VMs.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click **+Add All**, and then click **Next**.

> **Note:** You cannot add more than 100 VMs in a single migration plan.

You can filter the VMs by name in the **Filter** bar. You can also sort the VM types by selecting the appropriate column. The VMs for which migration has failed are not displayed. To show the entire list of VMs, select **All VMs** from the drop-down list. To show the configuration of VMs, select **Configuration** from the drop-down

list. A question mark icon appears beside the unavailable VM, which displays more information about that VM and indicate why the VM cannot be migrated.

Note:

- If the source VM has RDM disks in physical compatibility mode, then those disks are converted to virtual compatibility mode during source VM preparation. By default, power cycle is enabled for the VMs with physical RDM disks. Move performs the following:

  1. Shut down the source VM.
  2. Convert physical RDM disks to virtual compatibility mode.
  3. Start the VM.

  If power cycle is not enabled, then Move converts the physical RDM disks to virtual with VM in powered on state.

- Migrate VMs retain their power state on the destination cluster.

The selected VMs are displayed in the sidebar.

7. In the **Network Configuration** screen, select the target network from the drop-down.

   - [Applicable only if Prism Central is used] Move provides the option to select VPC-based or VLAN-based target subnets. Based on the selection of a VPC or VLAN ID as the target network, the respective subnets are listed in the target subnet drop-down menu. Select the required subnet from the drop-down menu.

     Note: Overlay subnets which do not have IP address pool(s) associated will be disabled in the subnet drop-down menu.

   For performing test migration, refer to Creating a Test Capable VM Migration Plan on page 248 section.

   Click **Next**.

8. In the **VM Preparation** screen, select one of the following VM preparation modes.

   » **Automatic**. Move automatically runs scripts on the source VMs to prepare them for migration.

   » **Manual**. Move displays the VM preparation scripts for Windows and Linux VMs.

   These scripts prepare the source VMs by performing the following.

   - Installs NTNX VirtIO driver.

   - Runs the IP address retention script.

   - Runs Set SAN policy script.

   - Runs uninstall VMware tools script.

   If you do not want to uninstall VMware tools after migration, refer to Step 9b.

   » **Mixed**. Move allows you to select the VM preparation mode for each VM. This setting allows you to manually prepare one set of VMs and automatically prepare another set of VMs in the same migration plan. If you want to have such a hybrid combination of **Automatic** and **Manual** VM preparation modes, proceed to the next step.

   Note: The preparation mode automatically switches to **Mixed** if you change the mode of preparation for a set of VMs under **Change Settings** and have a mix of **Automatic** and **Manual** VM preparation modes. You cannot manually select the **Mixed** option from the **Preparation Mode** drop-down list.

   To perform data-only migration, refer to Performing Data-Only Migration on page 50.

9. (Optional) Under the **Guest Operations** section in the **VM Preparation** screen, do one or more of the settings:

   a. **Retain static IP addresses from source VMs**: Retains static IP addresses from the source VMs to the migrated VMs on AHV. By default, this option is enabled. Clear the checkbox if you do not want to retain the static IP addresses from the source VMs on the target VMs. If you disable this option, the static network is converted to DHCP network.

      For **Manual** preparation mode, if you do not want to do not want to retain the static IP addresses from the source VMs on the target VMs, do the following:

      - For Windows VMs, change the argument from `$retainIP = $true` to `$retainIP = $false` in the script.
      - For Linux VMs, remove the argument --retain-ip from the script.

   b. **Uninstall VMware tools on target VMs**: Uninstalls the VMware tools from the migrated VMs on AHV after migration. By default, this option is enabled. Clear the checkbox if you do not want to uninstall the VMware tools.

      > **Note:** Uninstalling VMware tools from the migrated VMs is on a best-effort basis.

      For **Manual** preparation mode, if you do not want to uninstall the VMware tools, do the following:

      - For Windows VMs, change the argument from `$uninstallVMwareTools = $true` to `$uninstallVMwareTools = $false` in the script.
      - For Linux VMs, remove the argument --uninstall-vmware-tools from the script.

   c. (Only for Automatic preparation mode) **Bypass Guest Operations on Source VMs**: Select this check box to bypass the guest operating system changes.

      You can select this option to override your migration to data-only migration.

      > **Note:** If you bypass guest operations on source VMs, **Retain static IP addresses from source VMs** and **Uninstall VMware Tools** configuration will not be applicable.

   d. (Only for Manual preparation mode) **Re-Generate Script**: This option gets enabled only when there is a change in the selection of the above-mentioned settings under the **Guest Operations** section. If enabled, click this option to re-generate the VM preparation scripts for both Windows and Linux VMs.

10. Do one of the following based on whether the preparation mode selected:

    » Automatic preparation mode: Provide the credentials of the source VMs under **Windows VMs** or **Linux VMs**, depending on the type of the source VM.

    Note:

    - For Windows VMs - Move supports only username-password sign-in option for authentication. It does not support username-PIN sign-in option.
    - For Linux VMs - Apart from credentials, Move supports PEM file for authentication. Select the **Use Private (.PEM) file to authenticate** option and upload the private key.
    - If you want to retain the static IP addresses, provide a common set of credentials for your selected Windows or Linux VMs.
    - If you face any issues while using the .PEM file for authentication, refer to KB 7090.
    - Currently, the default location where the preparation scripts are stored is the /tmp folder.

      If the /tmp folder is mounted as noexec, then Move will fallback to the /var/tmp folder. If the /var/tmp folder is also mounted as noexec, then Move will fallback to the /usr/tmp folder.

    » Manual preparation mode: Copy the scripts and manually run them on the respective source VMs, and then click **Next**.

    Note:

    - For running the script, use the Windows built-in administrator credentials for the Windows VMs and use root user for the Linux VMs.
    - If you have not run the preparation script in the source VMs, Move performs data-only migration.

      For more information, refer to Performing Data-Only Migration on page 50.

11. In the **Override individual VM Preparation** section, click **Change Settings** to override the **Guest Operations** settings (configured in the above steps) for the individual VMs. You can also edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

    Note: If you do any of the following for a VM, then copy the newly generated scripts of that specific VM and run them on the source VM:

    - Change the **Mode of Preparation** of a VM to **Manual**.
    - Change any of the guest operation settings of a VM with the preparation mode set to **Manual** (an icon appears next to the VM Name prompting to regenerate a new guest script).

    Access the newly generated VM preparation script for that VM by selecting the Copy Scripts icon under the **Actions** column.

    After making the required changes, click **Done**.

    Then, click **Next**.

12. (Optional) In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

    a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

    b. **Timezone**: Set the timezone as the hardware clock of the VMs in target.

    If set to default, Move configures UTC timezones for Linux VMs and cluster timezones for Windows VMs.

    > **Note: Timezone** is available only if **Prism Central** was selected as the target in the **Source & Target** screen.

    c. **Retain MAC Addresses from the Source VMs**: Select this check box to retain the MAC addresses from the source VMs.

    d. **Skip CDROM addition on target VMs**: Select this check box to skip the CDROM addition on the target VMs.

    e. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

    Only those categories that have values are available for selection.

    > **Note: Category Settings** is available only if **Prism Central** was selected as the target in the **Source & Target** screen.

    f. **Enable Memory Overcommit**: Select this option to enable memory overcommit on the target VM.

    For more information on memory overcommit deployment, refer to AHV Administration Guide.

    g. **VM Migration Type**: Select one of the following VM migration types.

    At the VM level, some of the target VM properties can be customized manually after the migration plan is created. For information on manually customizing the target VM configuration, refer to Customizing the Target VM Configuration on page 255.

    - **Configure Target VM Properties**: The target VM synchronizes with the source VM properties at the time of migration plan creation. Selecting this option allows you to edit the target VM properties at the

VM level (during migration). For information on editing the target VM properties, refer to Customizing the Target VM Configuration on page 255.

> **Note:** At the VM level, only the following properties can be edited:
>
> - Target VM name
> - Number of vCPUs
> - Number of cores per vCPU
> - Memory
> - Power state

- **Retain Source VM Properties**: The target VM synchronizes with the source VM properties whenever Move refreshes the source VM configuration details. Only the customizable properties are refreshed on the target. Selecting this option does not allow you to edit the target VM properties at the VM level.

  > **Note:**
  >
  > - The source VM properties are refreshed in the following ways.
  >   - (Manually) When you click the **Refresh Source VM Properties** button.
  >   - (Automatically) When you start a migration plan.
  >   - (Automatically) When you initiate a cutover.
  > - When you start a migration plan, Move refreshes both source VM and target VM properties by default. However, it will not refresh the target VM properties at the start of a migration plan if you modified the target VM properties after migration plan creation.

h. **Settings for individual VMs**: Click **Change Settings** to configure settings such as timezone, retain MAC addresses, VM priority, and skip CDROM addition for individual VMs. You can also search the VM by typing the name of the VM and change the settings.

i. **Schedule Data Seeding**: Check this check box to select the date and time for migration.

> **Note:** For migration of VMs with multiple NICs, static IP address configurations are applied correctly if the MAC address retention is applied from source VMs, otherwise best effort IP address configuration is done by mapping one NIC at a time.

This action does not affect the VM data migration but requires you to manually prepare the guest operating system with the necessary AHV drivers prior to cutover. In addition, if you bypass the guest operations, you have to take care of the static IP address retention separately.

13. In the **Summary** screen, choose one of the following, and then proceed review the VM migration summary.

    » **Back**: To edit the information, click this option.

    » **Save**: To save the migration plan, click this option.

    For more information about how to start the migration later, and check the migrated VM status and details, refer to Environments and Migration Plan Management on page 259.

    » **Save and Start**: Click this option to save the migration plan and begin the migration immediately.

    The seeding process for migration begins. You can monitor this information by selecting **Status** for the migration plan.

    > **Note:** The seeding process can take several minutes depending on the number of VMs.

What to do next

- To customize the target VM configuration at the VM level, refer to Customizing the Target VM Configuration on page 255.
- If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover on page 51.

Performing Data-Only Migration

Move performs data-only migration when you select **Automatic** preparation mode while creating a migration plan, and bypass the guest operations or does not provide the source VM credentials while preparing a migration plan. Or when you select **Manual** preparation mode while creating a migration plan, and do not run the preparation script in the source VMs. In data-only migration, Move skips the source VM guest operating system preparation tasks which includes installing VirtIO driver and copying of the scripts to retain the IP address. It also skips the uninstallation of VMware tools after migration.

About this task

> **Note:** Data-only migration is only supported for the following migrations:
>
> - From ESXi to AHV and ESXi to NC2 on AWS
> - From Hyper-V to AHV and Hyper-V to NC2 on AWS
> - From Hyper-V to ESXi
> - From AHV to AHV

To perform data-only migration, do the following:

Procedure

1. In the **VM Preparation** screen, if you select **Automatic**, then proceed without providing the credentials for the source VMs or select the **Bypass Guest Operations on Source VMs** check box or if you select **Manual**, do not run the preparation script in the source VMs.
   The following message appears when the **Automatic** preparation mode is selected,

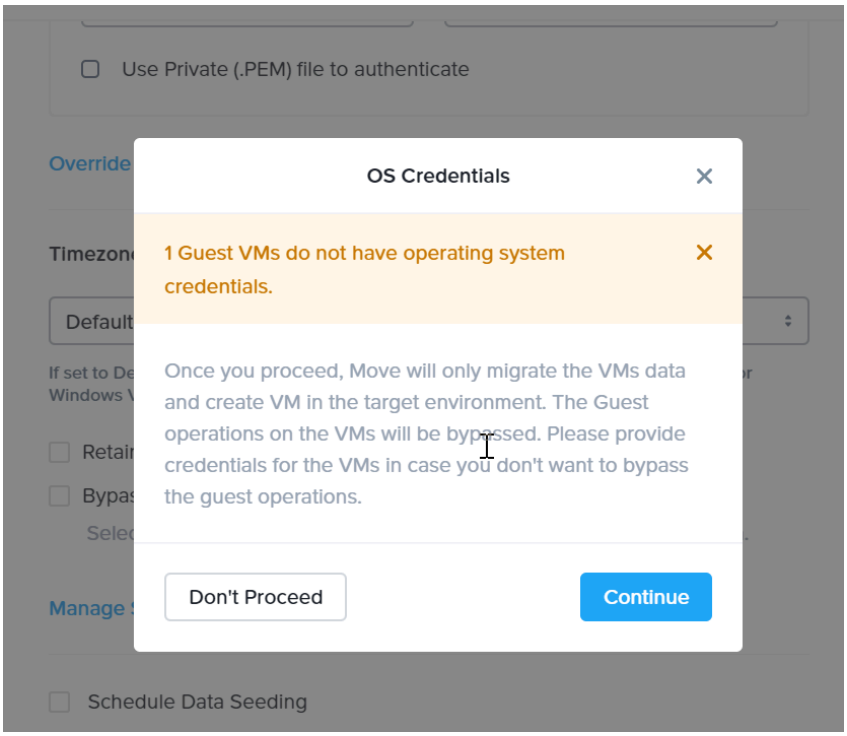

Figure 7: OS Credentials Dialog Box

2. Click **Continue**.
   Move migrates the VMs data and creates a VM in the target, and bypasses the operating system operations.

Performing a Migration Cutover

When the migration plan is started and the seeding process is complete, you can cut over the selected VMs to the NC2 on AWS and AHV cluster. You can monitor the VM migration progress by clicking the **Status** link.

About this task

Note:

- You can perform this operation only when the VM status is Ready to Cutover.
- Recommended that cutover should be performed within one week of initial data seeding.
- If the initial data seeding finishes in less than 10 minutes, the Move VM continues to wait for 10 minutes to take the incremental snapshot; however, you can trigger the cutover immediately.
- The cutover process increments in absolute numbers.

To perform the migration cutover, do the following:

Procedure

1. In the Move UI, click the **Ready to Cutover** status to display the list of available VMs.
2. To perform a cutover, select the VMs or group of VMs.
3. Click **Cutover**.
   The cutover process performs the following VM actions.

   - Shuts down the VM
   - Takes the final snapshots for the VM and copying the final changes to NC2 on AWS and AHV
   - Adds a note in the VM in the vCenter.
   - Disconnects the source VM network interfaces
   - Creates a VM in the target NC2 on AWS and AHV cluster
   - Attaches the replicated disks to the VM
   - Powers on or off the VM (depends on the initial power state)
   - Runs the scripts to set the static IP address

   The cutover process begins immediately and takes a few minutes. Once cutover is complete, the VM is ready for use in the new NC2 on AWS and AHV cluster.

What to do next

You can manage the migration plan once the migration plan is ready. For more information, refer to Environments and Migration Plan Management on page 259

Performance Matrix for Large Data Migration

Move performs end-to-end migration of large VMs. The scenarios are tested based on the following parameters. The following tables show the performance numbers from the Move lab.

Table 6: Performance Numbers of Large Data Migration (ESXi to AHV)

Total migration size	Number of VMs	Size for each vDisks	Number of vDisks	Network bandwidth	Migration time taken		Throughpı	Platform
					Data seeding	Cutover		
3.5 TB	1	3.5 TB	1	10 G	285 minutes	Less than 5 minutes	193 MBps	NX-5155-G6
1 TB	1	1 TB	1	10 G	98 minutes	Less than 5 minutes	180 MBps	NX-5155-G6
1 TB	1	256 GB	4	10 G	80 minutes	Less than 5 minutes	218 MBps	NX-5155-G6
280 GB	8	35 GB	8	10 G	20 minutes	Less than 5 minutes	240 MBps	NX-5155-G6

Prism Central Migration to AHV

Migration of single and scaled out Prism Central validated from ESXi to AHV.

With Prism Central VM migration, IP address must be same before and after the migration for proper connectivity between the Prism Central and the Prism Element.

If the source Prism Central VM have DHCP address configured, Move cannot ensure that the IP address will be retained. It is recommended that the Prism Central VM has static IP configured. Also, disable cloud-init in the source Prism Central VM(s) to keep the static network configuration on target after migration.

While creating the migration plan, select the **Retain MAC address** option. Also if you choose automatic preparation, select **Retain static IP address** from the source VMs option.

For manual preparation, run the preparation script provided by Move in the source Prism Central VM.

Note:

- The migration is qualified with the PC VM without having any other Prism Central features like Flows, Projects, etc.
- Move does not support the migration of PC VMs which are CMSP enabled.

ESXI TO ESXI

You can prepare and migrate ESXi VMs to ESXi on Nutanix environment by using Move.

Move supports the following:

- Migration of vDisk partitions encrypted through cryptsetup in the Linux Unified Key Setup-on-disk-format (LUKS) format.
- Migration of vTPM-enabled VM from ESXi to ESXi.

  After the migration, the configuration of the vTPM-enabled VM at the source is retained at the target. However, the data that was stored in the vTPM VM at the source is not retained. For example, if the vTPM is enabled at the source, it will remain enabled at the target after migration.

  The following configurations are also preserved in the VM after migration:

  - Virtualization-based security (VBS)
  - Input/output memory management unit (I/O MMU)
  - Hardware virtualization

  Note:

  - ESXi as a source refers to ESXi running on any environment.
  - For vTPM-enabled VM migration, if there are multiple key providers on the target, then the default key provider is used for selecting the encryption key.

Migration Considerations

You must consider the supported guest operating systems, requirements, recommendations, unsupported features, and limitations provided in this section before starting the migration process.

Supported Guest Operating Systems (ESXi to ESXi)

Move provides full migration and data-only migration support for the following guest operating systems for ESXi to ESXi migration.

Note: Guest operating systems outside this list are not supported.

Fully Supported

- Windows 7, 8, 8.1, 10, 11
- Windows Server 2008 R2, 2012, 2012 R2, 2016, 2019, 2022
- RHEL 6.3 (32-bit and 64-bit supported) to 6.10, 7.0–8.2
- CentOS 6.3 (32-bit and 64-bit supported) to 6.9, 7.0–7.7, 8.0-8.2
- Ubuntu Server and Desktop 12.04.5, 14.04.x, 16.04.x, 16.10 (32-bit and 64-bit supported)
- Ubuntu Server 12.0.4, 18.04, 19.04, 20.04
- FreeBSD 9.3 and 11.0
- SUSE 11 SP3 / SP4, SUSE 12, SUSE 12 SP1 / SP2 / SP3 / SP4, SUSE 15

- Oracle Linux 6.4 and later, 7.x
- Debian 9.4

Data-Only Support

- Windows with UAC enabled
- RHEL 5.11 with SATA disk controller (32-bit and 64-bit supported)
- CentOS 5.11 (32-bit and 64-bit supported)

Supported Operating Systems for UEFI Enabled VMs (ESXi to ESXi)

Move supports the following operating systems for UEFI enabled VMs.

Table 7: Supported Operating Systems

Operating systems
Windows 7, 10, 11
Windows Server 2008, 2012, 2016, 2019, 2022
Windows Server 2022 secure boot
RHEL 7, 7.1, 7.5, 7.6, 8
CentOS 7, 7.1, 7.5, 7.6, 8

Requirements (ESXi to ESXi)

Before attempting to migrate VMs running on ESXi using Move, make sure to conform to the requirements listed here.

General Requirements for ESXi to ESXi Migration

Ensure to conform to the following requirements for ESXi to ESXi migration.

- Supported browser: Google Chrome
- Ensure you have PowerShell version 4.0 or later.
- VMware Tools must be installed and up-to-date on the guest VMs for migration.
- Ensure to add ESXi (on Nutanix) as the target AOS environment.
- The VMs hardware version should be 7 or above to support the Changed Block Tracking (CBT) feature.
- Source VMs must support Changed Block Tracking (CBT).

  For more information, refer to *VMware KB 1020128, Changed Block Tracking (CBT) on virtual machines*.
- Disks must be either sparse or flat format and must have a minimum version of 2.
- ESXi version must be minimum 5.1.
- Hosts must not be in maintenance mode.
- vCenter reachable from Move on port TCP 443.

  Note: If vCenter is running on different port, make sure https://`ip-address`:port/sdk is accessible from the Move VM.

- ESXi hosts should be reachable from Move on ports TCP 443 and TCP 902.
- Every VM must have a UUID.
- The configuration file (.vmx) of the VMs to be migrated should be present in the ESXi host.
- The VMs must be compatible with the multiple (more than one) snapshots taken by Move.
- Allow ports (TCP and UDP) 2049 and 111 between the Move network and the AHV CVM network.
- Accounts used for performing in-guest operations require **Login as Batch Job** rights in the local security policy on Windows or within the group policy. Administrator users do not have sufficient rights.

  This requirement is only applicable if the VM preparation mode is automatic.
- Ensure that you are an administrator for Windows source VMs or a root for Linux source VMs to run the source VM preparation scripts.
- If local built-in administrator user performs guest preparation, admin approval mode should be disabled. By default, admin approval mode is in disabled state. If the admin approval mode is in enabled state, refer to KB 7672 in the Nutanix Support Portal.
- Ensure that the Move user must belong in a group with **Restore files and directories** security policy.
- Network Security Services (NSS) 3.44 is required for Linux VMs.
- Before you initiate a migration, ensure to disable backups.

Prerequisites for Linux guest VMs:

- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

Service accounts

Move requires the following service accounts with admin privileges.

- vCenter Server
- Prism Element UI for the ESXi on Nutanix cluster

Privileges Required on Source vCenter User for VM Migration

Following are the privileges that you must grant to the source vCenter user for VM migration. For more information about how to grant privileges, refer to the following topics:

- Granting Privileges to a User for a vCenter Server Inventory on page 35
- Granting Privileges to a User for Specific vCenter Server Inventory Objects on page 35

> **Note:** Update the privileges on the user level, not on the group level.

Cryptographic Operations

- Direct Access

Global

- Disable methods
- Enable methods

Sessions

- Validate session
- View and stop sessions

Virtual machine

- Change Configuration
  - Add existing disk
  - Advanced configuration
  - Change Settings
  - Configure Raw device
  - Modify device settings
  - Remove disk
  - Set annotation
  - Toggle disk change tracking
- Guest operations
  - Guest operation modifications
  - Guest operation program execution
  - Guest operation queries
- Interaction
  - Connect devices
  - Power off
  - Power on
- Provisioning
  - Allow read-only disk access
  - Allow virtual machine download
- Snapshot management
  - Create snapshot
  - Remove snapshot

Privileges Required on Target vCenter USer for ESXi on Nutanix

Following are the privileges that you must grant to the target vCenter user for VM migration. For more information about how to grant privileges, refer to the following topics:

- Granting Privileges to a User for a vCenter Server Inventory on page 35
- Granting Privileges to a User for Specific vCenter Server Inventory Objects on page 35

> **Note:** Update the privileges on the user level, not on the group level.

Cryptographic Operations

- Clone
- Encrypt
- Encrypt new
- Migrate
- Register VM

Global

- Disable methods
- Enable methods

Sessions

- Validate session
- View and stop sessions

Virtual machine

- Change Configuration
  - Add or remove device
  - Advanced configuration
  - Modify device settings
  - Memory
  - Change resource
  - Settings

Granting Privileges to a User for a vCenter Server Inventory

You need to grant privileges to the source and target vCenter users for successful VM migration using Move. This topic details the procedure to grant the necessary privileges to a user for a vCenter Server Inventory.

Before you begin

Ensure that you have added a vCenter single sign-on (SSO) user.

For information on adding a vCenter SSO user, refer to the *Add vCenter Single Sign-On Users* section in *VMware vSphere product documentation*.

About this task

To grant the privileges (for VM migration) to a user for a vCenter Server inventory, do the following:

Procedure

1. Login to the vCenter Server using the vSphere Client.

2. Create a role that has all the privileges required for VM migration using Move.

> **Note:**
>
> - For information on creating a role in vCenter, refer to *Create a vCenter Server Custom Role* section in *VMware vSphere product documentation*.
> - To know about the privileges required for a vCenter user for VM migration in Move, refer to Requirements (ESXi to ESXi) on page 55 or Requirements on page 31.

3. From the menu bar, select **Menu** > **Hosts and Clusters**.
4. In the vSphere Client object navigator, select the vCenter Server object for which you want to modify the permissions.
5. Go to the **Permissions** tab and click the **Plus (+)** icon.
   Add Permission window appears.
6. In the **User** field, select the user to whom you want to grant the privileges for VM migration.
7. From the **Role** dropdown menu, select the role that you created in 2 on page 59.
8. Select the **Propagate to children** checkbox and click **OK**.
   The user is assigned the selected role for the selected vCenter Server object and the data centers under it. All the privileges associated with that role are granted to the user.

Granting Privileges to a User for Specific vCenter Server Inventory Objects

You need to grant privileges to the source and target vCenter users for successful VM migration using Move. This topic details the procedure to grant the necessary privileges to a user for specific vCenter Server inventory objects.

Before you begin

Ensure that you have added a vCenter single sign-on (SSO) user.

For information on adding a vCenter SSO user, refer to the *Add vCenter Single Sign-On Users* section in *VMware vSphere product documentation*.

About this task

To grant the privileges (for VM migration) to a user for specific vCenter Server inventory objects, do the following:

Procedure

1. Login to the vCenter Server using the vSphere Client.
2. Create a role that has only the following privileges.
   - Global: Disable methods
   - Global: Enable methods
   - Sessions: Validate session
   - Sessions: View and stop sessions

> **Note:** For information on creating a role in vCenter, refer to *Create a vCenter Server Custom Role* section in *VMware vSphere product documentation*.

3. Create another role that has all the privileges required for VM migration using Move.

   Note: To know about the privileges required for a vCenter user for VM migration in Move, refer to Requirements (ESXi to ESXi) on page 55 or Requirements on page 31.

4. From the menu bar, select **Menu** > **Hosts and Clusters**.

5. In the vSphere Client object navigator, select the vCenter Server object for which you want to modify the permissions.

6. Go to the **Permissions** tab and click the **Plus (+)** icon.
   Add Permission window appears.

7. In the **User** field, select the user to whom you want to grant the privileges for VM migration.

8. From the **Role** dropdown menu, select the role that you created in 2 on page 59.

9. Do not select the **Propagate to children** checkbox and click **OK**.
   The user is assigned the selected role for the selected vCenter Server object only. The privileges associated with that role are granted to the user.

10. Under the selected vCenter Server object in the vSphere Client object navigator, select the data center object for which you want to modify the permissions.

11. In the **Permissions** tab, click the **Plus (+)** icon.
    Add Permission window appears.

12. Select the same user that was selected in 7 on page 60.

13. From the **Role** dropdown menu, select the role that you created in 3 on page 60.

14. Select the **Propagate to children** checkbox and click **OK**.
    The user is assigned the selected role for the selected data center object only. The privileges associated with that role are granted to the user.

What to do next

Under the same vCenter Server, if you have multiple data centers for which you want to modify the permissions, then perform 10 on page 60 to 14 on page 60 for each of those data centers.

Recommendations (ESXi to ESXi)

Nutanix recommends the following for optimal VM migration from ESXi.

Recommendations

- Clear all the VM alerts in vCenter, if any.
- Refresh the inventory.
- Convert the templates to VMs.
- Enable access to VMs from vCenter.
- Ensure that all VMs are connected through vCenter.
- Ensure that all VMs are valid in vCenter.
- Recover VMs fully, if any are orphaned.
- Ensure the disk compatibility with *Changed Block Tracking*(CBT).

- Disable fault tolerance for VMs, if any in a fault tolerance pair.
- Install the latest version of VMware Tools on the VMs to be migrated.
- Ensure that the CBT-enabled VMs have fewer than 30 snapshots in the inventory.

> **Note:** Move migrates maximum of 8 disks from single ESXi hosts in parallel. The other VMs for migration from the same ESXi hosts are queued and only progress as and when the earlier disk data seeding completes. The limit is 16 disks in parallel at the appliance level. Refer to KB-9460 for further information.

Unsupported Features (ESXi to ESXi)

This section lists the unsupported features for migration from ESXi to ESXi.

- PCIE pass-through
- Independent disks
- VMs with multi-writer disks attached
- VMs with 2 GB sparse disk attached
- VMs with SCSI controllers with a SCSI bus sharing attached

> **Note:** Change SCSI bus controller to `None`.

- Migration of Windows VMs with dynamic disks
- Migration of VMs from ESXi standalone hosts with free license

Limitations (ESXi to ESXi)

This section lists the limitations for migration from ESXi to ESXi.

ESXi to ESXi Migration Limitations

In addition to Migration Limitations on page 12, the following are the limitations while performing migration from ESXi to ESXi.

- For standalone ESXi migrations, you should not use *vMotion* or *Storage vMotion* on the VMs under migration.
- Containers that are mounted on all the ESXi hosts are available for selection as the container for the ESXi target.
- The target ESXi must be registered with vCenter while adding it as a target and must be registered for the entire duration of its existence as an entity with Move.
- The virtual hardware version of VMs defaults to the ESXi target version and VMs fails to start after migration.

  Workaround: Reinstall the VMware Tools in the target.
- VMs with UEFI boot configuration do not boot after migration for AOS versions less than 5.19.

  Workaround: Manually update the VM configuration. Refer to the VMware KB article *Enable or Disable UEFI Secure Boot for a Virtual Machine*.
- VMware tools are not installed at the target VM after migration from higher version of ESXi to lower version of ESXi.

  Workaround: Reinstall the VMware tools in the target.
- After migration, Windows VMs with connected NICs only will retain their IP address. Those with disconnected NICs will not retain their IP address.

Adding vCenter Server or Standalone ESXi Host Environment

While creating a migration plan for migration from ESXi to any target, ensure to add at least one vCenter Server or standalone ESXi host environment for migration.

About this task

> **Note:** This procedure is only applicable for migration from ESXi.

To add a vCenter Server or standalone ESXi host environment, do the following:

Procedure

1. Log on to the Move UI.
2. Click **+ Add Environment** under **Environments**.
   The **Add Source Environment** window appears.

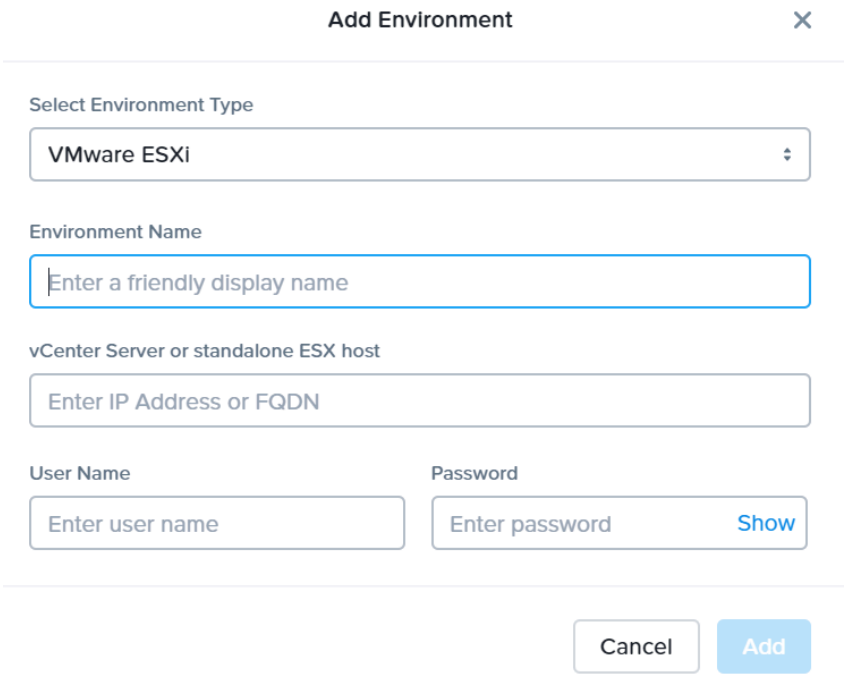

Figure 8: Add VMware ESXi Environment Dialog Box

3. Select **VMware ESXi** as the environment type.

4. Complete the indicated fields and click **Add**.

   a. **Environment Name**: Enter a name for the ESXi environment.

   b. **vCenter Server or standalone ESX host**: Enter the IP address or the FQDN of the vCenter Server, or the IP address of the ESXi host.

      If you do not have a standard port, use the custom port for vCenter in the `vCenter IP address:custom port number` format. For example, if you are using IP address, use the format `10.136.72.150:8443` or if you are using FQDN, use the format `vcenter.nutanix.com:8443`.

   c. **User Name**: Enter the username or User Principal Name (UPN) for logging on to the vCenter Server.

   d. **Password**: Enter the password for logging on to vCenter Server.

   The VMware ESXi environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

You can now add Nutanix AHV cluster environment. For more information, refer to Adding a Nutanix AOS Cluster Environment on page 39

Adding a Nutanix AOS Cluster Environment

While creating a migration plan, AHV AOS cluster can be added as both source or target. If you want to use ESXi on Nutanix cluster as target, then add the corresponding AOS cluster environment for ESXi.

About this task

> **Note:** When you add a AOS cluster, Move VM IP address with the subnet 255.255.255.255 is added to the global NFS allowlist.
>
> > **Caution:** Modifying the NFS allowlist of the destination container disables inheritance of the NFS allowlist at the global level and lead to issues with Move operations.

To add a Nutanix AOS cluster environment, do the following:

Procedure

1. Log on to the Move UI.

2. Click **+ Add Environment** under **Environments**.

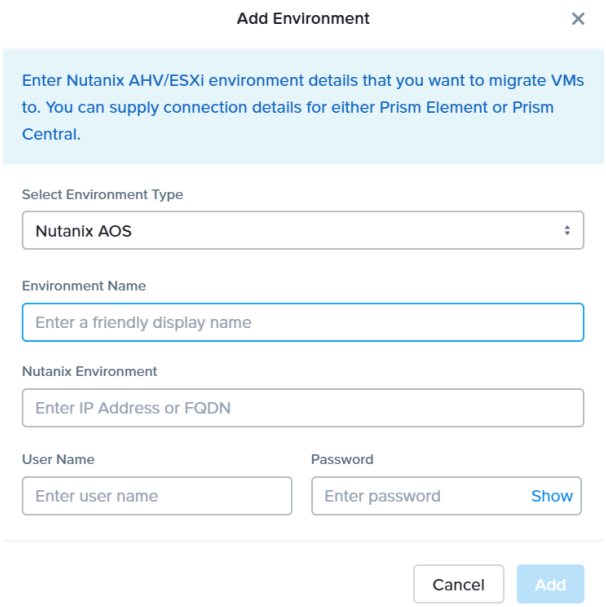

Figure 9: Add AOS Environment Dialog Box

The **Add Environment** window appears.

3. Select **Nutanix AOS** as the target environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the NC2 on AWS and AHV or VMware ESXi on Nutanix environment.
   b. **Nutanix Environment**: Enter the IP address or the FQDN for the target Prism Central or Prism Element.
   c. **User Name**: Enter the username for logging on to the target Nutanix environment.
   d. **Password**: Enter the password for logging on to the target Nutanix environment.

5. (Only for ESXi to ESXi migration) Enter credentials for registered vCenter.

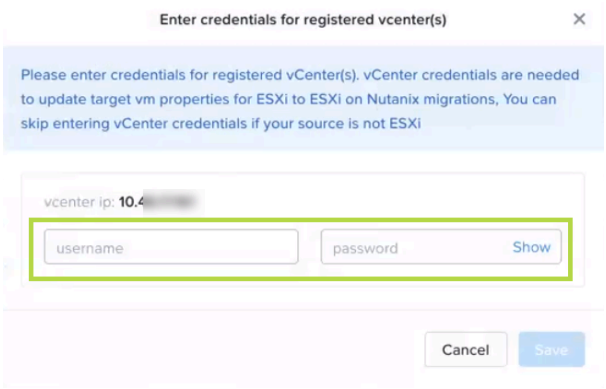

Figure 10: vCenter Credentials Dialog box

vCenter credentials are required to update the target VM properties for ESXi to ESXi on Nutanix migration.

Note:

- You can skip adding vCenter credentials if your source is not ESXi.
- To retain the VM properties on the target VM after migration, be sure to provide the target vCenter credentials. The following properties will be retained:
  - SCSI controller types
  - Network adaptor type
  - MAC address
  - Video card
  - Memory overcommit variables
  - Tool upgrade flag
  - Sync time with host flag
  - Disable acceleration flag
  - Enable logging flag

The AOS environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

Once you have added the environments, you can proceed to create the migration plan. For more information, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan (Hyper-V to ESXi) on page 103 or Creating a Migration Plan (Azure to AHV) on page 171 or Creating a Migration Plan (Azure to ESXi) on page 184

Creating a Migration Plan (ESXi to ESXi)

You can create a migration plan to seed the data, cutover, and monitor the VMs. You can create the migration plan in Move UI without initiating the cutover process.

About this task

Note:

- This procedure is only applicable for migration from ESXi to ESXi.
- If you are logging in for the first time, log on to the Move UI with your default credentials.
- You must have admin user credentials to complete the migration process.
- If you restart the management server, scheduled VM migration does not begin automatically.
- If the source boot is set to UEFI, set up the boot device manually in the VM post migration for the following operating systems.
  - CentOS 7.4 2, 6.8 3, 8, 8.1, 8.2
  - Ubuntu 12.04 4, 19.04
  - OEL 7 5
  - RHEL 6.8, 8.1

  For more information about setting up the boot device, refer to AHV Administration Guide.
- When Move encounters a large source VM disk of size greater than 2 TB that belongs to one ESXi host, Move prioritize the migration of this disk and does not migrate any other disks from the same host in parallel. Only after the large VM disk migration gets completed, Move migrates the other disks from the same host. Meanwhile, Move migrates the VM disks belongs to other source ESXi hosts in parallel.
- If you cancel or discard the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.

  For information on canceling an ongoing migration, see Pausing or Canceling a VM Migration on page 257.

  For information on performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

To create a migration plan, do the following:

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.
3. On the Move dashboard, click **Create a Migration Plan**.

   Note: If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.

4. Enter the new migration plan name, and then click **OK**.

> **Note:** You can edit the migration plan name by clicking the pencil icon next to the migration plan name.

The **New Migration Plan** window appears.

5. Complete the following fields, and then click **Next**.

a. **Select a Source**: Select any vCenter Server or standalone ESXi host source for migration. Once you select the source, an appropriate target appears.

> **Note:** You might at times see a message relating to inventory collection as shown below. During this time, the environments undergoing Refresh will not be available for selection. Such environments do not automatically show up for selection once the inventory collection is over. In case you wish to select one of the environments undergoing Refresh (not available for selection), you will need to select **Cancel** and wait for inventory collection to get over and then create the migration plan again.

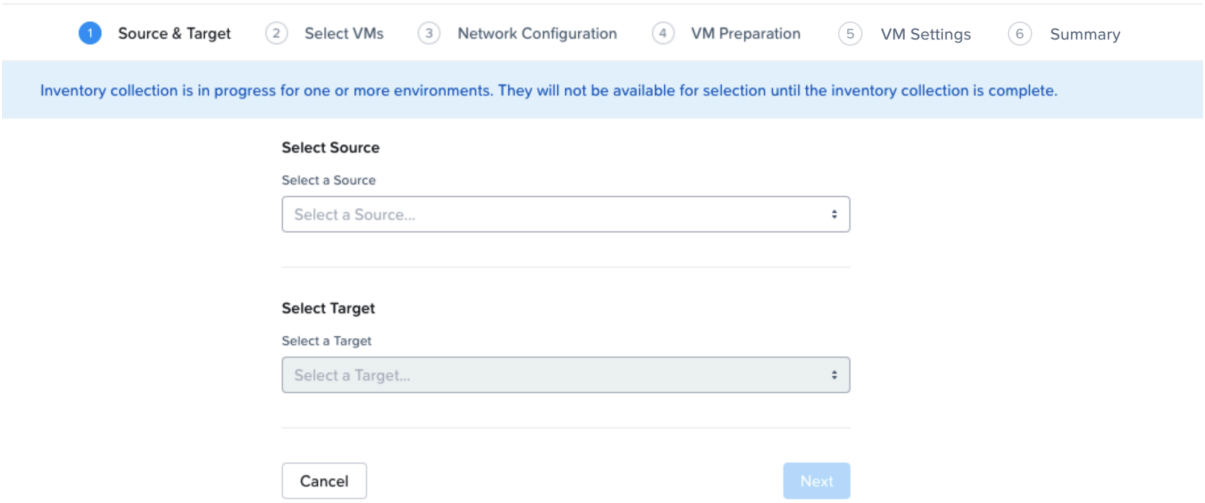

Figure 11: Inventory Collection Message

b. **Select a Target**: Select Prism Central or ESXi Prism Element as target for the VM migration.

c. **Target Project** (optional): Select the project you want as the target.
This field is available only if Prism Central is selected as the target.

d. **Target Cluster**: Select the cluster on which you will migrate the VMs.
This field is available only if Prism Central is selected as the target.

e. **Target Containers**: Select the container on which you will migrate the VMs.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click **+Add All**, and then click **Next**.

> **Note:** You cannot add more than 50 VMs in a single migration plan.

You can filter the VMs by name in the **Filter** bar. You can also sort the VM types by selecting the appropriate column. The VMs for which migration has failed are not displayed. To show the entire list of VMs, select **All VMs** from the drop-down list. To show the configuration of VMs, select **Configuration** from the drop-down

list. A question mark icon appears beside an unavailable VM, which displays more information about that VM and might indicate why the VM cannot be migrated.

Note:

- If the source VM has RDM disks in physical compatibility mode, then those disks are converted to virtual compatibility mode during source VM preparation. By default, power cycle is enabled for the VMs with physical RDM disks. Move performs the following:

  1. Shut down the source VM.
  2. Convert physical RDM disks to virtual compatibility mode.
  3. Start the VM.

  If power cycle is not enabled, then Move converts the physical RDM disks to virtual with VM in powered on state.

- Migrate VMs retain their power state on the destination cluster.

The selected VMs are displayed in the sidebar.

7. In the **Network Configuration** screen, select the target network from the drop-down, and then click **Next**.
   For performing test migration, refer to Creating a Test Capable VM Migration Plan on page 248 section.

8. In the **VM Preparation** screen, select one of the following VM preparation modes.

   » **Automatic**. Move automatically runs scripts on the source VMs to prepare them for migration.

   » **Manual**. Move displays the VM preparation scripts for Windows and Linux VMs.

   These scripts prepare the source VMs by performing the following:

   - Runs the IP address retention script.

   - Runs Set SAN policy script.

   Note: If you have not run the preparation script in the source VM, Move performs data-only migration. For more information, refer to Performing Data-Only Migration on page 50.

   » **Mixed**. Move allows you to select the VM preparation mode for each VM. This setting allows you to manually prepare one set of VMs and automatically prepare another set of VMs in the same migration plan. If you want to have such a hybrid combination of **Automatic** and **Manual** VM preparation modes, proceed to the next step.

   Note: The preparation mode automatically switches to **Mixed** if you change the mode of preparation for a set of VMs under **Change Settings** and have a mix of **Automatic** and **Manual** VM preparation modes. You cannot manually select the **Mixed** option from the **Preparation Mode** drop-down list.

9. (Optional) Under the **Guest Operations** section in the **VM Preparation** screen, do one or more of the settings.

   a. **Retain static IP addresses from source VMs**: Retains static IP addresses from the source VMs to the migrated VMs on AHV. By default, this option is enabled. Clear the checkbox if you do not want to retain

the static IP addresses from the source VMs on the target VMs. If you disable this option, the static network is converted to DHCP network.

For **Manual** preparation mode, if you do not want to do not want to retain the static IP addresses from the source VMs on the target VMs, do the following:

- For Windows VMs, change the argument from `$retainIP = $true` to `$retainIP = $false` in the script.
- For Linux VMs, remove the argument --retain-ip from the script.

b. (Only for Automatic preparation mode) **Bypass Guest Operations on Source VMs**: Select this check box to bypass the guest operating system changes.

You can select this option to override your migration to data-only migration.

> **Note:** If you bypass guest operations on source VMs, **Retain static IP addresses from source VMs** configuration will not be applicable.

c. (Only for Manual preparation mode) **Re-Generate Script**: This option gets enabled only when there is a change in the selection of the above-mentioned settings under the **Guest Operations** section. If enabled, click this option to re-generate the VM preparation scripts for both Windows and Linux VMs.

10. Do one of the following based on whether the preparation mode selected:

» Automatic preparation mode: Provide the credentials of the source VMs under **Windows VMs** or **Linux VMs**, depending on the type of the source VM.

> **Note:**
>
> - For Windows VMs - Move supports only username-password sign-in option for authentication. It does not support username-PIN sign-in option.
> - For Linux VMs - Apart from credentials, Move supports PEM file for authentication. Select the **Use Private (.PEM) file to authenticate** option and upload the private key.
> - If you want to retain the static IP addresses, provide a common set of credentials for your selected Windows or Linux VMs.
> - If you face any issues while using the .PEM file for authentication, refer to KB 7090.
> - Currently, the default location where the preparation scripts are stored is the /tmp folder.
>
> If the /tmp folder is mounted as noexec, then Move will fallback to the /var/tmp folder. If the /var/tmp folder is also mounted as noexec, then Move will fallback to the /usr/tmp folder.

» Manual preparation mode: Copy the scripts and manually run them on the respective source VMs, and then click **Next**.

> **Note:**
>
> - For running the script, use the Windows built-in administrator credentials for the Windows VMs and use root user for the Linux VMs.
> - If you have not run the preparation script in the source VMs, Move performs data-only migration.
>
> For more information, refer to Performing Data-Only Migration on page 50.

11. In the **Override individual VM Preparation** section, click **Change Settings** to override the **Guest Operations** settings (configured in the above steps) for the individual VMs. You can also edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

    Note: If you do any of the following for a VM, then copy the new generated scripts of that specific VM and run them on the source VM:

    - Change the **Mode of Preparation** of a VM to **Manual**.
    - Change any of the guest operation settings of a VM with the preparation mode set to **Manual** (an icon appears next to the VM Name prompting to regenerate a new guest script).

    Access the newly generated VM preparation script for that VM by selecting the Copy Scripts icon under the **Actions** column.

    After making the required changes, click **Done**.

12. (Optional) In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

    a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

    b. **Timezone**: Set the timezone as the hardware clock of the VMs in target.

    If set to default, Move configures UTC timezones for Linux VMs and cluster timezones for Windows VMs.

    Note: Timezone is available only if **Prism Central** was selected as the target in the **Source & Target** screen.

    c. **Retain MAC Addresses from the Source VMs**: Select this check box to retain the MAC addresses from the source VMs.

    d. **Skip CDROM addition on target VMs**: Select this check box to skip the CDROM addition on the target VMs.

    e. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

    Only those categories that have values are available for selection.

    Note: Category Settings is available only if **Prism Central** was selected as the target in the **Source & Target** screen.

    f. **VM Migration Type**: Select one of the following VM migration types.

    At the VM level, some of the target VM properties can be customized manually after the migration plan is created. For information on manually customizing the target VM configuration, refer to Customizing the Target VM Configuration on page 255.

    - **Configure Target VM Properties**: The target VM synchronizes with the source VM properties at the time of migration plan creation. Selecting this option allows you to edit the target VM properties at the

VM level (during migration). For information on editing the target VM properties, refer to Customizing the Target VM Configuration on page 255.

> **Note:** At the VM level, only the following properties can be edited:
>
> - Target VM name
> - Number of vCPUs
> - Number of cores per vCPU
> - Memory
> - Power state

- **Retain Source VM Properties**: The target VM synchronizes with the source VM properties whenever Move refreshes the source VM configuration details. Only the customizable properties are refreshed on the target. Selecting this option does not allow you to edit the target VM properties at the VM level.

  > **Note:**
  >
  > - The source VM properties are refreshed in the following ways.
  >   - (Manually) When you click the **Refresh Source VM Properties** button.
  >   - (Automatically) When you start a migration plan.
  >   - (Automatically) When you initiate a cutover.
  > - When you start a migration plan, Move refreshes both source VM and target VM properties by default. However, it will not refresh the target VM properties at the start of a migration plan if you modified the target VM properties after migration plan creation.

g. **Settings for individual VMs**: Click **Change Settings** to configure settings such as timezone, retain MAC addresses, VM priority, and skip CDROM addition for individual VMs. You can also search the VM by typing the name of the VM and change the settings.

h. **Schedule Data Seeding**: Check this check box to select the date and time for migration.

> **Note:** For migration of VMs with multiple NICs, static IP address configurations are applied correctly if the MAC address retention is applied from source VMs, otherwise best effort IP address configuration is done by mapping one NIC at a time.

This action does not affect the VM data migration but requires you to manually prepare the guest operating system with the necessary AHV drivers prior to cutover. In addition, if you bypass the guest operations, you have to take care of the static IP address retention separately.

13. In the **Summary** screen, choose one of the following, and then proceed review the VM migration summary.

    » **Back**: Click this option to edit the information.

    » **Save**: Click this option to save the migration plan.

    For more information about how to start the migration later, and check the migrated VM status and details, refer to Environments and Migration Plan Management on page 259.

    » **Save and Start**: Click this option to save the migration plan and begin the migration immediately

    The seeding process for migration begins. You can monitor this information by selecting **Status** for the migration plan.

    > **Note:** The seeding process can take several minutes depending on the number of VMs.

What to do next

- To customize the target VM configuration at the VM level, refer to Customizing the Target VM Configuration on page 255.
- If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover (ESXi to ESXi) on page 72.

Performing a Migration Cutover (ESXi to ESXi)

When the migration plan is started and the seeding process is complete, you can cut over the selected VMs. You can monitor the VM migration progress by clicking the **Status** link.

About this task

> **Note:**
>
> - You can perform this operation only when the VM status is Ready to Cutover.
> - Recommended that cutover should be performed within one week of initial data seeding.
> - If the initial data seeding finishes in less than 10 minutes, Move continues to wait for 10 minutes to take the incremental snapshot; however, you can trigger the cutover immediately.
> - The cutover process increments in absolute numbers.

Procedure

1. In the Move UI, click the **Ready to Cutover** status to display the list of available VMs.
2. To perform cutover, select the VMs or group of VMs.

3. Click **Cutover**.

   The cutover process performs the following VM actions.

   - Shuts down the VM
   - Takes the final snapshots for the VM and copying the final changes to ESXi on Nutanix cluster
   - Adds a note in the VM in the vCenter
   - Disconnects the source VM network interfaces
   - Creates a VM in the target ESXi on Nutanix cluster
   - Attaches the replicated disks to the VM
   - Powers on or off the VM (depends on the initial power state)
   - Runs the scripts to set the static IP address

   The cutover process begins immediately and takes a few minutes. Once cutover is complete, the VM is ready for use in the target.

What to do next

You can manage the migration plan once the migration plan is ready. For more information, refer to Environments and Migration Plan Management on page 259

Performance Matrix for Large Data Migration (ESXi to ESXi)

Move performs end-to-end migration of large VMs. The scenarios are tested based on the following parameters. The following tables show the performance numbers from the Move lab.

Table 8: Performance Numbers of Large Data Migration (ESXi to ESXi)

Total Migration size	Number of VMs	Size for each Vdisks	Number of Vdisks	I/O on the source	Data churn	Network bandwidth	Migration time taken	
							Data seeding	Cutover
2 TB	1	2 TB	1	No	No data churn.	10 G	190 minutes	Less than 5 minutes
1 TB	1	1 TB	1	Yes	10 GB	1 G	418 minutes	28 minutes

HYPER-V TO AHV AND HYPER-V TO NUTANIX CLOUD CLUSTERS (NC2) ON AWS

You can prepare and migrate VMs running on Hyper-V hypervisor to Nutanix Cloud Clusters (NC2) on AWS and Hyper-V hypervisor to AHV by using Move.

Note:

- For migration from any source (ESXi, Hyper-V, and AWS) to AHV target and from any source (ESXi, Hyper-V, and AWS) to NC2 on AWS target, Move should be deployed on the same destination target cluster where the VMs need to be migrated.
- For NC2 on AWS, static IP retention is not enabled by default.

Migration Considerations

You must consider the supported guest operating systems, requirements, recommendations, unsupported features, and limitations provided in this section before starting the migration process.

Supported Guest Operating Systems

Move supports some common operating systems from Hyper-V to AHV.

Note: The Move UI does not display a message or warning during the creation of a migration plan if you are attempting to migrate VMs running on an unsupported guest operating system.

For more information about the supported guest operating systems on AHV, refer to Compatibility and Interoperability Matrix. It also indicates whether an operating system is community-supported, legacy, or deprecated on AHV.

Table 9: Supported Guest Operating Systems

Operating systems	Gen 1 Support	Gen 2 Support
Windows Server 2022	Yes	Yes **
Windows Server 2019	Yes	Yes**
Windows Server 2016	Yes	Yes **
Windows Server 2012 R2	Yes	Yes **
Windows Server 2012	Yes	Yes **
Windows Server 2008 SP2 (32 bit)	NA - target	NA - source
Windows Server 2008 SP2 (64 bit)	NA - target	NA - source
Windows Server 2008 R2 SP1	Yes	NA - source
Windows Server 2003 SP2 (32 bit)	NA - target	NA - source

Operating systems	Gen 1 Support	Gen 2 Support
Windows 7 (32 bit)	Yes	NA - source
Windows 7 (64 bit)	Yes	NA - source
Windows 8 (32 bit)	Yes	NA - source
Windows 8 (64 bit)	Yes	Yes **
Windows 10 (32 bit)	Yes	NA - source
Windows 10 (64 bit)	Yes	Yes **
RHEL 6.5–6.9, 7.0–7.5 (64 bit)	Yes	NA -Target
RHEL 7.6–7.7	Yes	Yes **
RHEL 7.8–8.0	Yes	Yes
RHEL 8.1–8.4	Yes	Not Certified
RHEL 9.1, 9.2	Yes	Yes
CentOS 6.5–6.9, 7.0–7.5 (64 bit), 8.1-8.2	Yes	NA -Target
CentOS 7.3	Yes	Yes **
CentOS 8.0, 8.4	Yes	Yes
Ubuntu 14, 16, 18, 19.04, 20.04	Yes	Yes #
FreeBSD 11	Yes	NA - source
Debian 9.4	Yes	Not Certified
Oracle Linux 7.x	Yes	Not Certified
Oracle Linux 7.8, 9.1	Yes	Yes **
SUSE 12 SP3	Not Certified	Yes
SUSE 12 SP5	Yes	Yes
SUSE 15 SP1	Yes	Yes **

Table 10: Legend

1. **Yes**. Guest operating system is supported for migration.
2. **NA - source**. Guest operating system is not supported on Microsoft Hyper-V for given generation.
3. **NA - target**. Guest operating system is not supported on AHV.
4. **Yes ****. Generation 2 guest operating system is supported for migration. For Generation 2 VM, during cutover process, Secure Boot feature is disabled for VM on the added Hyper-V source until cutover is completed. Once cutover is completed, the setting is reverted on the added Hyper-V source.

   Current available versions of AHV do not support Generation 2 VMs directly. Therefore, during the cutover process, Move automatically runs the following command on a migrated VM on the target cluster.

   ```
   acli vm.update <vm_id> uefi_boot=True
   ```

   This command passes the target information, IP address, and credentials to Move. This command might fail due to the following reasons.

   - The added target is Prism Central.
   - Unable to SSH to Prism Element with the credentials provided while adding Prism Element as target.

   Perform the following manual steps on the target cluster if the migration of a Generation 2 VM fails with the following error message: Failed to update UEFI flag for <VmName>

   1. Shut down the VM.
   2. Run the following aCLI command.

      ```
      acli vm.update <vm_id> uefi_boot=True
      ```
   3. Start the VM.

   After you perform the preceding steps and verify the VM on the target, discard the migration in Move for that VM.
5. **Yes #** . In addition to the VM preparation, run the following commands on the source Ubuntu 14 and 16 Generation 2 VMs before migration.

   1. Log on to source VM.
   2. Change to bash shell.

      ```
      sudo bash
      ```
   3. Change the directory.

      ```
      cd /boot/efi/EFI
      ```
   4. Copy the files to boot directory.

      ```
      cp -r ubuntu/ boot
      ```
   5. Change the directory to boot.

      ```
      cd boot
      ```
   6. Rename the file.

      ```
      mv shimx64.efi bootx64.efi
      ```

Support for UEFI with Secure Boot Enabled VMs

Move supports UEFI with secure boot enabled VMs.

Table 11: Supported Guest Operating Systems

Operating systems
Windows Server 2016, 2019, 2022
Windows Desktop 10
RHEL 7.7, 9.1, 9.2
CentOS 7.3, 8.0, 8.4
Ubuntu 22.04
Oracle Linux 9.1

> **Note:** Non-LVM (Logical Volume Management) home partition is lost with the migration of RHEL 9.x and OEL 9.x.
>
> For workaround, refer to KB 13157.

Requirements

Before attempting to migrate VMs running on Hyper-V using Move, make sure to conform to the requirements listed here.

General Requirements

Make sure to conform to the following requirements for Hyper-V to NC2 on AWS and AHV migration:

- Supported browser: Google Chrome.
- Ensure you have PowerShell version 4.0 or later.
- Ensure guest VMs have connectivity with Move.
- Ensure that the guest VMs have integration services installed and an IP address is present in the **Networking** section of the Hyper-V manager for automatic preparation of the source VMs.

  For more information, refer to *Manage Hyper-V Integration Services* in Microsoft documentation for Hyper-V.
- Ensure the following:
  - WinRM is configured on Hyper-V servers and Windows guest VM.
  - The following inbound and outbound ports using the TCP protocol are enabled for the Windows Remote Management (WinRM) feature to work.
    - WinRM-HTTPS: 5986
    - WinRM-HTTP: 5985

    > **Note:** This requirement is applicable only for automatic VM preparation mode.
- If the local built-in administrator performs guest preparation, then admin approval mode must be disabled. By default, admin approval mode is in disabled state. If the admin approval mode is in enabled state, refer to KB 7672 in the Nutanix Support Portal.
- The source has sufficient space for the snapshot disks created during migration.
- Network Security Services (NSS) 3.44 is required for Linux VMs.
- Before you initiate a migration, ensure to disable backups.

> **Note:** If the Microsoft Hyper-V Server version is 2016 or higher, Nutanix recommends configuring the checkpoints as production checkpoints for better migration performance. For information about the procedure to change checkpoints to production checkpoints, refer to Microsoft documentation.

Prerequisites for Linux guest VMs:

- SSH service must be up and running.
- The credentials provided must have root or sudo user permission.
- Guest VM must have curl utility installed.

Service Accounts

For successful migration of VMs from Hyper-V to AHV and Hyper-V to NC2 on AWS, service accounts for the following must have admin privileges.

- Source Hyper-V Server (standalone or cluster)
- Prism Element UI for the AHV cluster
- Source VMs planned for migration

  A user with administrator role for Windows source VMs or a root user for Linux source VMs to run the source VM preparation scripts.

Steps to Migrate VMs with Unordered Disks

To specify the boot disk for VMs with unordered boot disk information, perform the following steps:

1. SSH into the Move VM with the admin credentials. For more information, refer to Accessing Move VM with SSH on page 21.
2. Switch to the root user by entering the password for the admin user.

   ```
   admin@move on ~ $ rs
   [sudo] password for admin:
   ```

3. Browse to the directory /opt/xtract-vm/conf.

   ```
   cd /opt/xtract-vm/conf
   ```

4. Create or open the file tgtagent.json.

   ```
   vi tgtagent.json
   ```

5. Create the JSON content.

   Following is an example of tgtagent.json. You can use this example JSON file, and change or remove the values of the flags as necessary.

   ```
   {
     "AHVTargetBootConfig": {
       "ConvertIDEtoSCSI": false
     }
   }
   ```

   > **Note:** By default, Move converts IDE disks to SCSI. If you want to retain IDE disks and not convert them to SCSI, then set ConvertIDEtoSCSI to `false`.

6. Restart the tgtagent container.

   ```
   restart-tgtagent
   ```

   > **Note:** You can create this JSON file while the migration is in progress or before initiating the migration.

Limitations

This section lists the limitations for migration from Hyper-V to AHV and Hyper-V to NC2 on AWS.

In addition to Migration Limitations on page 12, the support is constrained by the following.

- Migration of guest VMs might fail if you delete an existing user snapshot or checkpoint during the migration.
- VM names with single and double quotes are not supported for migration.
- If VMs are configured with dynamic memory, Move takes only the start-up memory configuration while creating the target VM on AHV.
- Automatic installation of the Move Hyper-V agent is not supported for Windows Server 2012 (standalone & cluster).
- VMs protected by Hyper-V replica are not supported.

  To allow migration, disable the Hyper-V replication for the required VMs.
- For Hyper-V to AHV migrations, if VMs have non-aligned disks, then migrated VMs might not power on.

  Use Convert-VHD command to create VHD, and set the alignment before starting migration for such disks.

  For more information, refer to Nutanix KB article 10980.
- VMs migrated with more than one network interfaces might not retain all the IP addresses.

  Workaround: Manually assign IP addresses after migration.
- For IP addresses retention to work for VMs having both static NIC and virtual switches, you need to select the **Retain MAC** check box. Otherwise, IP addresses may not be retained on the migrated VMs.
- After migration, Windows VMs with connected NICs only will retain their IP address. Those with disconnected NICs will not retain their IP address.
- (For Hyper-V to NC2) Retention of static IP address is not supported for VM migrations.

Adding Hyper-V Environment

While creating a migration plan for migration from Hyper-V to any target, be sure to add at least one Hyper-V environment for migration.

About this task

> **Note:** This procedure is only applicable for migration from Hyper-V.

To add a Hyper-V environment, do the following:

Procedure

1. Log on to Move UI.

2. Click **+ Add Environment** under Environments.
   The **Add Environment** appears.

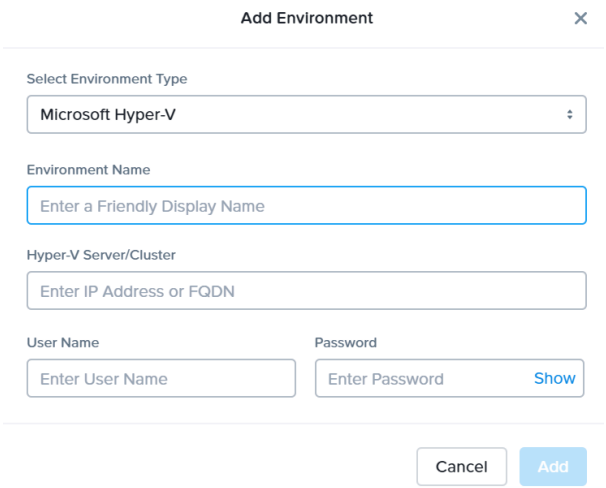

Figure 12: Add Hyper-V Environment Dialog Box

3. Select **Microsoft Hyper-V** as the source environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the Hyper-V environment.
   b. **Hyper-V Server**: Enter the IP address or FQDN of the Hyper-V source server.
   c. **Username**: Enter a Windows account with administrator privileges of the Hyper-V source server.
   d. **Password**: Enter the password for this Windows account.

   Note:

   - Hyper-V user account must have administrator privileges.
   - Enter cluster IP address (or failover IP address) to discover all clustered VMs from Hyper-V cluster environment.
   - For cluster environment, enter the domain credentials the administrator privileges.
   - You must either add the cluster IP address or the standalone IP address of the node.

   The Hyper-V environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

You can also add a Nutanix AHV environment. For more information, refer to Adding a Nutanix AOS Cluster Environment on page 80

Adding a Nutanix AOS Cluster Environment

While creating a migration plan, if you need to add AHV as the source or target, or ESXi on Nutanix as a target, then you have to add at least one AOS cluster environment for migration.

About this task

> **Note:** When you add a AOS cluster, Move VM IP address with the subnet 255.255.255.255 is added to the global NFS allowlist.
>
> **Caution:** Modifying the NFS allowlist of the destination container disables inheritance of the NFS allowlist at the global level and lead to issues with Move operations.

To add a Nutanix AOS cluster environment, do the following:

Procedure

1. Log on to the Move UI.
2. Click **+ Add Environment** under **Environments**.

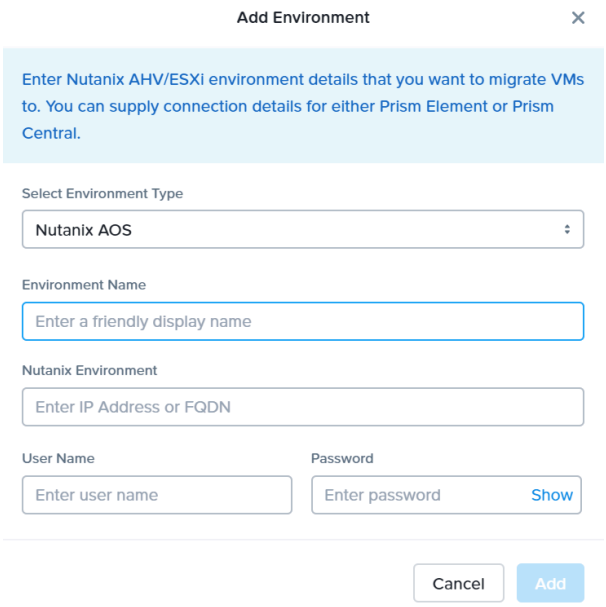

Figure 13: Add AOS Environment Dialog Box

The **Add Environment** window appears.

3. Select **Nutanix AOS** as the target environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the NC2 on AWS and AHV or VMware ESXi on Nutanix environment.
   b. **Nutanix Environment**: Enter the IP address or the FQDN for the target Prism Central or Prism Element.
   c. **User Name**: Enter the username for logging on to the target Nutanix environment.
   d. **Password**: Enter the password for logging on to the target Nutanix environment.

5. (Only for ESXi to ESXi migration) Enter credentials for registered vCenter.

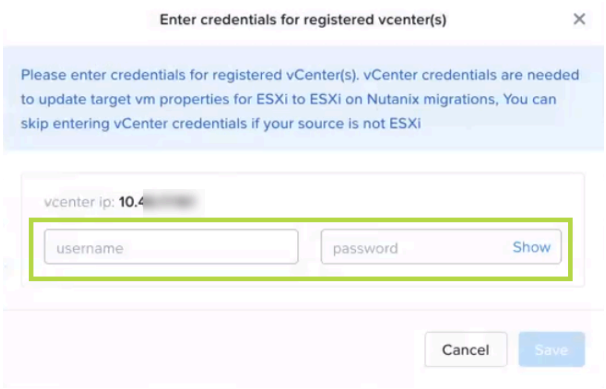

Figure 14: vCenter Credentials Dialog box

vCenter credentials are required to update the target VM properties for ESXi to ESXi on Nutanix migration.

> **Note:** You can skip adding vCenter credentials if your source is not ESXi.

The AOS environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

Once you have added the environments, you can proceed to create the migration plan. For more information, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan (Hyper-V to ESXi) on page 103 or Creating a Migration Plan (Azure to AHV) on page 171 or Creating a Migration Plan (Azure to ESXi) on page 184

Deploying the Move Agent on Hyper-V Host

You can manually or automatically deploy the Move agent on the source Hyper-V host.

Manual Deployment

To download and install the Move agent on each of the source Hyper-V host machine to support VM discovery and migrations, do the following:

1. In the Hyper-V host, download **move-agent-installer.exe** from **http://<nutanix-move-ip>/downloads/agents/move-agent-installer.exe** .

   Replace `<nutanix-move-ip>` with the IP address of the Move VM.
2. Go to the location where you have downloaded the agent, and copy or move the downloaded file under C:\users \Administrator.
3. Launch the command prompt with **Run as Administrator** to run the following command from C:\users \Administrator.

   ```
   move-agent-installer.exe -o [operation] -ip [move ip] -u [user]
   ```

   Example: `move-agent-installer.exe -o install -ip 10.5.244.55 -u user`

   User can be either a domain or local user with administrator privileges.
4. Enter the password when the command prompt requests for it.

Note:

- Do not use **-p** option in the command to include the password if the password contains special characters. The PowerShell console throws an exception when a password containing special characters is included in the command using **-p** option.
- By default, the Move agent is installed in the user directory. If you want to change the location, use the **-d** option.
- To uninstall the Move agent, use the command `move-agent-installer.exe -o remove`.
- Enter help along with any command for information to display options for that command.
- Move agent service installation will add inbound firewall rule to open port 8087, which is required for Move VM and Hyper-V interactions. The Move Hyper-V agent service running on Hyper-V uses the 8087 port for interactions with Move. This service requires only the 8087 port and you cannot customize this service to use any other port.
- Ensure that the service manager console is not opened during installation or removal of the Move agent.
- For using different IP address for Move agent installation on Hyper-V host, refer to Using Different IP Address for Move Agent Installation on Hyper-V Host on page 83.

Automatic Deployment

Note:

- Before the deployment, the source Hyper-V server must have WinRM enabled over HTTP or HTTPS. For more information, refer to Enabling WinRM on page 92.
- Automatic installation of the Move Hyper-V agent is not supported for Windows Server 2012 (standalone and cluster).
- The Move Hyper-V agent service running on Hyper-V uses the 8087 port for interactions with Move. This service requires only the 8087 port and you cannot customize this service to use any other port.

Move Hyper-V agent is pushed and installed automatically to the Hyper-V source when the Hyper-V server is added as a source in the Move UI.

Using Different IP Address for Move Agent Installation on Hyper-V Host

You can use different IP address for Move agent installation on Hyper-V host other than Move eth0 IP address.

About this task

To use different IP address for Move agent installation on Hyper-V host, do the following:

Procedure

1. Go to the path /opt/xtract-vm/conf/.

```
cd /opt/xtract-vm/conf/
```

2. Create srcagent.json.
3. Add the following content to the file.

```
{
    "HyperVProviderConfig":
```

```
        {
            "MoveIPForHyperVAgentInstallation": "Move interface IP"
        }
}
```

Replace `Move interface IP` with the IP address to be used for Move agent installation.

4. Restart **srcagent** service.

   For multiple interfaces on Move, refer to *Nutanix KB article* 7399.

Creating a Migration Plan

You can create a migration plan to seed the data, cutover, and monitor the VMs. You can create the migration plan in Move without initiating the cutover process.

About this task

Note:

- This procedure is only applicable for migration from Hyper-V to NC2 on AWS and AHV.
- If the source boot is set to UEFI, set up the boot device manually in the VM post migration for the following operating systems.
  - CentOS 7.4 2, 6.8 3, 8, 8.1, 8.2
  - Ubuntu 12.04 4, 19.04
  - OEL 7 5
  - RHEL 6.8, 8.1

  For more information about setting up the boot device, refer to AHV Administration Guide.
- The Move UI does not display a message or warning during the creation of a migration plan if you are attempting to migrate VMs running on an unsupported guest operating systems.
- If you are logging in for the first time, log on to the Move UI with your defaults credentials.
- If you cancel or discard the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.

  For information on canceling an ongoing migration, see Pausing or Canceling a VM Migration on page 257.

  For information on performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

To create a migration plan, do the following:

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.

3. On the Move dashboard, click **Create a Migration Plan**.

   > **Note:** If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.

4. Enter the new migration plan name, and then click **OK**.

   > **Note:** You can edit the migration plan name by clicking the pencil icon next to the migration plan name.

   The **New Migration Plan** window appears.

5. Complete the following fields, and then click **Next**.

   a. **Select a Source**: Select an added Hyper-V source for migration.
   Once you select the source, an appropriate target appears.

   > **Note:** You might at times see a message relating to inventory collection as shown below. During this time, the environments undergoing Refresh will not be available for selection. Such environments do not automatically show up for selection once the inventory collection is over. In case you wish to select one of the environments undergoing Refresh (not available for selection), you will need to select **Cancel** and wait for inventory collection to get over and then create the migration plan again.

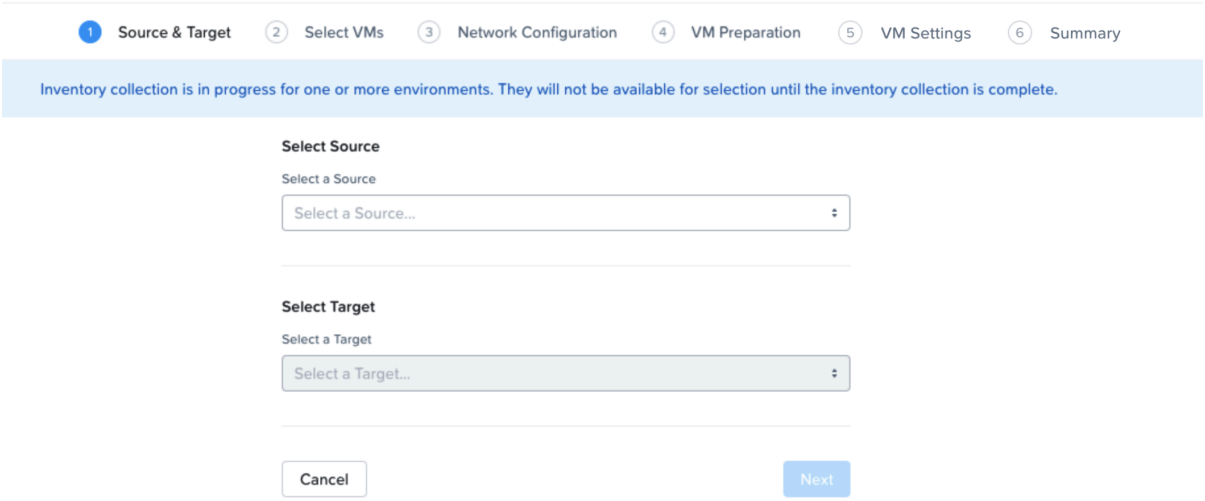

Figure 15: Inventory Collection Message

   b. **Select a Target**: Select the target for the migrating VMs.

   c. **Target Project** (optional): Select the project you want as the target.
   This field is available only with Prism Central and when AHV is the target.

   d. **Target Owners**: Select the owners for the selected target project.
   This field is available only when a target project is selected.

   e. **Target Cluster**: Select the cluster on which you will migrate the VMs.

   f. **Target Container**: Select the container on which you will migrate the VMs.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click the **+** icon next to **Name**, and then click **Next**.

   > **Note:** You cannot add more than 50 VMs in a single migration plan.

   You can filter the VMs by name in the **Filter** bar. You can also sort the VM types by selecting the appropriate column.

   > **Note:** The VMs for which migration has failed are not displayed. To show the entire list of VMs, select **All VMs** from the drop-down list. A question mark icon appears beside an unavailable VM, which displays more information about that VM and indicates why the VM cannot be migrated.

   > **Note:** Migrate VMs retain their power state on AHV.

   The selected VMs are displayed in the sidebar.

7. In the **Network Configuration** screen, select the target network from the drop-down.

   - [Applicable only if Prism Central is used] Move provides the option to select VPC-based or VLAN-based target subnets. Based on the selection of a VPC or VLAN ID as the target network, the respective subnets are listed in the target subnet drop-down menu. Select the required subnet from the drop-down menu.

     > **Note:** Overlay subnets which do not have IP address pool(s) associated will be disabled in the subnet drop-down menu.

   For performing test migration, refer to Creating a Test Capable VM Migration Plan on page 248 section.

   Click **Next**.

8. In the **VM Preparation** screen, under the **Preparation Mode** drop-down, select **Automatic**. You can also select one of the following VM preparation modes.

   » **Automatic**. Move automatically runs scripts on the source VMs to prepare them for migration. Provide the credentials of the source VMs under **Windows VMs** or **Linux VMs**, depending on the type of the source VM.

   > **Note:**
   >
   > - For Windows VMs, Move supports only username-password sign-in option for authentication. It does not support username-PIN sign-in option.
   > - Currently, the default location where the preparation scripts are stored is the /tmp folder.
   >
   >   If the /tmp folder is mounted as noexec, then Move will fallback to the /var/tmp folder. If the /var/tmp folder is also mounted as noexec, then Move will fallback to the /usr/tmp folder.

   For more information, refer to Automatic VM Preparation for Hyper-V on page 90.

   » **Manual**. Move displays the VM preparation scripts for Windows and Linux VMs.

   » **Mixed**. Move allows you to select the VM preparation mode for each VM. This setting allows you to manually prepare one set of VMs and automatically prepare another set of VMs in the same migration plan.

If you want to have such a hybrid combination of **Automatic** and **Manual** VM preparation modes, proceed to step 10.

> **Note:** The preparation mode automatically switches to **Mixed** if you change the mode of preparation for a set of VMs under **Change Settings** and have a mix of **Automatic** and **Manual** VM preparation modes. You cannot manually select the **Custom** option from the **Preparation Mode** drop-down list.

To perform data-only migration, refer to Performing Data-Only Migration on page 93.

9. (Optional) Under the **Guest Operations** section in the **VM Preparation** screen, do one or more of the settings:

   a. **Retain static IP addresses from source VMs**: Retains static IP addresses from the source VMs to the migrated VMs on AHV. By default, this option is enabled. Clear the checkbox if you do not want to retain the static IP addresses from the source VMs on the target VMs. If you disable this option, the static network is converted to DHCP network.

   For **Manual** preparation mode, if you do not want to do not want to retain the static IP addresses from the source VMs on the target VMs, do the following:

   - For Windows VMs, change the argument from `$retainIP = $true` to `$retainIP = $false` in the script.
   - For Linux VMs, remove the argument --retain-ip from the script.

   b. (Only for Manual preparation mode) **Re-Generate Script**: This option gets enabled only when there is a change in the selection of the above-mentioned settings under the **Guest Operations** section. If enabled, click this option to re-generate the VM preparation scripts for both Windows and Linux VMs.

10. (For manual preparation mode only) To manually download and run migration preparation software, select **Manual**, and then run the scripts provided in **VM Preparation** on the respective guest VMs.

    > **Note:** This step is a mandatory. Do not start the VM migration if the script execution fails.
    >
    > Ensure the following.
    >
    > - If you are preparing the VM one more time, the `C:\Nutanix` folder must not be present on the guest Windows VM before running the VM preparation scripts.
    > - The guest VM must be reachable from Move.
    > - Run the preparation scripts with administrator or root permissions.

    These scripts prepare the VMs by installing Nutanix VirtIO device drivers.

11. In the **Override individual VM Preparation** section, click **Change Settings** to override the **Guest Operations** settings (configured in the above steps) for the individual VMs. You can also edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

    > **Note:** If you do any of the following for a VM, then copy the new generated scripts of that specific VM and run them on the source VM:
    >
    > - Change the **Mode of Preparation** of a VM to **Manual**.
    > - Change any of the guest operation settings of a VM with the preparation mode set to **Manual** (an icon appears next to the VM Name prompting to regenerate a new guest script).
    >
    > Access the newly generated VM preparation script for that VM by selecting the Copy Scripts icon under the **Actions** column.

    After making the required changes, click **Done**.

12. (Optional) In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

    a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

    b. **Timezone**: Set the timezone as the hardware clock of the VMs in target.

    If set to default, Move configures UTC timezones for Linux VMs and cluster timezones for Windows VMs.

    c. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

    Only those categories which have values are available for selection.

    d. **Retain MAC Addresses from the Source VMs**: Select this check box to retain the MAC addresses from the source VMs.

    e. **Skip CDROM addition on target VMs**: Select this check box to skip the CDROM addition on the target VMs.

    f. **Enable Memory Overcommit**: Select this option to enable memory overcommit on the target VM.

    For more information on memory overcommit deployment, refer to AHV Administration Guide.

    g. **VM Migration Type**: Select one of the following VM migration types.

    At the VM level, some of the target VM properties can be customized manually after the migration plan is created. For information on manually customizing the target VM configuration, refer to Customizing the Target VM Configuration on page 255.

    - **Configure Target VM Properties**: The target VM synchronizes with the source VM properties at the time of migration plan creation. Selecting this option allows you to edit the target VM properties at the

VM level (during migration). For information on editing the target VM properties, refer to Customizing the Target VM Configuration on page 255.

> **Note:** At the VM level, only the following properties can be edited:
>
> - Target VM name
> - Number of vCPUs
> - Number of cores per vCPU
> - Memory
> - Power state

- **Retain Source VM Properties**: The target VM synchronizes with the source VM properties whenever Move refreshes the source VM configuration details. Only the customizable properties are refreshed on the target. Selecting this option does not allow you to edit the target VM properties at the VM level.

  > **Note:**
  >
  > - The source VM properties are refreshed in the following ways.
  >   - (Manually) When you click the **Refresh Source VM Properties** button.
  >   - (Automatically) When you start a migration plan.
  >   - (Automatically) When you initiate a cutover.
  > - When you start a migration plan, Move refreshes both source VM and target VM properties by default. However, it will not refresh the target VM properties at the start of a migration plan if you modified the target VM properties after migration plan creation.

h. **Settings for individual VMs**: Click **Change Settings** to configure settings such as timezone, retain MAC addresses, VM priority, and skip CDROM addition for individual VMs. You can also search the VM by typing the name of the VM and change the settings.

i. **Schedule Data Seeding**: Check this check box to select the date and time for migration.

> **Note:** For migration of VMs with multiple NICs, static IP address configurations are applied correctly if the MAC address retention is applied from source VMs, otherwise best effort IP address configuration is done by mapping one NIC at a time.

This action does not affect the VM data migration but requires you to manually prepare the guest operating system with the necessary AHV drivers prior to cutover. In addition, if you bypass the guest operations, you have to take care of the static IP address retention separately.

13. In the **Summary** screen, choose one of the following, and then proceed review the VM migration summary.

    » **Back**: Click this option to edit the information.

    » **Save**: Click this option to save the migration plan.

    For more information about how to start the migration later, and check the migrated VM status and details, refer to Environments and Migration Plan Management on page 259.

    » **Save and Start**: Click this option to save the migration plan and begin the migration immediately

    The seeding process for migration begins. You can monitor this information by selecting **Status** for the migration plan.

    > **Note:** The seeding process can take several minutes depending on the number of VMs.

What to do next

- To customize the target VM configuration at the VM level, refer to Customizing the Target VM Configuration on page 255.
- If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover on page 94.

Automatic VM Preparation for Hyper-V

You can automate the guest VM preparation.

Before you begin

General prerequisites

- The VM IP address must be present in the Hyper-V manager in the **Networking** section.

  If the IP address is not present, then check the Hyper-V integration services status.
- Guest VM must be reachable from the Move VM.

Prerequisites for Windows guest VMs

- WinRM should be configured on the guest VM.

  For more information, refer to Enabling WinRM on page 92.
- The credentials provided must have administrator privileges.
- Make sure that the folder c:\Nutanix is not present on the guest VM before initiating the migration if the user is preparing the VM one more time.

Prerequisites for Linux guest VMs

- SSH service should be up and running.
- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

About this task

> **Note:**

- For powered off guest VM, Move does not validate the guest VM credentials in the preparation phase. The credentials are validated during the migration phase. If the credentials are incorrect, then you cannot retry the migration of the guest VM. You must discard the migration of the guest VM.
- If UAC is enabled or the automatic VM preparation fails for certain VMs, you can choose to use manual preparation to prepare such VMs.

  For more information about manual VM preparation, refer to Creating a Migration Plan on page 84.

To automatically prepare the VMs, do the following:

Procedure

1. In the **Preparation Mode** drop-down, select **Automatic**.

   To perform data-only migration, refer to Performing Data-Only Migration on page 93.

2. (Optional) Under the **Guest Operations** section in the **VM Preparation** screen, do the following:

   a. **Retain static IP addresses from source VMs**: Retains static IP addresses from the source VMs to the migrated VMs on AHV. By default, this option is enabled. Clear this checkbox if you do not want to retain the static IP addresses from the source VMs on the target VMs. If you disable this option, the static network is converted to DHCP network.

3. Complete the fields in **Credentials for Source VMs**. If VMs in the migration plan have same credentials, then enter the user name and password for the guest VMs to allow Move to install the necessary drivers.

4. In the **Override individual VM Preparation** section, click **Change Settings** to override the **Guest Operations** settings (configured in the above steps) for the individual VMs. You can also edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

5. (Optional) In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

   a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

   b. **Timezone**: Set the timezone as the hardware clock of the VMs in target.

      If set to default, Move configures UTC timezones for Linux VMs and cluster timezones for Windows VMs.

   c. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

      Only those categories which have values are available for selection.

   d. **Retain MAC Addresses from the Source VMs**: Select this check box to retain the MAC addresses from the source VMs.

   e. **Skip CDROM addition on target VMs**: Select this check box to skip the CDROM addition on the target VMs.

   f. **Enable Memory Overcommit**: Select this option to enable memory overcommit on the target VM.

      For more information on memory overcommit deployment, refer to AHV Administration Guide.

   g. **Settings for individual VMs**: Click **Change settings** to configure settings such as timezone, retain MAC addresses, VM priority, and skip CDROM addition for individual VMs. You can also search the VM by typing the name of the VM and change the settings.

      > **Note:** For migration of VMs with multiple NICs, static IP address configurations are applied correctly if the MAC address retention is applied from source VMs, otherwise best effort IP address configuration is done by mapping one NIC at a time.

      This action does not affect the VM data migration but requires you to manually prepare the guest operating system with the necessary AHV drivers prior to cutover. In addition, if you bypass the guest operations, you have to take care of the static IP address retention separately.

   h. **Schedule Data Seeding**: Check this check box to select the date and time for migration.

   The credentials of VMs are validated. Once the validation is successful, the Guest Tools are downloaded and installed in all the VMs of the migration plan. Then, the VMs are validated for readiness.

   > **Note:** If the validation of credentials or Guest Tools installation fails, you can update the credentials or remove the VM from the migration plan and proceed by clicking **Next**.

What to do next

If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover on page 94.

Enabling WinRM

Enable WinRM to install the Guest Tools on Windows Hyper-V VMs.

About this task

> **Note:**
>
> - This method is a prerequisite for automatic VM preparation to work with Windows Hyper-V VMs.
> - Ensure that the ingress ports 5985 and 5986 are enabled.

To enable WinRM, do the following:

Procedure

1. Open PowerShell in Windows VM.
2. Run the script to enable WinRM for Windows Hyper-V VMs.

```
> winrm quickconfig -q
winrm set winrm/config/winrs '@{MaxMemoryPerShellMB="300"}'
winrm set winrm/config '@{MaxTimeoutms="1800000"}'
winrm set winrm/config/service '@{AllowUnencrypted="true"}'
winrm set winrm/config/service/auth '@{Basic="true"}'

netsh advfirewall firewall add rule name="WinRM 5985" protocol=TCP dir=in
 localport=5985 action=allow
netsh advfirewall firewall add rule name="WinRM 5986" protocol=TCP dir=in
 localport=5986 action=allow

net stop winrm
cmd /c 'sc config winrm start= auto'
net start winrm
```

3. Run the script to enable Secure Sockets Layer (SSL).

```
> $c = New-SelfSignedCertificate -DnsName "$(hostname)" -CertStoreLocation cert:
\LocalMachine\My
winrm create winrm/config/Listener?Address=*+Transport=HTTPS
 "@{Hostname=`"$(hostname)`";CertificateThumbprint=`"$($c.ThumbPrint)`"}"
```

Performing Data-Only Migration

Move performs data-only migration when you select **Automatic** preparation mode while creating a migration plan, and bypass the guest operations or does not provide the source VM credentials while preparing a migration plan. Or when you select **Manual** preparation mode while creating a migration plan, and do not run the preparation script in the source VMs. In data-only migration, Move skips the source VM guest operating system preparation tasks which includes installing VirtIO driver and copying of the scripts to retain the IP address. It also skips the uninstallation of VMware tools after migration.

About this task

Note: Data-only migration is only supported for the following migrations:

- From ESXi to AHV and ESXi to NC2 on AWS
- From Hyper-V to AHV and Hyper-V to NC2 on AWS
- From Hyper-V to ESXi
- From AHV to AHV

To perform data-only migration, do the following:

Procedure

1. In the **VM Preparation** screen, if you select **Automatic**, then proceed without providing the credentials for the source VMs or select the **Bypass Guest Operations on Source VMs** check box or if you select **Manual**, do not run the preparation script in the source VMs.
   The following message appears when the **Automatic** preparation mode is selected,

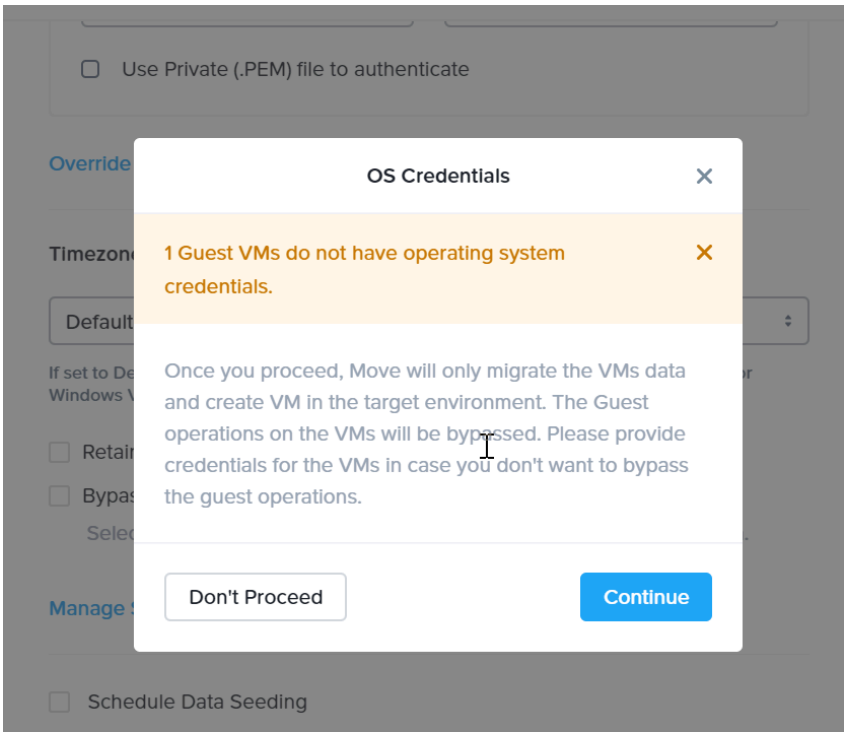

Figure 16: OS Credentials Dialog Box

2. Click **Continue**.
   Move migrates the VMs data and creates a VM in the target, and bypasses the operating system operations.

Performing a Migration Cutover

When the migration plan is started and the seeding process is complete, you can cut over the selected VMs in the AHV cluster. You can monitor the VM migration progress by clicking the **Status** link.

About this task

Note:

- You can perform this operation only when the VM status is Ready to Cutover.
- Recommended that cutover should be performed within one week of initial data seeding.
- If the initial data seeding finishes in less than 10 minutes, the Move continues to wait for 10 minutes to take the incremental snapshot; however, you can trigger the cutover immediately.
- The cutover process increments in absolute numbers.
- For cutover of Hyper-V source VMs, it is recommended to go with maximum 30 VMs at a time.

To perform a cutover, do the following:

Procedure

1. In the Move UI, click the **Ready to Cutover** status to display the list of available VMs.
2. Select the VMs or group of VMs to cut over.
3. Click **Cutover**.
   The cutover process performs the following VM actions.

   - Shuts down the VM
   - Takes the final snapshots for the VM and copying the final changes to AHV after the last 10 minute interval
   - Adds a note in the VM in the Hyper-V manager
   - Disconnects the source VM network interfaces
   - Creates a VM in the target AHV cluster
   - Attaches the replicated disks to the VM
   - Powers on or off the VM (depends on the initial power state)
   - Runs the scripts to set the static IP address

   The cutover process begins immediately and might take a few minutes. Once cutover is complete, the VM is ready for use in the new AHV cluster.

What to do next

You can manage the migration plan once the migration plan is ready. For more information, refer to Environments and Migration Plan Management on page 259

Performance Matrix for Large Data Migration

Move performs end-to-end migration of large VMs. The scenarios are tested based on the following parameters. The following tables show the performance numbers from the Move lab.

Table 12: Performance Numbers of Large Data Migration (Hyper-V to AHV)

Total Migration size	Number of VMs	Size of each Vdisks	Number of Vdisks	I/O on the source	Data churn	Network bandwidth	Migration time taken	
							Data seeding	Cutover
2 TB	1	2 TB	1	No	No I/O	10 G	430 minutes	Less than 5 minutes
2 TB	1	1 TB	2	No	No I/O	10 G	260 minutes	Less than 5 minutes
2 TB	1	512 GB	4	No	No I/O	10 G	215 minutes	Less than 5 minutes
2 TB	1	256 GB	8	No	No I/O	10 G	225 minutes	Less than 5 minutes
2 TB	20	100 GB	20	No	No I/O	10 G	210 minutes	Less than 5 minutes

HYPER-V TO ESXI

You can prepare and migrate VMs running on Hyper-V hypervisor to ESXi on Nutanix by using Move.

> **Note:** ESXi as a target refers to ESXi running on Nutanix appliances.

Migration Considerations

You must consider the supported guest operating systems, requirements, recommendations, unsupported features, and limitations provided in this section before starting the migration process.

Supported Guest Operating Systems

Move supports some common operating systems for migration from Hyper-V to ESXi on Nutanix.

> **Note:** The Move UI does not display a message or warning during the creation of a migration plan if you are attempting to migrate VMs running on an unsupported guest operating systems.

Generation 1 VMs Support

- Windows 2008 R2, Windows 2012 R2, Windows 2016
- RHEL 6, 7
- CentOS 6, 7
- SUSE Linux Enterprise Server 12 SP2
- Ubuntu Server 12.0.4

Generation 2 VMs Support (UEFI Enabled VMs)

- Windows 2012 R2, Windows 2016
- RHEL 7
- CentOS 7
- SUSE Linux Enterprise Server 12 SP2

Support for UEFI with Secure Boot Enabled VMs

Move supports UEFI with secure boot enabled VMs.

Table 13: Supported Guest Operating Systems

Operating systems
Windows 2012 R2, Windows 2016

Requirements

Before attempting to migrate VMs running on Hyper-V using Move, make sure to conform to the requirements listed here.

General Requirements

Make sure to conform to the following requirements for Hyper-V to ESXi migration:

- Supported browser: Google Chrome.
- Ensure you have PowerShell version 4.0 or later.
- Ensure guest VMs have connectivity with Move.
- Ensure that the guest VMs have integration services installed and an IP address is present in the **Networking** section of the Hyper-V manager for automatic preparation of the source VMs.

  For more information, refer to *Manage Hyper-V Integration Services* section in the Microsoft documentation for Hyper-V.
- Ensure the following:
  - WinRM is configured on Hyper-V servers and Windows guest VM.
  - The following inbound and outbound ports using the TCP protocol are enabled for the Windows Remote Management (WinRM) feature to work.
    - WinRM-HTTPS: 5986
    - WinRM-HTTP: 5985

    Note: This requirement is applicable only for automatic VM preparation mode.
- If a local built-in administrator user performs guest preparation, admin approval mode should be disabled. By default, admin approval mode is in disabled state. If the admin approval mode is in enabled state, refer to KB 7672 in the Nutanix Support Portal.
- The source should have sufficient space for the snapshot disks created during migration.
- Network Security Services (NSS) 3.44 is required for Linux VMs.
- Before you initiate a migration, ensure to disable backups.

  Note: If the Microsoft Hyper-V Server version is 2016 or higher, Nutanix recommends to configure the checkpoints as production checkpoints for better migration performance. For information about the procedure to change checkpoints to production checkpoints, refer to Microsoft documentation.

Prerequisites for Linux guest VMs:

- SSH service should be up and running.
- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

Service Accounts

For successful migration of VMs from Hyper-V to ESXi, service accounts for the following must have admin privileges.

- Source Hyper-V Server (standalone or cluster)
- vCenter Server
- Prism Element UI for the ESXi on Nutanix cluster
- Source VMs planned for migration

  A user with administrator role for Windows source VMs or a root user for Linux source VMs to run the source VM preparation scripts.

Limitations

This section lists the limitations for migration from Hyper-V to ESXi.

In addition to Migration Limitations on page 12, the support is constrained by the following.

- Migration of guest VMs might fail if you delete an existing user snapshot or checkpoint during the migration.
- VM names with single and double quotes are not supported for migration.
- If VMs are configured with dynamic memory, Move takes only the start-up memory configuration while creating the target VM on ESXi.
- Automatic installation of the Move Hyper-V agent is not supported for Windows Server 2012 (standalone & cluster).
- VMs protected by Hyper-V replica are not supported.

  To allow migration, disable the Hyper-V replication for the required VMs.
- VMs migrated with more than one network interfaces might not retain all the IP addresses.

  Workaround: Manually assign IP addresses after migration.
- After migration, Windows VMs with connected NICs only will retain their IP address. Those with disconnected NICs will not retain their IP address.

Adding Hyper-V Environment

While creating a migration plan for migration from Hyper-V to any target, be sure to add at least one Hyper-V environment for migration.

About this task

> **Note:** This procedure is only applicable for migration from Hyper-V.

To add a Hyper-V environment, do the following:

Procedure

1. Log on to Move UI.

2. Click **+ Add Environment** under Environments.
   The **Add Environment** appears.

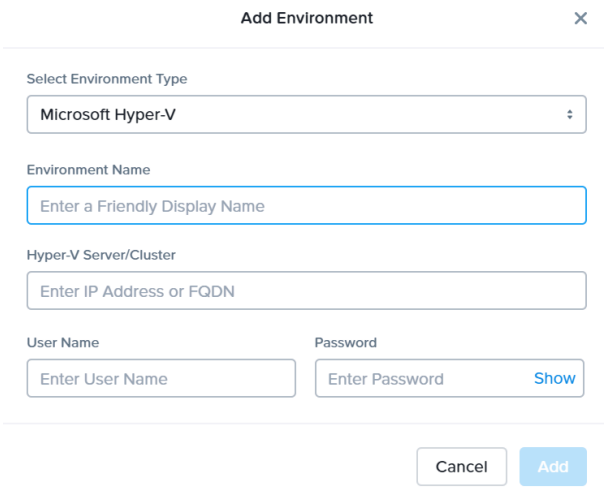

Figure 17: Add Hyper-V Environment Dialog Box

3. Select **Microsoft Hyper-V** as the source environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the Hyper-V environment.
   b. **Hyper-V Server**: Enter the IP address or FQDN of the Hyper-V source server.
   c. **Username**: Enter a Windows account with administrator privileges of the Hyper-V source server.
   d. **Password**: Enter the password for this Windows account.

   Note:

   - Hyper-V user account must have administrator privileges.
   - Enter cluster IP address (or failover IP address) to discover all clustered VMs from Hyper-V cluster environment.
   - For cluster environment, enter the domain credentials the administrator privileges.
   - You must either add the cluster IP address or the standalone IP address of the node.

   The Hyper-V environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

You can also add a Nutanix AHV environment. For more information, refer to Adding a Nutanix AOS Cluster Environment on page 99

Adding a Nutanix AOS Cluster Environment

While creating a migration plan, if you need to add AHV as the source or target, or ESXi on Nutanix as a target, then you have to add at least one AOS cluster environment for migration.

About this task

> **Note:** When you add a AOS cluster, Move VM IP address with the subnet 255.255.255.255 is added to the global NFS allowlist.
>
> > **Caution:** Modifying the NFS allowlist of the destination container disables inheritance of the NFS allowlist at the global level and lead to issues with Move operations.

To add a Nutanix AOS cluster environment, do the following:

Procedure

1. Log on to the Move UI.
2. Click **+ Add Environment** under **Environments**.

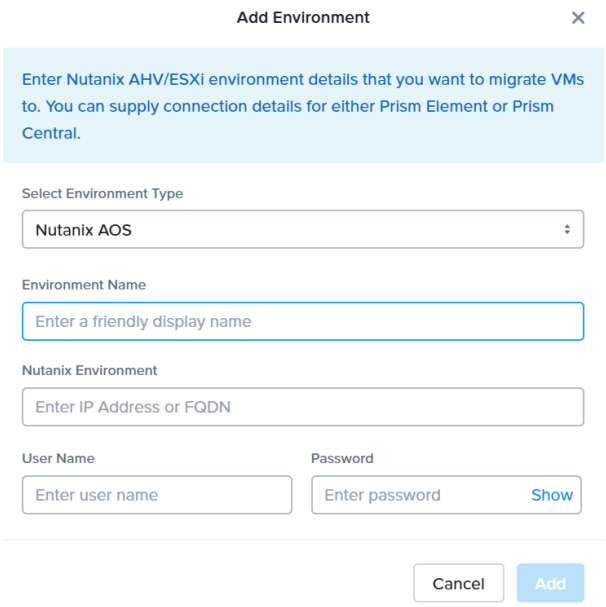

Figure 18: Add AOS Environment Dialog Box

The **Add Environment** window appears.

3. Select **Nutanix AOS** as the target environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the NC2 on AWS and AHV or VMware ESXi on Nutanix environment.
   b. **Nutanix Environment**: Enter the IP address or the FQDN for the target Prism Central or Prism Element.
   c. **User Name**: Enter the username for logging on to the target Nutanix environment.
   d. **Password**: Enter the password for logging on to the target Nutanix environment.

5. (Only for ESXi to ESXi migration) Enter credentials for registered vCenter.

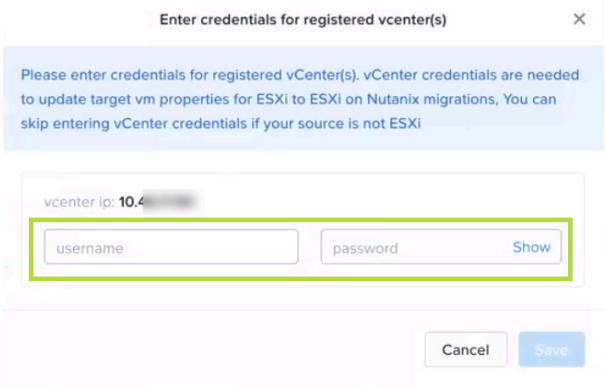

Figure 19: vCenter Credentials Dialog box

vCenter credentials are required to update the target VM properties for ESXi to ESXi on Nutanix migration.

> **Note:** You can skip adding vCenter credentials if your source is not ESXi.

The AOS environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

Once you have added the environments, you can proceed to create the migration plan. For more information, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan (Hyper-V to ESXi) on page 103 or Creating a Migration Plan (Azure to AHV) on page 171 or Creating a Migration Plan (Azure to ESXi) on page 184

Deploying the Move Agent on Hyper-V Host

You can manually or automatically deploy the Move agent on the source Hyper-V host.

Manual Deployment

To download and install the Move agent on each of the source Hyper-V host machine to support VM discovery and migrations, do the following:

1. In the Hyper-V host, download **move-agent-installer.exe** from **http://<nutanix-move-ip>/downloads/agents/move-agent-installer.exe** .

   Replace `<nutanix-move-ip>` with the IP address of the Move VM.
2. Go to the location where you have downloaded the agent, and copy or move the downloaded file under C:\users \Administrator.
3. Launch the command prompt with **Run as Administrator** to run the following command from C:\users \Administrator.

   ```
   move-agent-installer.exe -o [operation] -ip [move ip] -u [user]
   ```

   Example: `move-agent-installer.exe -o install -ip 10.5.244.55 -u user`

   User can be either a domain or local user with administrator privileges.
4. Enter the password when the command prompt requests for it.

Note:

- Do not use **-p** option in the command to include the password if the password contains special characters. The PowerShell console throws an exception when a password containing special characters is included in the command using **-p** option.
- By default, the Move agent is installed in the user directory. If you want to change the location, use the **-d** option.
- To uninstall the Move agent, use the command `move-agent-installer.exe -o remove`.
- Enter help along with any command for information to display options for that command.
- Move agent service installation will add inbound firewall rule to open port 8087, which is required for Move VM and Hyper-V interactions. The Move Hyper-V agent service running on Hyper-V uses the 8087 port for interactions with Move. This service requires only the 8087 port and you cannot customize this service to use any other port.
- Ensure that the service manager console is not opened during installation or removal of the Move agent.
- For using different IP address for Move agent installation on Hyper-V host, refer to Using Different IP Address for Move Agent Installation on Hyper-V Host on page 102.

Automatic Deployment

Note:

- Before the deployment, the source Hyper-V server must have WinRM enabled over HTTP or HTTPS. For more information, refer to Enabling WinRM on page 111.
- Automatic installation of the Move Hyper-V agent is not supported for Windows Server 2012 (standalone and cluster).
- The Move Hyper-V agent service running on Hyper-V uses the 8087 port for interactions with Move. This service requires only the 8087 port and you cannot customize this service to use any other port.

Move Hyper-V agent is pushed and installed automatically to the Hyper-V source when the Hyper-V server is added as a source in the Move UI.

Using Different IP Address for Move Agent Installation on Hyper-V Host

You can use different IP address for Move agent installation on Hyper-V host other than Move eth0 IP address.

About this task

To use different IP address for Move agent installation on Hyper-V host, do the following:

Procedure

1. Go to the path /opt/xtract-vm/conf/.

```
cd /opt/xtract-vm/conf/
```

2. Create srcagent.json.
3. Add the following content to the file.

```
{
    "HyperVProviderConfig":
```

```
        {
            "MoveIPForHyperVAgentInstallation": "Move interface IP"
        }
}
```

Replace `Move interface IP` with the IP address to be used for Move agent installation.

4. Restart **srcagent** service.

   For multiple interfaces on Move, refer to *Nutanix KB article* 7399.

Creating a Migration Plan (Hyper-V to ESXi)

You can create a migration plan to seed the data, cutover, and monitor the VMs. You can create the migration plan in Move without initiating the cutover process.

About this task

Note:

- This procedure is only applicable for migration from Hyper-V to ESXi.
- The Move UI does not display a message or warning during the creation of a migration plan if you are attempting to migrate VMs running on an unsupported guest operating systems.
- If you are logging in for the first time, log on to the Move UI with your defaults credentials.
- If you cancel or discard the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.

  For information on canceling an ongoing migration, see Pausing or Canceling a VM Migration on page 257.

  For information on performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

To create a migration plan, do the following:

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.
3. On the Move dashboard, click **Create a Migration Plan**.

   Note: If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.
4. Enter the new migration plan name, and then click **OK**.

   Note: You can edit the migration plan name by clicking the pencil icon next to the migration plan name.

   The **New Migration Plan** window appears.

5. Complete the following fields, and then click **Next**.

   a. **Select a Source**: Select an added Hyper-V source for migration.
   Once you select the source, an appropriate target appears.

   > **Note:** You might at times see a message relating to inventory collection as shown below. During this time, the environments undergoing Refresh will not be available for selection. Such environments do not automatically show up for selection once the inventory collection is over. In case you wish to select one of the environments undergoing Refresh (not available for selection), you will need to select **Cancel** and wait for inventory collection to get over and then create the migration plan again.

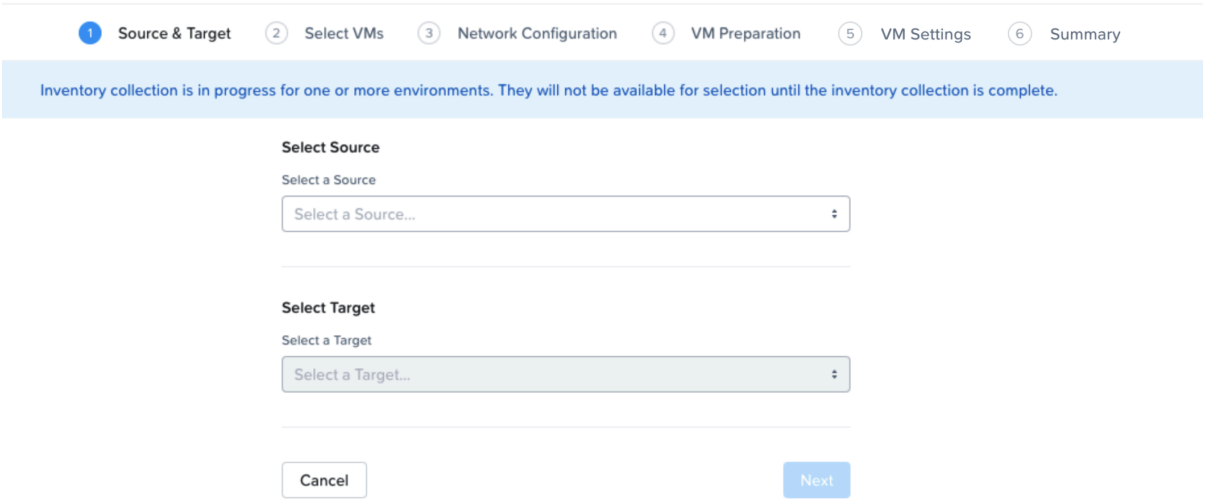

Figure 20: Inventory Collection Message

   b. **Select a Target**: Select the ESXi target for the migrating VMs.

   c. **Target Containers**: Select the container on which you will migrate the VMs.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click the **+** icon next to **Name**, and then click **Next**.

   > **Note:** You cannot add more than 50 VMs in a single migration plan.

   You can filter the VMs by name in the **Filter** bar. You can also sort the VM types by selecting the appropriate column.

   > **Note:** The VMs for which migration has failed are not displayed. To show the entire list of VMs, select **All VMs** from the drop-down list. A question mark icon appears beside an unavailable VM, which displays more information about that VM and indicates why the VM cannot be migrated.

   > **Note:** Migrate VMs retain their power state on AHV.

   The selected VMs are displayed in the sidebar.

7. In the **Network Configuration** screen, select the target network from the drop-down, and then click **Next**.
   For performing test migration, refer to Creating a Test Capable VM Migration Plan on page 248 section.

8. In the **VM Preparation** screen, under the **Preparation Mode** drop-down, select **Automatic**. You can also select one of the following VM preparation modes.

   » **Automatic**. Move automatically runs scripts on the source VMs to prepare them for migration.

   » **Manual**. Move displays the VM preparation scripts for Windows and Linux VMs. Copy the scripts and manually run them on the source VMs.

   » **Mixed**. Move allows you to select the VM preparation mode for each VM. This setting allows you to manually prepare one set of VMs and automatically prepare another set of VMs in the same migration plan. If you want to have such a hybrid combination of **Automatic** and **Manual** VM preparation modes, proceed to step 9.

   > **Note:** The preparation mode automatically switches to **Mixed** if you change the mode of preparation for a set of VMs under **Change Settings** and have a mix of **Automatic** and **Manual** VM preparation modes. You cannot manually select the **Custom** option from the **Preparation Mode** drop-down list.

   To perform data-only migration, refer to Performing Data-Only Migration (Hyper-V to ESXi) on page 112.

9. Under the **Guest Operations** section in the **VM Preparation** screen, do one or more of the settings:

   a. **Retain static IP addresses from source VMs**: Retains static IP addresses from the source VMs to the migrated VMs on AHV. By default, this option is enabled. Clear the checkbox if you do not want to retain the static IP addresses from the source VMs on the target VMs. If you disable this option, the static network is converted to DHCP network.

      For **Manual** preparation mode, if you do not want to do not want to retain the static IP addresses from the source VMs on the target VMs, do the following:

      - For Windows VMs, change the argument from `$retainIP = $true` to `$retainIP = $false` in the script.
      - For Linux VMs, remove the argument --retain-ip from the script.

   b. (Only for Automatic preparation mode) **Bypass Guest Operations on Source VMs**: Select this check box to bypass the guest operating system changes.

      You can select this option to override your migration to data-only migration.

      > **Note:** If you bypass guest operations on source VMs, **Retain static IP addresses from source VMs** and **Uninstall VMware Tools** configuration will not be applicable.

   c. (Only for Manual preparation mode) **Re-Generate Script**: This option gets enabled only when there is a change in the selection of the above-mentioned settings under the **Guest Operations** section. If enabled, click this option to re-generate the VM preparation scripts for both Windows and Linux VMs.

10. (For Automatic preparation mode only) You must provide the credentials of the source VMs under **Windows VMs** or **Linux VMs**, depending on the type of the source VM.

    > **Note:** For Windows VMs, Move supports only username-password sign-in option for authentication. It does not support username-PIN sign-in option.

    For more information, refer to Automatic VM Preparation for Hyper-V on page 109.

11. (For Manual preparation mode only) To manually download and run migration preparation software, select **Manual**, and then run the scripts provided in **VM Preparation** on the respective guest VMs.

    > **Note:** This step is a mandatory. Do not start the VM migration if the script execution fails.
    >
    > Ensure the following.
    >
    > - If you are preparing the VM one more time, the `C:\Nutanix` folder must not be present on the guest Windows VM before running the VM preparation scripts.
    > - The guest VM must be reachable from Move.
    > - Run the preparation scripts with administrator or root permissions.

    These scripts prepare the VMs by installing Nutanix VirtIO device drivers.

12. (For Mixed preparation mode only) If you want to have a combination of **Automatic** and **Manual** VM preparation modes, do the following.

    a. Click **Change Settings** under **Override individual VM Preparation** section.

    b. For the VMs whose mode of preparation you want to change, select the mode of preparation (**Automatic** or **Manual**) and click **Done**.

    > **Note:** If you select **Manual**, you must copy the displayed scripts and run them on the source VMs.

13. In the **Override individual VM Preparation** section, click **Change Settings** to override the **Guest Operations** settings (configured in the above steps) for the individual VMs. You can also edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

    > **Note:** If you do any of the following for a VM, then copy the new generated scripts of that specific VM and run them on the source VM:
    >
    > - Change the **Mode of Preparation** of a VM to **Manual**.
    > - Change any of the guest operation settings of a VM with the preparation mode set to **Manual** (an icon appears next to the VM Name prompting to regenerate a new guest script).
    >
    > Access the newly generated VM preparation script for that VM by selecting the Copy Scripts icon under the **Actions** column.

    After making the required changes, click **Done**.

    Then, click **Next**.

14. In the **VM Settings** screen, do the following, and then click **Next**.

    a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

    b. **Timezone**: Set the timezone as the hardware clock of the VMs in target.

    If set to default, Move configures UTC timezones for Linux VMs and cluster timezones for Windows VMs.

    c. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

    Only those categories which have values are available for selection.

    d. **VM Migration Type**: Select one of the following VM migration types.

    At the VM level, some of the target VM properties can be customized manually after the migration plan is created. For information on manually customizing the target VM configuration, refer to Customizing the Target VM Configuration on page 255.

    - **Configure Target VM Properties**: The target VM synchronizes with the source VM properties at the time of migration plan creation. Selecting this option allows you to edit the target VM properties at the

VM level (during migration). For information on editing the target VM properties, refer to Customizing the Target VM Configuration on page 255.

> **Note:** At the VM level, only the following properties can be edited:
>
> - Target VM name
> - Number of vCPUs
> - Number of cores per vCPU
> - Memory
> - Power state

- **Retain Source VM Properties**: The target VM synchronizes with the source VM properties whenever Move refreshes the source VM configuration details. Only the customizable properties are refreshed on the target. Selecting this option does not allow you to edit the target VM properties at the VM level.

  > **Note:**
  >
  > - The source VM properties are refreshed in the following ways.
  >   - (Manually) When you click the **Refresh Source VM Properties** button.
  >   - (Automatically) When you start a migration plan.
  >   - (Automatically) When you initiate a cutover.
  > - When you start a migration plan, Move refreshes both source VM and target VM properties by default. However, it will not refresh the target VM properties at the start of a migration plan if you modified the target VM properties after migration plan creation.

e. **Settings for individual VMs**: Click **Change Settings** to configure settings such as timezone and VM priority for individual VMs. You can also search the VM by typing the name of the VM and change the settings.

f. **Schedule Data Seeding**: Select this checkbox to select the date and time for migration.

> **Note:** For migration of VMs with multiple NICs, static IP address configurations are applied correctly if the MAC address retention is applied from source VMs, otherwise best effort IP address configuration is done by mapping one NIC at a time.

This action does not affect the VM data migration but requires you to manually prepare the guest operating system with the necessary AHV drivers prior to cutover. In addition, if you bypass the guest operations, you have to take care of the static IP address retention separately.

15. In the **Summary** screen, choose one of the following, and then proceed review the VM migration summary.

    » **Back**: Click this option to edit the information.

    » **Save**: Click this option to save the migration plan.

    For more information about how to start the migration later, and check the migrated VM status and details, refer to Environments and Migration Plan Management on page 259.

    » **Save and Start**: Click this option to save the migration plan and begin the migration immediately

    The seeding process for migration begins. You can monitor this information by selecting **Status** for the migration plan.

    > **Note:** The seeding process can take several minutes depending on the number of VMs.

What to do next

- To customize the target VM configuration at the VM level, refer to Customizing the Target VM Configuration on page 255.
- If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover on page 113.

Automatic VM Preparation for Hyper-V

You can automate the guest VM preparation.

Before you begin

General prerequisites

- The VM IP address must be present in the Hyper-V manager in the **Networking** section.

  If the IP address is not present, then check the Hyper-V integration services status.
- Guest VM must be reachable from the Move VM.

Prerequisites for Windows guest VMs

- WinRM should be configured on the guest VM.

  For more information, refer to Enabling WinRM on page 111.
- The credentials provided must have administrator privileges.
- Make sure that the folder c:\Nutanix is not present on the guest VM before initiating the migration if the user is preparing the VM one more time.

Prerequisites for Linux guest VMs

- SSH service should be up and running.
- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

About this task

> **Note:**

- For powered off guest VM, Move does not validate the guest VM credentials in the preparation phase. The credentials are validated during the migration phase. If the credentials are incorrect, then you cannot retry the migration of the guest VM. You must discard the migration of the guest VM.
- If UAC is enabled or the automatic VM preparation fails for certain VMs, you can choose to use manual preparation to prepare such VMs.

  For more information about manual VM preparation, refer to Creating a Migration Plan (Hyper-V to ESXi) on page 103.

To automatically prepare the VMs, do the following:

Procedure

1. In the **Preparation Mode** drop-down, select **Automatic**.
   To perform data-only migration, refer to Performing Data-Only Migration (Hyper-V to ESXi) on page 112.
2. Select the settings as necessary under the **Guest Operations** section in the **VM Preparation** screen.
3. Complete the fields in **Credentials for Source VMs**. If VMs in the migration plan have same credentials, then enter the username and password for the guest VMs to allow Move to install the necessary drivers.
4. In the **Override Individual VM Settings** section, do the following, and then click **Next**.
   a. Click **Change settings** to override the settings for the individual VMs. You can update the VM preparation credentials of the VMs, remove VMs, or update the VM preparation mode.

      Note: If you select **Manual**, you must copy the displayed scripts and run them on the source VMs.

5. In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

   a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

   b. **Timezone**: Set the timezone as the hardware clock of the VMs in target.

      If set to default, Move configures UTC timezones for Linux VMs and cluster timezones for Windows VMs.

   c. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

      Only those categories which have values are available for selection.

   d. **Settings for individual VMs**: Click **Change settings** to configure settings such as timezone and VM priority for individual VMs. You can also search the VM by typing the name of the VM and change the settings.

      > **Note:** For migration of VMs with multiple NICs, static IP address configurations are applied correctly if the MAC address retention is applied from source VMs, otherwise best effort IP address configuration is done by mapping one NIC at a time.

      This action does not affect the VM data migration but requires you to manually prepare the guest operating system with the necessary AHV drivers prior to cutover. In addition, if you bypass the guest operations, you have to take care of the static IP address retention separately.

   e. **Schedule Data Seeding**: Select this checkbox to select the date and time for migration.

   The credentials of VMs are validated. Once the validation is successful, the Guest Tools are downloaded and installed in all the VMs of the migration plan. Then, the VMs are validated for readiness.

   > **Note:** If the validation of credentials or Guest Tools installation fails, you can update the credentials or remove the VM from the migration plan and proceed by clicking **Next**.

What to do next

If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover on page 113.

Enabling WinRM

Enable WinRM to install the Guest Tools on Windows Hyper-V VMs.

About this task

> **Note:**
>
> - This method is a prerequisite for automatic VM preparation to work with Windows Hyper-V VMs.
> - Ensure that the ingress ports 5985 and 5986 are enabled.

To enable WinRM, do the following:

Procedure

1. Open PowerShell in Windows VM.
2. Run the script to enable WinRM for Windows Hyper-V VMs.

```
> winrm quickconfig -q
winrm set winrm/config/winrs '@{MaxMemoryPerShellMB="300"}'
```

```
winrm set winrm/config '@{MaxTimeoutms="1800000"}'
winrm set winrm/config/service '@{AllowUnencrypted="true"}'
winrm set winrm/config/service/auth '@{Basic="true"}'

netsh advfirewall firewall add rule name="WinRM 5985" protocol=TCP dir=in
 localport=5985 action=allow
netsh advfirewall firewall add rule name="WinRM 5986" protocol=TCP dir=in
 localport=5986 action=allow

net stop winrm
cmd /c 'sc config winrm start= auto'
net start winrm
```

3. Run the script to enable Secure Sockets Layer (SSL).

```
> $c = New-SelfSignedCertificate -DnsName "$(hostname)" -CertStoreLocation cert:
\LocalMachine\My
winrm create winrm/config/Listener?Address=*+Transport=HTTPS
 "@{Hostname=`"$(hostname)`";CertificateThumbprint=`"$($c.ThumbPrint)`"}"
```

Performing Data-Only Migration (Hyper-V to ESXi)

Move performs data-only migration when you select **Automatic** preparation mode while creating a migration plan, and bypass the guest operations or does not provide the source VM credentials while preparing a migration plan. Or when you select **Manual** preparation mode while creating a migration plan, and do not run the preparation script in the source VMs. In data-only migration, Move skips the source VM guest operating system preparation tasks which includes installing VirtIO driver and copying of the scripts to retain the IP address. It also skips the uninstallation of VMware tools after migration.

About this task

Note: Data-only migration is only supported for the following migrations:

- From ESXi to AHV and ESXi to NC2 on AWS
- From Hyper-V to AHV and Hyper-V to NC2 on AWS
- From Hyper-V to ESXi
- From AHV to AHV

To perform data-only migration, do the following:

Procedure

1. In the **VM Preparation** screen, if you select **Automatic**, then proceed without providing the credentials for the source VMs or select the **Bypass Guest Operations on Source VMs** check box or if you select **Manual**, do not run the preparation script in the source VMs.
   The following message appears when the **Automatic** preparation mode is selected,

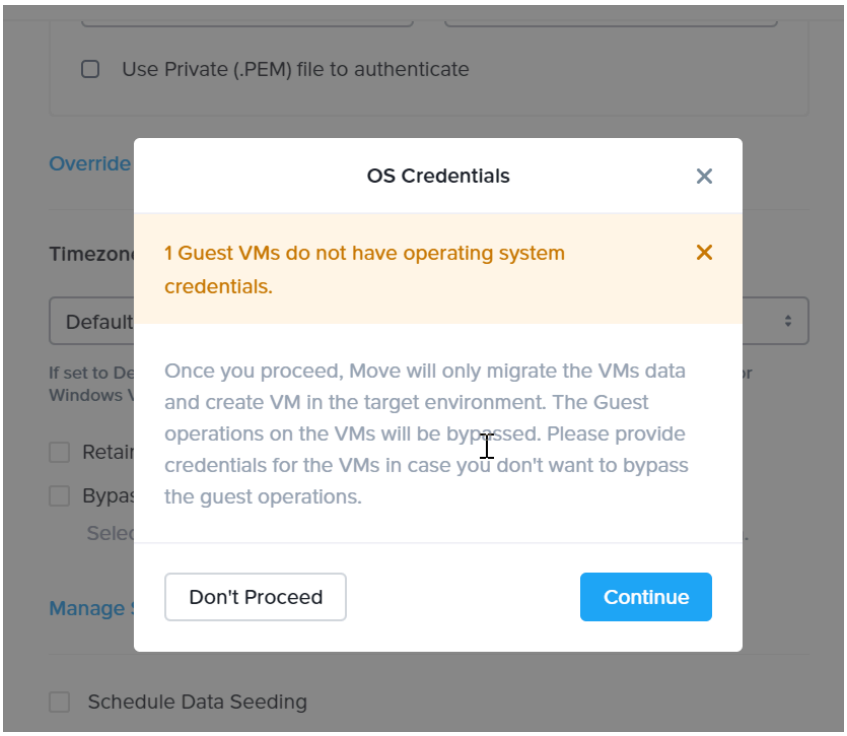

Figure 21: OS Credentials Dialog Box

2. Click **Continue**.
   Move migrates the VMs data and creates a VM in the target, and bypasses the operating system operations.

Performing a Migration Cutover

When the migration plan is started and the seeding process is complete, you can cut over the selected VMs in the ESXi cluster. You can monitor the VM migration progress by clicking the **Status** link.

About this task

Note:

- You can perform this operation only when the VM status is Ready to Cutover.
- Recommended that cutover should be performed within one week of initial data seeding.
- If the initial data seeding finishes in less than 10 minutes, the Move continues to wait for 10 minutes to take the incremental snapshot; however, you can trigger the cutover immediately.
- The cutover process increments in absolute numbers.
- For cutover of Hyper-V source VMs, it is recommended to go with maximum 30 VMs at a time.

To perform a cutover, do the following:

Procedure

1. In the Move UI, click the **Ready to Cutover** status to display the list of available VMs.
2. Select the VMs or group of VMs to cut over.
3. Click **Cutover**.
   The cutover process performs the following VM actions.

   - Shuts down the VM
   - Takes the final snapshots for the VM and copying the final changes to ESXi after the last 10 minute interval
   - Adds a note in the VM in the Hyper-V manager
   - Disconnects the source VM network interfaces
   - Creates a VM in the target ESXi cluster
   - Attaches the replicated disks to the VM
   - Powers on or off the VM (depends on the initial power state)
   - Runs the scripts to set the static IP address

   The cutover process begins immediately and might take a few minutes. Once cutover is complete, the VM is ready for use in the new ESXi cluster.

What to do next

You can manage the migration plan once the migration plan is ready. For more information, refer to Environments and Migration Plan Management on page 259

Performance Matrix for Large Data Migration

Move performs end-to-end migration of large VMs. The scenarios are tested based on the following parameters. The following tables show the performance numbers from the Move lab.

Table 14: Performance Numbers of Large Data Migration (Hyper-V to ESXi)

Total Migration size	Number of VMs	Size of each Vdisks	Number of Vdisks	I/O on the source	Data churn	Network bandwidth	Migration time taken	
							Data seeding	Cutover
2 TB	1	2 TB	1	No	No I/O	10 G	430 minutes	Less than 5 minutes
2 TB	1	1 TB	2	No	No I/O	10 G	260 minutes	Less than 5 minutes
2 TB	1	512 GB	4	No	No I/O	10 G	215 minutes	Less than 5 minutes
2 TB	1	256 GB	8	No	No I/O	10 G	225 minutes	Less than 5 minutes
2 TB	20	100 GB	20	No	No I/O	10 G	210 minutes	Less than 5 minutes

AWS TO AHV AND AWS TO NUTANIX CLOUD CLUSTERS (NC2) ON AWS

You can prepare and migrate AWS EC2 instances to Nutanix Cloud Clusters (NC2) on AWS and AWS EC2 instances to AHV by using Move.

Note:

- For migration from any source (ESXi, Hyper-V, and AWS) to AHV target and from any source (ESXi, Hyper-V, and AWS) to NC2 on AWS target, Move should be deployed on the same destination target cluster where the VMs need to be migrated.
- For NC2 on AWS, static IP retention is not enabled by default.

Migration Considerations

You must consider the supported guest operating systems, requirements, recommendations, unsupported features, and limitations provided in this section before starting the migration process.

Supported Guest Operating Systems

Move supports some common operating systems. Unless otherwise specified, Nutanix has qualified the following 64-bit guest operating system versions.

For more information about the supported guest operating systems on AHV, refer to Compatibility and Interoperability Matrix. It also indicates whether an operating system is community-supported, legacy, or deprecated on AHV.

Fully Supported

Note: Some of the operating systems do not support AWS System Manager and thereby cannot used for Automatic preparation. These are marked for Manual preparation only in the following list:

- Windows 7, 8, 8.1, 10 (Manual preparation only)
- Windows Server 2008 R2, 2012, 2012 R2, 2016, 2019
- RHEL 6.3 (32-bit and 64-bit supported) (Manual preparation only), 6.5-6.10, 7.0–7.7, and 8.1

  Note: RHEL 6.3 is supported only with IDE as the disk controller.
- CentOS 7.8, 7.9
- CentOS 6.3 (32-bit and 64-bit supported) to 6.9, 7.0–8.2

  Note: CentOS 6.3 is supported only with IDE as the disk controller.
- Ubuntu Server and Desktop 12.04.5, 14.04.x, 16.04.x, 16.10 (32-bit and 64-bit supported)
- Ubuntu Server 12.0.4, 18.04, 19.04 (Manual preparation only)
- SUSE 11 SP3/SP4 (Manual preparation only), SUSE 12, SUSE 12 SP1 / SP2 / SP3 / SP4

- Oracle Linux 6.4 and later (Manual preparation only), 7.5-7.9, 8.8

  Note: If you face kernel panic issue on Oracle Linux versions after migration to AHV, then refer and apply the KB article 000004604 for these Oracle Linux VMs.

- Debian 9.4

Requirements

Before attempting to migrate VMs running on AWS using Move, make sure to conform to the requirements listed here.

General Requirements

Ensure to conform to the following requirements for AWS EC2 instances to NC2 on AWS and AWS EC2 instances to AHV migration.

- Supported browser: Google Chrome
- Ensure you have PowerShell version 4.0 or later.
- Ensure that you enable all outbound ports in the source VM.
- The VMs that are being migrated must be able to connect to the public S3 buckets.
- Ensure TCP 443 connection to AWS endpoint for operations in AWS.
- For Windows source VMs, ensure to disable UAC for Windows administrator user.
- For Linux source VMs, ensure that the VM has Internet connectivity during initial preparation to download the required packages. Install the following packages along with their dependencies: *wget*, *curl*, *jq*, *bash*, *sudo*.
- For automated guest preparation, make sure that the guest VM has enabled and is managed by the AWS SSM Agent.
- Network Security Services (NSS) 3.44 is required for Linux VMs.
- Before you initiate a migration, ensure to disable backups.
- The AWS account provided while adding an AWS source must have the set of permissions as provided in the following JSON to do end-to-end migration of an EC2 instance.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
```

```
        "ebs:GetSnapshotBlock"
      ],
      "Resource": "*"
    }
  ]
}
```

- The AWS account provided while adding AWS as both source and target must have the set of permissions as provided in the following JSON to do end-to-end migration of an EC2 instance.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
        "ebs:GetSnapshotBlock",
        "ec2:*KeyPair*",
        "ec2:RunInstances",
        "ec2:TerminateInstances",
        "ec2:CreateVolume",
        "ec2:DeleteVolume",
        "ec2:AttachVolume",
        "ec2:DetachVolume",
        "ec2:ModifyInstanceAttribute",
        "ebs:StartSnapshot",
        "ebs:PutSnapshotBlock",
        "ebs:CompleteSnapshot",
        "ec2:RegisterImage",
        "ec2:DeregisterImage"
      ],
      "Resource": "*"
    }
  ]
}
```

- Move should have all the permissions listed above for the following regions.

  Note: There can be restrictions for certain regions. For more information, refer to the next point.

```
"af-south-1"     // Africa (Cape Town).
"ap-east-1"      // Asia Pacific (Hong Kong).
"ap-northeast-1" // Asia Pacific (Tokyo).
"ap-northeast-2" // Asia Pacific (Seoul).
"ap-northeast-3" // Asia Pacific (Osaka).
"ap-south-1"     // Asia Pacific (Mumbai).
"ap-south-2"     // Asia Pacific (Hyderabad).
```

```
"ap-southeast-1" // Asia Pacific (Singapore).
"ap-southeast-2" // Asia Pacific (Sydney).
"ap-southeast-3" // Asia Pacific (Jakarta).
"ap-southeast-4" // Asia Pacific (Melbourne).
"ca-central-1"   // Canada (Central).
"eu-central-1"   // Europe (Frankfurt).
"eu-central-2"   // Europe (Zurich).
"eu-north-1"     // Europe (Stockholm).
"eu-south-1"     // Europe (Milan).
"eu-south-2"     // Europe (Spain).
"eu-west-1"      // Europe (Ireland).
"eu-west-2"      // Europe (London).
"eu-west-3"      // Europe (Paris).
"me-central-1"   // Middle East (UAE).
"me-south-1"     // Middle East (Bahrain).
"sa-east-1"      // South America (Sao Paulo).
"us-east-1"      // US East (N. Virginia).
"us-east-2"      // US East (Ohio).
"us-west-1"      // US West (N. California).
"us-west-2"      // US West (Oregon).
```

- Policies needed for explicit deny by region:

  Move should always have the following permissions for the region `us-east-1`:

  - `iam:GetUser`
  - `iam:SimulatePrincipalPolicy`

  Examples of the JSON are provided below.

  When restricting access to specific regions, the AWS account provided while adding an AWS source must have the set of permissions as provided in the following JSON to do end-to-end migration.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
        "ebs:GetSnapshotBlock"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "<AWS-REGION-1>",
            "<AWS-REGION-2>",
            ...
          ]
```

```
        }
      }
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": [
        "iam:GetUser",
        "iam:SimulatePrincipalPolicy"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "us-east-1",
            "<AWS-REGION-1>",
            "<AWS-REGION-2>",
            ...
          ]
        }
      }
    }
  ]
}
```

When restricting access to specific regions, the AWS account provided while adding AWS as source and target must have the set of permissions as provided in the following JSON to do end-to-end migration.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
        "ebs:GetSnapshotBlock",
        "ec2:*KeyPair*",
        "ec2:RunInstances",
        "ec2:TerminateInstances",
        "ec2:CreateVolume",
        "ec2:DeleteVolume",
        "ec2:AttachVolume",
        "ec2:DetachVolume",
        "ec2:ModifyInstanceAttribute",
        "ebs:StartSnapshot",
        "ebs:PutSnapshotBlock",
        "ebs:CompleteSnapshot",
        "ec2:RegisterImage",
        "ec2:DeregisterImage"
```

```
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "<AWS-REGION-1>",
            "<AWS-REGION-2>",
            ...
          ]
        }
      }
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": [
        "iam:GetUser",
        "iam:SimulatePrincipalPolicy"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "us-east-1",
            "<AWS-REGION-1>",
            "<AWS-REGION-2>",
            ...
          ]
        }
      }
    }
  ]
}
```

- Policies needed for explicit deny by IP address/CIDR:

  When restricting access using IP address/CIDR, the AWS account provided while adding an AWS source must have the set of permissions as provided in the following JSON to do end-to-end migration.

  This following JSON is an example. Update the JSON as necessary based on the security policies in your organization.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AwsSrcPermsWithRestrictedIP",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
```

```
        "ebs:GetSnapshotBlock"
      ],
      "Resource": "*",
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": [
            "<CIDR-BLOCK-1>",
            "<CIDR-BLOCK-2>",
            ...
          ]
        },
        "Bool": {
          "aws:ViaAWSService": "false"
        }
      }
    }
  ]
}
```

Note:

- A CIDR block is a group of IP addresses that share the same network prefix and have the same number of bits.

  Example: `192.168.x.x/29`

- In the above JSON, replace `<CIDR-BLOCK-x>` with the appropriate CIDR block.

- No explicit deny policies should be defined for the account.

  Explicit deny policies restrict access to AWS based on various parameters such as source IP address, VPC, VPC endpoint, and so on. For example, explicit deny for source IP address denies access to AWS when a request comes from an IP address outside the specified range.

Note: If the EC2 instances have disks with Amazon EBS encryption, then the AWS account may need additional permissions. For more information regarding permissions, refer to *AWS documentation* on Amazon EBS encryption.

Prerequisites for Linux guest VMs:

- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

Service Accounts

For successful migration of VMs from AWS EC2 instances to NC2 on AWS and AWS EC2 instances to AHV, Move requires the following.

- Prism Element UI for the AHV cluster
- An administrator for Windows source VMs or a root for Linux source VMs to run the source VM preparation scripts.

Qualified Metrics

The section lists the qualified metrics for migration from AWS EC2 instances to NC2 on AWS and AWS EC2 instances to AHV.

Qualified Metrics for Migration from AWS

The following metrics are qualified for migration from AWS.

- Migration of VMs using all EC2 instance type.

  Note: You cannot migrate EC2 instances with more than 40 disks.

- Migration of large VMs, VMs with active workloads, and combination of Windows and Linux VMs with multiple disks
- Migration of VMs from 13 different regions in parallel using single Nutanix-Move
- Migration plan with five instances each having 13 disks in parallel
- Single migration having two instances with 500 GB disk
- Migration of more than 10 VMs in a plan
- Virtual Private Cloud (VPC) Endpoint for S3 is required for preparation of instances without internet access.

  For more information about VPC endpoint for S3, refer to VPC Endpoint for Amazon S3.

  Also, to use automatic preparation, these instances needs to be managed by Amazon Web Services Systems Manager (AWS SSM) using VPC endpoint for Systems Manager.

  For more information about VPC endpoint for Systems Manager, refer to VPC endpoint for Systems Manager.

  These instances can only be migrated from the following regions where Move has S3 buckets deployed.

  - af-south-1
  - ap-east-1
  - ap-northeast-1
  - ap-northeast-2
  - ap-south-1
  - ap-southeast-1
  - ap-southeast-2
  - ca-central-1
  - eu-central-1
  - eu-north-1
  - eu-south-1
  - eu-west-1
  - eu-west-2
  - eu-west-3
  - me-south-1
  - sa-east-1
  - us-east-1
  - us-east-2
  - us-west-1
  - us-west-2

- Time required to create the first snapshot for large disks is longer. Depending on the size of the volumes, it can take several hours (up to 24 hours) for the AMI-creation process to complete.

Unsupported Features

This section lists the unsupported features for migration from AWS EC2 instances to NC2 on AWS and AWS EC2 instances to AHV.

- Guest operating systems other than the supported operating systems.

  For more information, refer to Supported Guest Operating Systems for AWS Migration

- Migration of VMs with EFS backed storage.

  > **Note:** Move currently supports migration of VMs with EBS backed storage.

- IP address and MAC address retention.
- Certain types of spot instances (created using non-persistent spot requests) cannot be stopped. When you perform a cutover for a spot instance, the instance does not stop and the cutover snapshot includes all the changes up to the point when the cutover snapshot is taken. Any IOs after this point are not available on the target machine.

Limitations

The section lists the limitations and qualified metrics for migration from AWS.

AWS EC2 instances to NC2 on AWS and AWS EC2 instances to AHV

The support for migration is constrained by the following.

- Migrations of VMs with Instance Store Volumes are supported. Instance Store Disks are not migrated in this migration.
- Time taken for migration depends on the size of the VM, data churn rate within the VM during migration, and Internet connectivity between AWS and on-prem data center.
- AWS instances with CentOS 6.8 and 6.9 take longer time to get the IP address on the target AHV cluster after migration.

Warnings and Cautions

- Source VM disks attached to mix of PVSCSI and LSI adapters might get different device names (`sda`, `sdb`, and so on).

  > **Note:** For Linux VMs, manually edit `fstab`. Then, arrange the correct order for the UUIDs.

Adding an AWS Environment

While creating a migration plan, if you need to add an AWS source or target, you have to add at least one AWS environment for migration.

About this task

> **Note:** This procedure is only applicable for migration to and from an AWS environment.

To add an AWS environment, do the following:

Procedure

1. Log on to Move UI.

2. Click **+ Add Environment** under Environments.
   The **Add Environment** appears.

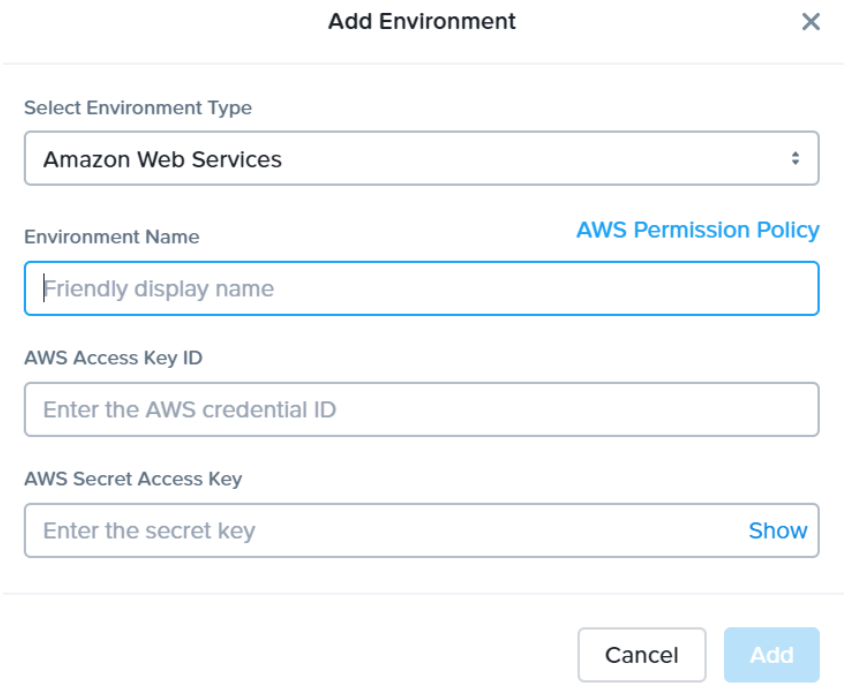

Figure 22: Add AWS Environment Dialog Box

3. Select **Amazon Web Services** as the source environment type.
4. Complete the indicated fields and click **Add**.
   a. **Source Name**: Enter a name for the AWS environment.
   b. **AWS Access Key ID**: Enter the access key ID of the AWS account.
   c. **AWS Secret Access Key**: Enter the key for logging on to AWS account.

   The source environment is added to Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

You can add a Nutanix AOS cluster environment. For more information, refer to Adding a Nutanix AOS Cluster Environment on page 39

Adding a Nutanix AOS Cluster Environment

While creating a migration plan, AHV AOS cluster can be added as both source or target. If you want to use ESXi on Nutanix cluster as target, then add the corresponding AOS cluster environment for ESXi.

About this task

Note: When you add a AOS cluster, Move VM IP address with the subnet 255.255.255.255 is added to the global NFS allowlist.

> **Caution:** Modifying the NFS allowlist of the destination container disables inheritance of the NFS allowlist at the global level and lead to issues with Move operations.

To add a Nutanix AOS cluster environment, do the following:

Procedure

1. Log on to the Move UI.
2. Click **+ Add Environment** under **Environments**.

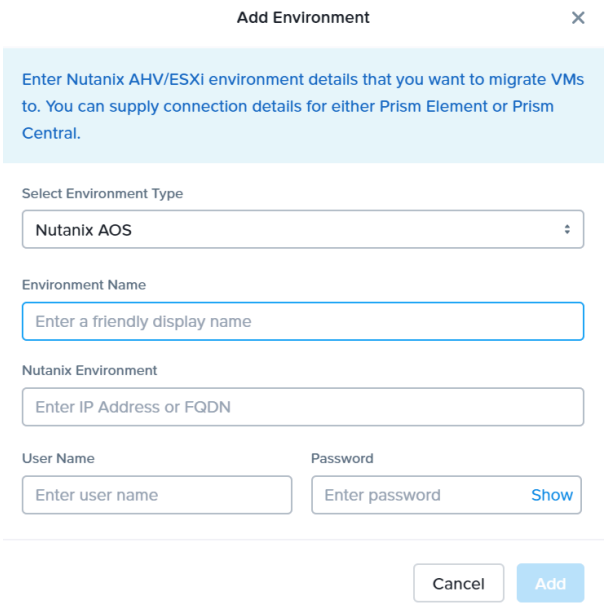

Figure 23: Add AOS Environment Dialog Box

   The **Add Environment** window appears.

3. Select **Nutanix AOS** as the target environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the NC2 on AWS and AHV or VMware ESXi on Nutanix environment.
   b. **Nutanix Environment**: Enter the IP address or the FQDN for the target Prism Central or Prism Element.
   c. **User Name**: Enter the username for logging on to the target Nutanix environment.
   d. **Password**: Enter the password for logging on to the target Nutanix environment.

5. (Only for ESXi to ESXi migration) Enter credentials for registered vCenter.

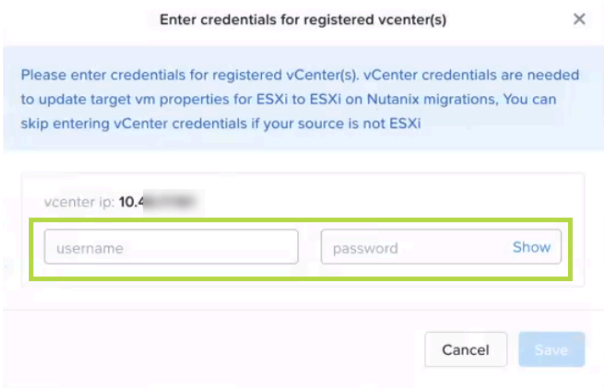

Figure 24: vCenter Credentials Dialog box

vCenter credentials are required to update the target VM properties for ESXi to ESXi on Nutanix migration.

Note:

- You can skip adding vCenter credentials if your source is not ESXi.
- To retain the VM properties on the target VM after migration, be sure to provide the target vCenter credentials. The following properties will be retained:
  - SCSI controller types
  - Network adaptor type
  - MAC address
  - Video card
  - Memory overcommit variables
  - Tool upgrade flag
  - Sync time with host flag
  - Disable acceleration flag
  - Enable logging flag

The AOS environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

Once you have added the environments, you can proceed to create the migration plan. For more information, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan (Hyper-V to ESXi) on page 103 or Creating a Migration Plan (Azure to AHV) on page 171 or Creating a Migration Plan (Azure to ESXi) on page 184

Creating a Migration Plan

You can create a migration plan to seed the data, cutover, and monitor the VMs. You can create the migration plan in Move without initiating the cutover process.

About this task

Note:

- This procedure is only applicable for migration from AWS EC2 instances to NC2 on AWS and AWS EC2 instances to AHV.
- In case of upgrade if the AWS provider is already present then you need to update the required AWS permissions. You can get these permissions from the Move User Guide and update it on AWS. Alternatively, you can also remove the AWS provider and try adding it back again. The attempt to add the AWS provider again will fail due to insufficient permissions, however, the list of new permissions is provided for you to update on AWS.
- For ongoing migrations from AWS to AOS, ensure the migration is in completed state or complete the migration before performing the upgrade.
- For AWS migrations, Move relies on the power state to identify if a VM is migratable. So, it is recommended that the VMs are started to check the correct status of migration.
- On your first logon, you can log on to Move with your defaults credentials.
- If you cancel or discard the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.

  For information on canceling an ongoing migration, see Pausing or Canceling a VM Migration on page 257.

  For information on performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

To create a migration plan, do the following:

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.
3. On the Move dashboard, click **Create a Migration Plan**.

   Note: If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.
4. Enter the new migration plan name, and then click **OK**.

   Note: You can edit the migration plan name by clicking the pencil icon next to the migration plan name.

   The **New Migration Plan** window appears.

   Note: You might at times see a message relating to inventory collection as shown below. During this time, the environments undergoing Refresh will not be available for selection. Such environments do not automatically

show up for selection once the inventory collection is over. In case you wish to select one of the environments undergoing Refresh (not available for selection), you will need to select **Cancel** and wait for inventory collection to get over and then create the migration plan again.

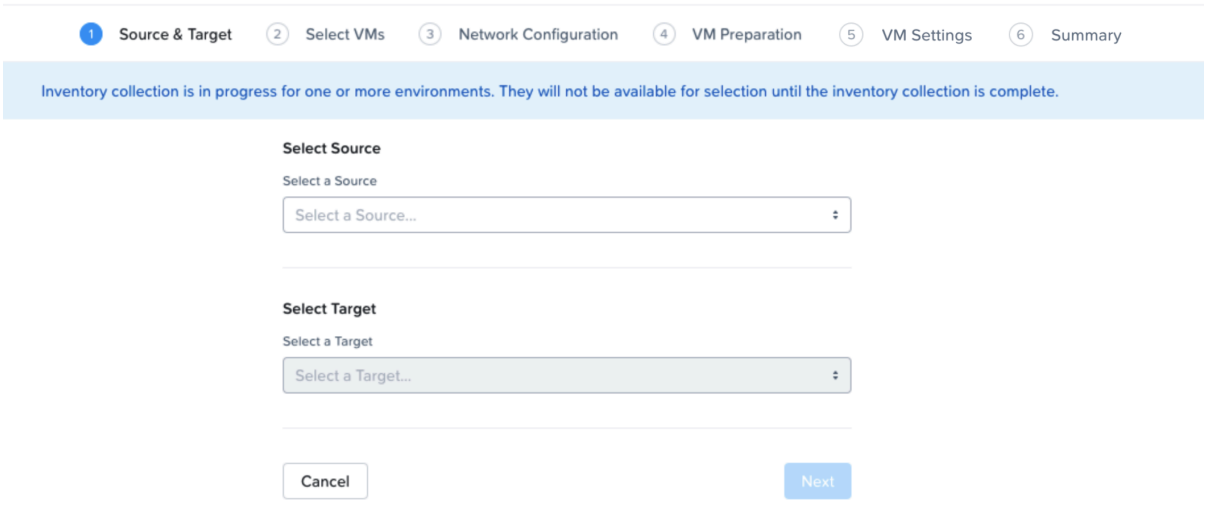

Figure 25: Inventory Collection Message

5. Complete the following fields, and then click **Next**.

   a. **Select a Source**: Select any AWS source for migration.

   b. **Region**: Select a region from where you will migrate the VMs.

   c. **Select a Target**: Select the target for migrating the VMs.

   d. **Target Project** (optional): Select the project you want as the target.
   This field is available only with Prism Central and when AHV is the target.

   e. **Target Owners**: Select the owners for the selected target project.
   This field is available only when a target project is selected.

   f. **Target Cluster**: Select the cluster on which you will migrate the VMs.

   g. **Target Container**: Select the container on which you will migrate the VMs.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click **+Add All**, and then click **Next**.

   Note: You cannot add more than 25 VMs in a single migration plan.

   You can filter the VMs by name in the **Filter** bar. You can also sort the VM types by selecting the appropriate column.

   Note: The VMs for which migration has failed are not displayed. To show the entire list of VMs, select **All VMs** from the drop-down list. To show the configuration of VMs, select **Configuration** from the drop-down list. A

question mark icon appears beside an unavailable VM that displays more information about that VM and indicate the reason of a failed VM migration.

Note: Migrated VMs retain their power state on the destination cluster.

The selected VMs are displayed in the sidebar.

7. In the **Network Configuration** screen, select the target network from the drop-down.

   - [Applicable only if Prism Central is used] Move provides the option to select VPC-based or VLAN-based target subnets. Based on the selection of a VPC or VLAN ID as the target network, the respective subnets are listed in the target subnet drop-down menu. Select the required subnet from the drop-down menu.

     Note: Overlay subnets which do not have IP address pool(s) associated will be disabled in the subnet drop-down menu.

   For performing test migration, refer to Creating a Test Capable VM Migration Plan on page 248 section.

   Click **Next**.

8. In the **VM Preparation** screen, select one of the following VM preparation modes.

   » **Automatic**. Select this option to automatically run scripts on the source VMs and prepare them for migration. Ensure that guest VM instance has enabled and is managed by the AWS SSM Agent.

   Note: For automatic Windows VM preparation, ensure the guest VM instance is managed by the AWS SSM Manager and the IAM profile on the instance should have the permissions mentioned in the following json. Alternatively, you can use the IAM profile provided by AWS Name "AmazonEC2RoleforSSM" and compare that with these permissions:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ssm:DescribeAssociation",
                "ssm:GetDeployablePatchSnapshotForInstance",
                "ssm:GetDocument",
                "ssm:DescribeDocument",
                "ssm:GetManifest",
                "ssm:GetParameters",
                "ssm:ListAssociations",
                "ssm:ListInstanceAssociations",
                "ssm:PutInventory",
                "ssm:PutComplianceItems",
                "ssm:PutConfigurePackageResult",
                "ssm:UpdateAssociationStatus",
                "ssm:UpdateInstanceAssociationStatus",
                "ssm:UpdateInstanceInformation"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "ssmmessages:CreateControlChannel",
                "ssmmessages:CreateDataChannel",
                "ssmmessages:OpenControlChannel",
                "ssmmessages:OpenDataChannel"
            ],
            "Resource": "*"
```

```
        },
        {
            "Effect": "Allow",
            "Action": [
                "ec2messages:AcknowledgeMessage",
                "ec2messages:DeleteMessage",
                "ec2messages:FailMessage",
                "ec2messages:GetEndpoint",
                "ec2messages:GetMessages",
                "ec2messages:SendReply"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "cloudwatch:PutMetricData"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeInstanceStatus"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "ds:CreateComputer",
                "ds:DescribeDirectories"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "logs:CreateLogGroup",
                "logs:CreateLogStream",
                "logs:DescribeLogGroups",
                "logs:DescribeLogStreams",
                "logs:PutLogEvents"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "s3:GetBucketLocation",
                "s3:PutObject",
                "s3:GetObject",
                "s3:GetEncryptionConfiguration",
                "s3:AbortMultipartUpload",
                "s3:ListMultipartUploadParts",
                "s3:ListBucket",
                "s3:ListBucketMultipartUploads"
            ],
            "Resource": "*"
        }
    ]
```

```
}
```

For more information on AWS SSM Agent, refer to Working with SSM Agent.

» **Manual**. Move displays the VM preparation scripts for Windows and Linux VMs. To manually download and run the migration preparation software, select **Manual**, and then run the scripts provided in the **VM Preparation** screen on the respective source EC2 instance.

These scripts prepare the instance by performing the NTNX VirtIO driver installation.

» **Mixed**. Move allows you to select the VM preparation mode for each VM. This setting allows you to manually prepare one set of VMs and automatically prepare another set of VMs in the same migration plan. If you want to have such a hybrid combination of **Automatic** and **Manual** VM preparation modes, proceed to the step 8.

> **Note:** The preparation mode automatically switches to **Mixed** if you change the mode of preparation for a set of VMs under **Change Settings** and have a mix of **Automatic** and **Manual** VM preparation modes. You cannot manually select the **Mixed** option from the **Preparation Mode** drop-down list.

9. In the **Override individual VM Preparation** section, click **Change Settings** to override settings for the individual VMs. You can edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

> **Note:** If you change the **Mode of Preparation** to **Manual** for a VM, then copy the new generated scripts of that specific VM and run them on the source VM.
>
> Access the newly generated VM preparation script for that VM by selecting the Copy Scripts icon under the **Actions** column.

After making the required changes, click **Done**.

Then, click **Next**.

10. In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

    a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

    b. **Timezone**: Set the timezone as the hardware clock of the VMs in target.

    If set to default, Move configures UTC timezones for Linux VMs and cluster timezones for Windows VMs.

    c. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

    Only those categories which have values are available for selection.

    d. **Skip CDROM addition on target VMs**: Select this check box to skip the CDROM addition on the target VMs.

    e. **Retain User Data**: Select this check box to retain the AWS user-data scripts on the target VM.

    Note:

    - AWS user-data is a set of commands or data that can be provided to an AWS EC2 instance at the time of launch.
    - To retain the user-data on AWS instances on the target environment, the size of the user-data should not exceed 16KiB.

    f. **Enable Memory Overcommit**: Select this option to enable memory overcommit on the target VM.

    For more information on memory overcommit deployment, refer to AHV Administration Guide.

    g. **Settings for individual VMs**: Click **Change Settings** to configure settings such as timezone, VM priority, retain user-data scripts, and skip CDROM addition for individual VMs. You can also search the VM by typing the name of the VM and change the settings.

    h. **Schedule Data Seeding**: Check this check box to select the date and time for migration.

11. In the **Summary** screen, choose one of the following, and then proceed review the VM migration summary.

    » **Back**: To edit the information, click this option.

    » **Save**: To save the migration plan, click this option.

    For more information about how to start the migration later, and check the migrated VM status and details, refer to Environments and Migration Plan Management on page 259.

    » **Save and Start**: Click this option to save the migration plan and begin the migration immediately

    The seeding process for migration begins. You can monitor this information by selecting **Status** for the migration plan.

    Note: The seeding process can take several minutes depending on the number of VMs.

What to do next

If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover on page 132

Performing a Migration Cutover

When the migration plan is started and the seeding process is complete, you can cut over the selected VMs in the AHV cluster. You can monitor the VM migration progress by clicking the **Status** link.

About this task

Note:

- You can perform this operation only when the VM status is Ready to Cutover.
- Recommended that cutover should be performed within one week of initial data seeding.
- If the initial data seeding finishes in less than 10 minutes, Move continues to wait for 10 minutes to take the incremental snapshot; however, you can trigger the cutover immediately.
- The cutover process increments in absolute numbers.

To perform a migration cutover, do the following:

Procedure

1. In the Move UI, click the **Ready to Cutover** status to display the list of available VMs.
2. To perform the cutover, select the VMs or group of VMs.
3. Click **Cutover**.
   The cutover process performs the following VM actions.

   - Shuts down the instance
   - Takes the final snapshots for the instance and copying the final changes to AHV
   - Creates a VM in the target AHV cluster
   - Attaches the replicated drives to the VM
   - Powers on or off the VM (depends on the initial power state)

   The cutover process begins immediately and might take a few minutes. Once cutover is complete, the VM is ready for use in the new AHV cluster.

What to do next

You can manage the migration plan once the migration plan is ready. For more information, refer to Environments and Migration Plan Management on page 259

Performance Matrix for Large Data Migration

Move performs end-to-end migration of large VMs. The scenarios are tested based on the following parameters. The following tables show the performance numbers from the Move lab.

Table 15: Performance Numbers of Large Data Migration (AWS to AHV)

Total migration size	Data churn	Region used	Target network bandwidth	First snapshot creation time	Migration time taken		Platform
					Data seeding	Cutover	

Total migration size	Data churn	Region used	Target network bandwidth	First snapshot creation time	Migration time taken		Platform
150 G	No	us-west-2 to US onprem Cluster	10 G	NA	52 minutes	Less than 4 minutes	NX-1065
1 TB	No	us-west-2 to US onprem Cluster	10 G	6 hours	3 hours 50 minutes	Less than 9 minutes	NX-1065
1 TB with 3 vDisks	5 GB	us-west-1 to US onprem Cluster	10 G	2 hours 28 minutes	1 hour 38 minutes	Less than 24 minutes	NX-1065

AWS TO ESXI

You can prepare and migrate AWS EC2 instances to ESXi by using Move.

> **Note:** ESXi as a target refers to ESXi running on Nutanix appliances.

Migration Considerations

You must consider the supported guest operating systems, requirements, recommendations, unsupported features, and limitations provided in this section before starting the migration process.

Supported Guest Operating Systems (AWS to ESXi)

Move supports some common operating systems. Unless otherwise specified, Nutanix has qualified the following 64-bit guest operating system versions.

Fully Supported

> **Note:** Some of the operating systems do not support AWS System Manager and thereby cannot used for Automatic preparation. These are marked for Manual preparation only in the following list:

- Windows 7, 8, 8.1, 10 (Manual preparation only)
- Windows Server 2012, 2012 R2, 2016, 2019
- RHEL 6.3 (32-bit and 64-bit supported) (Manual preparation only), 6.5-6.10, 7.0–7.7, and 8.1

  > **Note:** RHEL 6.3 is supported only with IDE as the disk controller.

- CentOS 6.3 (32-bit and 64-bit supported) to 6.9, 7.0–7.7, 8.0-8.2

  > **Note:** CentOS 6.3 is supported only with IDE as the disk controller.

- Ubuntu Server and Desktop 12.04.5, 14.04.x, 16.04.x, 16.10 (32-bit and 64-bit supported)
- Ubuntu Server 12.0.4, 18.04, 19.04 (Manual preparation only)

Requirements (AWS to ESXi)

Before attempting to migrate VMs running on AWS using Move, make sure to conform to the requirements listed here.

General Requirements

Ensure to conform to the following requirements for AWS to ESXi migration.

- Supported browser: Google Chrome.
- Ensure that you have PowerShell version 4.0 or later.
- Ensure that you enable all outbound ports in the source VM.
- The VMs that are being migrated must be able to connect to the public S3 buckets.
- For Windows source VMs, ensure to disable UAC for Windows administrator user.
- For Linux source VMs, ensure that the VM has Internet connectivity during initial preparation to download the required packages. Install the following packages along with their dependencies: *wget, curl, jq, bash, sudo*.

- For automated guest preparation, make sure that the guest VM has enabled and is managed by the AWS SSM Agent.
- Network Security Services (NSS) 3.44 is required for Linux VMs.
- Before you initiate a migration, ensure to disable backups.
- The AWS account provided while adding an AWS source must have the set of permissions as provided in the following JSON to do end-to-end migration of an EC2 instance.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
        "ebs:GetSnapshotBlock"
      ],
      "Resource": "*"
    }
  ]
}
```

- The AWS account provided while adding AWS as both source and target must have the set of permissions as provided in the following JSON to do end-to-end migration of an EC2 instance.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
```

```
        "ebs:GetSnapshotBlock",
        "ec2:*KeyPair*",
        "ec2:RunInstances",
        "ec2:TerminateInstances",
        "ec2:CreateSecurityGroup",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateVolume",
        "ec2:DeleteVolume",
        "ec2:AttachVolume",
        "ec2:DetachVolume",
        "ec2:ModifyInstanceAttribute",
        "ebs:StartSnapshot",
        "ebs:PutSnapshotBlock",
        "ebs:CompleteSnapshot",
        "ec2:RegisterImage",
        "ec2:DeregisterImage"
      ],
      "Resource": "*"
    }
  ]
}
```

- Move should have all the permissions listed above for the following regions.

  Note: There can be restrictions for certain regions. For more information, refer to the next point.

```
"af-south-1"     // Africa (Cape Town).
"ap-east-1"      // Asia Pacific (Hong Kong).
"ap-northeast-1" // Asia Pacific (Tokyo).
"ap-northeast-2" // Asia Pacific (Seoul).
"ap-northeast-3" // Asia Pacific (Osaka).
"ap-south-1"     // Asia Pacific (Mumbai).
"ap-south-2"     // Asia Pacific (Hyderabad).
"ap-southeast-1" // Asia Pacific (Singapore).
"ap-southeast-2" // Asia Pacific (Sydney).
"ap-southeast-3" // Asia Pacific (Jakarta).
"ap-southeast-4" // Asia Pacific (Melbourne).
"ca-central-1"   // Canada (Central).
"eu-central-1"   // Europe (Frankfurt).
"eu-central-2"   // Europe (Zurich).
"eu-north-1"     // Europe (Stockholm).
"eu-south-1"     // Europe (Milan).
"eu-south-2"     // Europe (Spain).
"eu-west-1"      // Europe (Ireland).
"eu-west-2"      // Europe (London).
"eu-west-3"      // Europe (Paris).
"me-central-1"   // Middle East (UAE).
"me-south-1"     // Middle East (Bahrain).
"sa-east-1"      // South America (Sao Paulo).
"us-east-1"      // US East (N. Virginia).
"us-east-2"      // US East (Ohio).
"us-west-1"      // US West (N. California).
"us-west-2"      // US West (Oregon).
```

- Policies needed for explicit deny by region:

  Move should always have the following permissions for the region `us-east-1`:

  - `iam:GetUser`
  - `iam:SimulatePrincipalPolicy`

  Examples of the JSON are provided below.

  When restricting access to specific regions, the AWS account provided while adding an AWS source must have the set of permissions as provided in the following JSON to do end-to-end migration.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
        "ebs:GetSnapshotBlock"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "<AWS-REGION-1>",
            "<AWS-REGION-2>",
            ...
          ]
        }
      }
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": [
        "iam:GetUser",
        "iam:SimulatePrincipalPolicy"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "us-east-1",
            "<AWS-REGION-1>",
            "<AWS-REGION-2>",
            ...
          ]
        }
      }
```

```
    }
  ]
}
```

When restricting access to specific regions, the AWS account provided while adding AWS as source and target must have the set of permissions as provided in the following JSON to do end-to-end migration.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
        "ebs:GetSnapshotBlock",
        "ec2:*KeyPair*",
        "ec2:RunInstances",
        "ec2:TerminateInstances",
        "ec2:CreateSecurityGroup",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateVolume",
        "ec2:DeleteVolume",
        "ec2:AttachVolume",
        "ec2:DetachVolume",
        "ec2:ModifyInstanceAttribute",
        "ebs:StartSnapshot",
        "ebs:PutSnapshotBlock",
        "ebs:CompleteSnapshot",
        "ec2:RegisterImage",
        "ec2:DeregisterImage"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "<AWS-REGION-1>",
            "<AWS-REGION-2>",
            ...
          ]
        }
      }
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": [
        "iam:GetUser",
        "iam:SimulatePrincipalPolicy"
      ],
      "Resource": "*",
```

```
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "us-east-1",
            "<AWS-REGION-1>",
            "<AWS-REGION-2>",
            ...
          ]
        }
      }
    }
  ]
}
```

- Policies needed for explicit deny by IP address/CIDR:

  When restricting access using IP address/CIDR, the AWS account provided while adding an AWS source must have the set of permissions as provided in the following JSON to do end-to-end migration.

  This following JSON is an example. Update the JSON as necessary based on the security policies in your organization.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AwsSrcPermsWithRestrictedIP",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
        "ebs:GetSnapshotBlock"
      ],
      "Resource": "*",
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": [
            "<CIDR-BLOCK-1>",
            "<CIDR-BLOCK-2>",
            ...
          ]
        },
        "Bool": {
          "aws:ViaAWSService": "false"
        }
      }
    }
  ]
```

```
}
```

Note:

- A CIDR block is a group of IP addresses that share the same network prefix and have the same number of bits.

  Example: `192.168.x.x/29`

- In the above JSON, replace `<CIDR-BLOCK-x>` with the appropriate CIDR block.

- No explicit deny policies should be defined for the account.

  Explicit deny policies restrict access to AWS based on various parameters such as source IP address, VPC, VPC endpoint, and so on. For example, explicit deny for source IP address denies access to AWS when a request comes from an IP address outside the specified range.

Note: If the EC2 instances have disks with Amazon EBS encryption, then the AWS account may need additional permissions. For more information regarding permissions, refer to *AWS documentation* on Amazon EBS encryption.

Prerequisites for Linux guest VMs:

- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

Service Accounts

For successful migration of VMs from AWS to ESXi, Move requires the following.

- Prism Web Console UI for the ESXi cluster
- An administrator for Windows source VMs or a root for Linux source VMs to run the source VM preparation scripts.

Qualified Metrics (AWS to ESXi)

The section lists the qualified metrics for migration from AWS.

The following metrics are qualified for migration from AWS.

- Migration of VMs using all EC2 instance type.

Note: You cannot migrate EC2 instances with more than 40 disks.

- Migration of large VMs, VMs with active workloads, and combination of Windows and Linux VMs with multiple disks
- Migration of VMs from 13 different regions in parallel using single Nutanix-Move
- Migration plan with five instances each having 13 disks in parallel
- Single migration having two instances with 500 GB disk
- Migration of more than 10 VMs in a plan

- Virtual Private Cloud (VPC) Endpoint for S3 is required for preparation of instances without internet access.

  For more information about VPC endpoint for S3, refer to VPC Endpoint for Amazon S3.

  Also, to use automatic preparation, these instances needs to be managed by Amazon Web Services Systems Manager (AWS SSM) using VPC endpoint for Systems Manager.

  For more information about VPC endpoint for Systems Manager, refer to VPC endpoint for Systems Manager.

  These instances can only be migrated from the following regions where Move has S3 buckets deployed.

  - af-south-1
  - ap-east-1
  - ap-northeast-1
  - ap-northeast-2
  - ap-south-1
  - ap-southeast-1
  - ap-southeast-2
  - ca-central-1
  - eu-central-1
  - eu-north-1
  - eu-south-1
  - eu-west-1
  - eu-west-2
  - eu-west-3
  - me-south-1
  - sa-east-1
  - us-east-1
  - us-east-2
  - us-west-1
  - us-west-2
- Time required to create the first snapshot for large disks is longer. Depending on the size of the volumes, it can take several hours (up to 24 hours) for the AMI-creation process to complete.

Unsupported Features (AWS to ESXi)

This section lists the unsupported features for migration from AWS.

- Guest operating systems other than the supported operating systems.

  For more information, refer to Supported Guest Operating Systems for AWS Migration

- Migration of VMs with EFS backed storage.

  > **Note:** Move currently supports migration of VMs with EBS backed storage.

- IP address and MAC address retention.

- Certain types of spot instances (created using non-persistent spot requests) cannot be stopped. When you perform a cutover for a spot instance, the instance does not stop and the cutover snapshot includes all the changes up to the point when the cutover snapshot is taken. Any IOs after this point are not available on the target machine.

Limitations (AWS to ESXi)

The section lists the limitations for migration from AWS.

AWS to ESXi Migration Limitations

The support for migration is constrained by the following.

- Migrations of VMs with Instance Store Volumes are supported. Instance Store Disks are not migrated in this migration.
- Time taken for migration depends on the size of the VM, data churn rate within the VM during migration, and Internet connectivity between AWS and on-prem data center.
- Move does not install VMware Tools on the VMs.
- AWS instances with CentOS 6.8 and 6.9 take longer time to get the IP address on the target AHV cluster after migration.

Warnings and Cautions

- Source VM disks attached to mix of PVSCSI and LSI adapters might get different device names (`sda`, `sdb`, and so on).

  Note: For Linux VMs, manually edit `fstab`. Then, arrange the correct order for the UUIDs.

Adding an AWS Environment

While creating a migration plan, if you need to add an AWS source or target, you have to add at least one AWS environment for migration.

About this task

Note: This procedure is only applicable for migration to and from an AWS environment.

To add an AWS environment, do the following:

Procedure

1. Log on to Move UI.

2. Click **+ Add Environment** under Environments.
   The **Add Environment** appears.

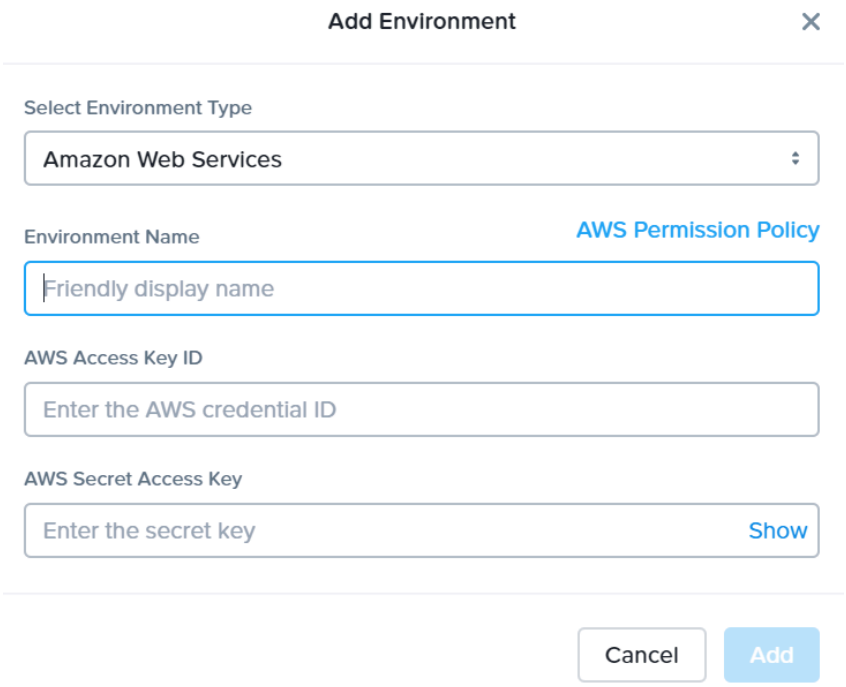

Figure 26: Add AWS Environment Dialog Box

3. Select **Amazon Web Services** as the source environment type.
4. Complete the indicated fields and click **Add**.
   a. **Source Name**: Enter a name for the AWS environment.
   b. **AWS Access Key ID**: Enter the access key ID of the AWS account.
   c. **AWS Secret Access Key**: Enter the key for logging on to AWS account.

   The source environment is added to Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

You can add a Nutanix AOS cluster environment. For more information, refer to Adding a Nutanix AOS Cluster Environment on page 39

Adding a Nutanix AOS Cluster Environment

While creating a migration plan, AHV AOS cluster can be added as both source or target. If you want to use ESXi on Nutanix cluster as target, then add the corresponding AOS cluster environment for ESXi.

About this task

Note: When you add a AOS cluster, Move VM IP address with the subnet 255.255.255.255 is added to the global NFS allowlist.

> **Caution:** Modifying the NFS allowlist of the destination container disables inheritance of the NFS allowlist at the global level and lead to issues with Move operations.

To add a Nutanix AOS cluster environment, do the following:

Procedure

1. Log on to the Move UI.
2. Click **+ Add Environment** under **Environments**.

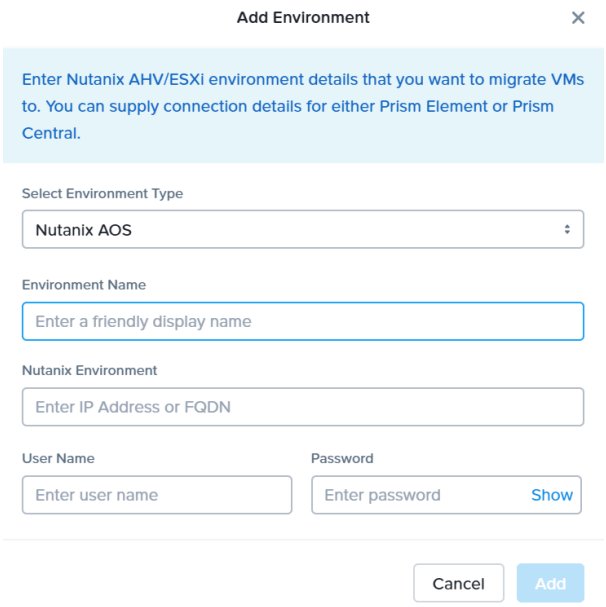

Figure 27: Add AOS Environment Dialog Box

   The **Add Environment** window appears.

3. Select **Nutanix AOS** as the target environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the NC2 on AWS and AHV or VMware ESXi on Nutanix environment.
   b. **Nutanix Environment**: Enter the IP address or the FQDN for the target Prism Central or Prism Element.
   c. **User Name**: Enter the username for logging on to the target Nutanix environment.
   d. **Password**: Enter the password for logging on to the target Nutanix environment.

5. (Only for ESXi to ESXi migration) Enter credentials for registered vCenter.

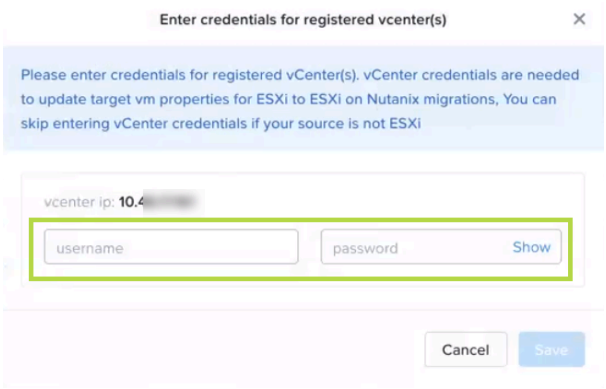

Figure 28: vCenter Credentials Dialog box

vCenter credentials are required to update the target VM properties for ESXi to ESXi on Nutanix migration.

Note:

- You can skip adding vCenter credentials if your source is not ESXi.
- To retain the VM properties on the target VM after migration, be sure to provide the target vCenter credentials. The following properties will be retained:
  - SCSI controller types
  - Network adaptor type
  - MAC address
  - Video card
  - Memory overcommit variables
  - Tool upgrade flag
  - Sync time with host flag
  - Disable acceleration flag
  - Enable logging flag

The AOS environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

Once you have added the environments, you can proceed to create the migration plan. For more information, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan (Hyper-V to ESXi) on page 103 or Creating a Migration Plan (Azure to AHV) on page 171 or Creating a Migration Plan (Azure to ESXi) on page 184

Creating a Migration Plan (AWS to ESXi)

You can create a migration plan to seed the data, cutover, and monitor the VMs. You can create the migration plan in Move without initiating the cutover process.

About this task

Note:

- This procedure is only applicable for migration from AWS to ESXi.
- In case of upgrade if the AWS provider is already present then you need to update the required AWS permissions. You can get these permissions from the Move User Guide and update it on AWS. Alternatively, you can also remove the AWS provider and try adding it back again. The attempt to add the AWS provider again will fail due to insufficient permissions, however, the list of new permissions is provided for you to update on AWS.
- For AWS migrations, Move relies on the power state to identify if a VM is migratable. So, it is recommended that the VMs are started to check the correct status of migration.
- For ongoing migrations from AWS to AOS, ensure the migration is in completed state or complete the migration before performing the upgrade.
- On your first logon, you can log on to Move with your defaults credentials.
- If you cancel or discard the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.

  For information on canceling an ongoing migration, see Pausing or Canceling a VM Migration on page 257.

  For information on performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

To create a migration plan, do the following:

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.
3. On the Move dashboard, click **Create a Migration Plan**.

   Note: If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.
4. Enter the new migration plan name, and then click **OK**.

   Note: You can edit the migration plan name by clicking the pencil icon next to the migration plan name.

   The **New Migration Plan** window appears.

   Note: You might at times see a message relating to inventory collection as shown below. During this time, the environments undergoing Refresh will not be available for selection. Such environments do not automatically show up for selection once the inventory collection is over. In case you wish to select one of the environments

undergoing Refresh (not available for selection), you will need to select **Cancel** and wait for inventory collection to get over and then create the migration plan again.

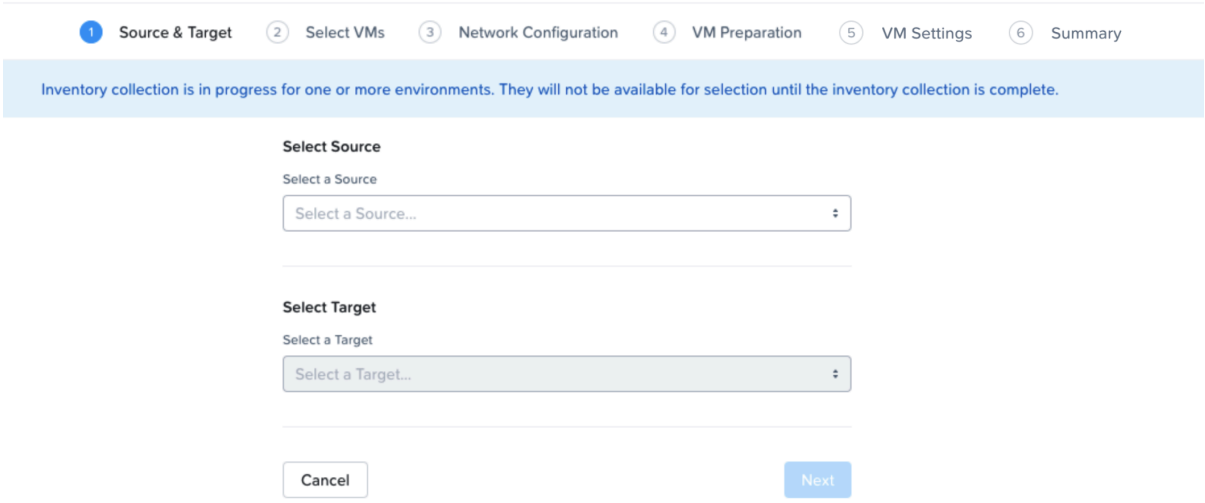

Figure 29: Inventory Collection Message

5. Complete the following fields, and then click **Next**.

   a. **Select a Source**: Select any AWS source for migration.

   b. **Region**: Select a region from where you will migrate the VMs.

   c. **Select a Target**: Select the target for migrating the VMs.

   d. **Target Cluster**: Select the cluster on which you will migrate the VMs.

   e. **Target Container**: Select the container on which you will migrate the VMs.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click **+Add All**, and then click **Next**.

   Note: You cannot add more than 25 VMs in a single migration plan.

   You can filter the VMs by name in the **Filter** bar. You can also sort the VM types by selecting the appropriate column.

   Note: The VMs for which migration has failed are not displayed. To show the entire list of VMs, select **All VMs** from the drop-down list. To show the configuration of VMs, select **Configuration** from the drop-down list. A question mark icon appears beside an unavailable VM that displays more information about that VM and indicate the reason of a failed VM migration.

   Note: Migrated VMs retain their power state on the destination cluster.

   The selected VMs are displayed in the sidebar.

7. In the **Network Configuration** screen, select the target network from the drop-down, and then click **Next**.

   For performing test migration, refer to Creating a Test Capable VM Migration Plan on page 248 section.

8. In the **VM Preparation** screen, select one of the following VM preparation modes.

   » **Automatic**. Select this option to automatically run scripts on the source VMs and prepare them for migration. Ensure that guest VM instance has enabled and is managed by the AWS SSM Agent.

   Note: For automatic Windows VM preparation, ensure the guest VM instance is managed by the AWS SSM Manager and the IAM profile on the instance should have the permissions mentioned in the following json. Alternatively, you can use the IAM profile provided by AWS Name "AmazonEC2RoleforSSM" and compare that with these permissions:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ssm:DescribeAssociation",
                "ssm:GetDeployablePatchSnapshotForInstance",
                "ssm:GetDocument",
                "ssm:DescribeDocument",
                "ssm:GetManifest",
                "ssm:GetParameters",
                "ssm:ListAssociations",
                "ssm:ListInstanceAssociations",
                "ssm:PutInventory",
                "ssm:PutComplianceItems",
                "ssm:PutConfigurePackageResult",
                "ssm:UpdateAssociationStatus",
                "ssm:UpdateInstanceAssociationStatus",
                "ssm:UpdateInstanceInformation"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "ssmmessages:CreateControlChannel",
                "ssmmessages:CreateDataChannel",
                "ssmmessages:OpenControlChannel",
                "ssmmessages:OpenDataChannel"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "ec2messages:AcknowledgeMessage",
                "ec2messages:DeleteMessage",
                "ec2messages:FailMessage",
                "ec2messages:GetEndpoint",
                "ec2messages:GetMessages",
                "ec2messages:SendReply"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "cloudwatch:PutMetricData"
            ],
            "Resource": "*"
        },
```

```
        {
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeInstanceStatus"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "ds:CreateComputer",
                "ds:DescribeDirectories"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "logs:CreateLogGroup",
                "logs:CreateLogStream",
                "logs:DescribeLogGroups",
                "logs:DescribeLogStreams",
                "logs:PutLogEvents"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "s3:GetBucketLocation",
                "s3:PutObject",
                "s3:GetObject",
                "s3:GetEncryptionConfiguration",
                "s3:AbortMultipartUpload",
                "s3:ListMultipartUploadParts",
                "s3:ListBucket",
                "s3:ListBucketMultipartUploads"
            ],
            "Resource": "*"
        }
    ]
}
```

For more information on AWS SSM Agent, refer to Working with SSM Agent.

» **Manual**. Move displays the VM preparation scripts for Windows and Linux VMs. Copy the scripts and manually run them on the source VMs. If you select this option, proceed to step 8.

» **Mixed**. Move allows you to select the VM preparation mode for each VM. This setting allows you to manually prepare one set of VMs and automatically prepare another set of VMs in the same migration plan. If you want to have such a hybrid combination of **Automatic** and **Manual** VM preparation modes, proceed to the step 8.

> **Note:** The preparation mode automatically switches to **Mixed** if you change the mode of preparation for a set of VMs under **Change Settings** and have a mix of **Automatic** and **Manual** VM preparation modes. You cannot manually select the **Custom** option from the **Preparation Mode** drop-down list.

9. In the **Override individual VM Preparation** section, click **Change Settings** to override settings for the individual VMs. You can edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

   > **Note:** If you change the **Mode of Preparation** to **Manual** for a VM, then copy the new generated scripts of that specific VM and run them on the source VM.
   >
   > Access the newly generated VM preparation script for that VM by selecting the Copy Scripts icon under the **Actions** column.

   After making the required changes, click **Done**.

   Then, click **Next**.

10. In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

    a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

    b. **Timezone**: Set the timezone as the hardware clock of the VMs in target.

       If set to default, Move configures UTC timezones for Linux VMs and cluster timezones for Windows VMs.

    c. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

       Only those categories which have values are available for selection.

    d. **Skip CDROM addition on target VMs**: Select this check box to skip the CDROM addition on the target VMs.

    e. **Settings for individual VMs**: Click **Change Settings** to configure settings such as timezone, VM priority, and skip CDROM addition for individual VMs. You can also search the VM by typing the name of the VM and change the settings.

    f. **Schedule Data Seeding**: Check this check box to select the date and time for migration.

11. In the **Summary** screen, choose one of the following, and then proceed to review the VM migration summary.

    » **Back**: To edit the information, click this option.

    » **Save**: To save the migration plan, click this option.

       For more information about how to start the migration later, and check the migrated VM status and details, refer to Environments and Migration Plan Management on page 259.

    » **Save and Start**: Click this option to save the migration plan and begin the migration immediately

    The seeding process for migration begins. You can monitor this information by selecting **Status** for the migration plan.

    > **Note:** The seeding process can take several minutes depending on the number of VMs.

What to do next

If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover (AWS to ESXi) on page 151

Performing a Migration Cutover (AWS to ESXi)

When the migration plan is started and the seeding process is complete, you can cut over the selected VMs in the AOS cluster. You can monitor the VM migration progress by clicking the **Status** link.

About this task

Note:

- You can perform this operation only when the VM status is Ready to Cutover.
- Recommended that cutover should be performed within one week of initial data seeding.
- If the initial data seeding finishes in less than 10 minutes, Move continues to wait for 10 minutes to take the incremental snapshot; however, you can trigger the cutover immediately.
- The cutover process increments in absolute numbers.

To perform a migration cutover, do the following:

Procedure

1. In the Move UI, click the **Ready to Cutover** status to display the list of available VMs.
2. To perform the cutover, select the VMs or group of VMs.
3. Click **Cutover**.
   The cutover process performs the following VM actions.

   - Shuts down the instance
   - Takes the final snapshots for the instance and copying the final changes to AOS
   - Creates a VM in the target AOS cluster
   - Attaches the replicated drives to the VM
   - Powers on or off the VM (depends on the initial power state)

   The cutover process begins immediately and might take a few minutes. Once cutover is complete, the VM is ready for use in the new AOS cluster.

What to do next

You can manage the migration plan once the migration plan is ready. For more information, refer to Environments and Migration Plan Management on page 259

AZURE TO AHV AND AZURE TO NUTANIX CLOUD CLUSTERS (NC2) ON AZURE

You can prepare and migrate Azure VMs to AHV and Azure to Nutanix Cloud Clusters (NC2) on Azure by using Move.

Migration Considerations

You must consider the supported guest operating systems, requirements, recommendations, unsupported features, and limitations provided in this section before starting the migration process.

Supported Guest Operating Systems

Move supports some common operating systems. Unless otherwise specified, Nutanix has qualified the following 64-bit guest operating system versions.

For more information about the supported guest operating systems on AHV, refer to Compatibility and Interoperability Matrix. It also indicates whether an operating system is community-supported, legacy, or deprecated on AHV.

Generation 1 VMs Support

- Windows 7 (TLS 1.2 and SHA2 should be enabled), 8, 8.1, 10
- Windows Server 2008 R2 (TLS 1.2 and SHA2 should be enabled), 2012, 2012 R2, 2016, 2019, 2022
- RHEL 6.3 (32-bit and 64-bit supported), 6.5-6.10, 7.0–7.7, and 8.0–8.2

  > **Note:** RHEL 6.3 is supported only with IDE as the disk controller.

- CentOS 6.3 (32-bit and 64-bit supported) to 6.9, 7.0–7.7, 8.0-8.2

  > **Note:** CentOS 6.3 is supported only with IDE as the disk controller.

- Ubuntu Server and Desktop 12.04.5, 14.04.x, 16.04.x, 16.10 (32-bit and 64-bit supported)
- Ubuntu Server 12.0.4, 18.04, 19.04

  > **Note:** For Ubuntu 12.0.4, SCSI is not supported. PCI disk controller is required on AOS target for this operating system version to boot correctly.

- SUSE Linux Enterprise Server 12 SP5, 15 SP1
- Oracle Linux 6.4–6.8, 7.5-8.0
- Debian 10

Generation 2 VMs Support (UEFI Enabled VMs)

- Windows 10 Pro, Windows 10 Enterprise
- Windows Server 2012, 2012 R2, 2016, 2019, 2022
- RHEL 7.0, 7.4–7.7, and 8.0–8.1
- CentOS 7.4–7.7, 8.0-8.1
- Ubuntu Server 16.04, 18.04, 19.04, 19.10

- SUSE Linux Enterprise Server 12 SP4, 15 SP1
- Oracle Linux 7.7, 7.7-ci

Requirements (Azure to AHV)

Before attempting to migrate VMs running on Azure using Move, make sure to conform to the requirements listed here.

General Requirements

Ensure to conform to the following requirements for Azure to AHV migration.

- Supported browser: Google Chrome
- Ensure that you enable all outbound ports in the source VM.
- Ensure that you have PowerShell version 4.0 or later.
- Ensure TCP 443 connection to Azure endpoint for operations in Azure.
- For Linux source VMs, install the following packages along with their dependencies: *wget*, *curl*, *jq*, *bash*, *sudo*.
- For automated guest VM preparation, make sure that the Azure Guest Agent is running on the Azure VMs.
- Ensure to register the Azure app and apply the required privileges for the subscription. You can do it through Azure UI, refer to Registering an App and Applying Privileges (Azure UI) on page 156 or Move CLI, refer to Registering the App in Azure and Assigning Custom Role (Move CLI) on page 164.
- Network Security Services (NSS) 3.44 is required for Linux VMs.
- Before you initiate a migration, ensure to disable backups.
- The resource provider `Microsoft.compute` should be registered for the subscription.
- While adding an Azure source, the Azure custom role assigned to the registered app must have the set of permissions as provided in the following JSON to do end-to-end migration of the VMs.

```
{ "permissions": [
           {
               "actions": [
                   "Microsoft.Authorization/roleAssignments/read",
                   "Microsoft.Authorization/roleDefinitions/read",
                   "Microsoft.Authorization/locks/read",
                   "Microsoft.Resources/subscriptions/locations/read",
                   "Microsoft.Compute/virtualMachines/read",
                   "Microsoft.Compute/virtualMachines/write",
                   "Microsoft.Compute/virtualMachines/vmSizes/read",
                   "Microsoft.Compute/locations/vmSizes/read",
                   "Microsoft.Compute/virtualMachines/instanceView/read",
                   "Microsoft.Compute/locations/runCommands/read",
                   "Microsoft.Compute/virtualMachines/runCommand/action",
                   "Microsoft.Network/networkInterfaces/read",
                   "Microsoft.Network/virtualNetworks/read",
                   "Microsoft.Compute/disks/read",
                   "Microsoft.Resources/subscriptions/resourceGroups/read",
                   "Microsoft.Resources/subscriptions/resourceGroups/write",
                   "Microsoft.Resources/subscriptions/resourceGroups/delete",
                   "Microsoft.Network/networkSecurityGroups/securityRules/write",
                   "Microsoft.Network/networkSecurityGroups/securityRules/delete",
                   "Microsoft.Compute/snapshots/write",
                   "Microsoft.Compute/snapshots/read",
                   "Microsoft.Compute/snapshots/delete",
                   "Microsoft.Compute/snapshots/beginGetAccess/action",
```

```
                    "Microsoft.Compute/snapshots/endGetAccess/action",
                    "Microsoft.Compute/virtualMachines/start/action",
                    "Microsoft.Compute/virtualMachines/powerOff/action",
                    "Microsoft.Compute/disks/beginGetAccess/action"
                ],
                "notActions": [],
                "dataActions": [],
                "notDataActions": []
            }
        ]
    }
```

- While adding Azure as both source and target, the Azure custom role assigned to the registered app must have the set of permissions as provided in the following JSON to do end-to-end migration of the VMs.

```
{
    "permissions": [
      {
        "actions": [
          "Microsoft.Authorization/roleAssignments/read",
          "Microsoft.Authorization/roleDefinitions/read",
          "Microsoft.Compute/disks/beginGetAccess/action",
          "Microsoft.Compute/disks/read",
          "Microsoft.Authorization/locks/read",
          "Microsoft.Compute/locations/runCommands/read",
          "Microsoft.Compute/locations/vmSizes/read",
          "Microsoft.Compute/snapshots/beginGetAccess/action",
          "Microsoft.Compute/snapshots/delete",
          "Microsoft.Compute/snapshots/endGetAccess/action",
          "Microsoft.Compute/snapshots/read",
          "Microsoft.Compute/snapshots/write",
          "Microsoft.Compute/virtualMachines/instanceView/read",
          "Microsoft.Compute/virtualMachines/powerOff/action",
          "Microsoft.Compute/virtualMachines/read",
          "Microsoft.Compute/virtualMachines/runCommand/action",
          "Microsoft.Compute/virtualMachines/start/action",
          "Microsoft.Compute/virtualMachines/vmSizes/read",
          "Microsoft.Compute/virtualMachines/write",
          "Microsoft.Network/networkInterfaces/read",
          "Microsoft.Network/networkSecurityGroups/securityRules/delete",
          "Microsoft.Network/networkSecurityGroups/securityRules/write",
          "Microsoft.Network/virtualNetworks/read",
          "Microsoft.Resources/subscriptions/locations/read",
          "Microsoft.Resources/subscriptions/resourceGroups/delete",
          "Microsoft.Resources/subscriptions/resourceGroups/read",
          "Microsoft.Resources/subscriptions/resourceGroups/write",
          "Microsoft.Compute/disks/delete",
          "Microsoft.Compute/disks/endGetAccess/action",
          "Microsoft.Compute/disks/write",
          "Microsoft.Network/networkInterfaces/delete",
          "Microsoft.Network/networkInterfaces/effectiveNetworkSecurityGroups/
action",
          "Microsoft.Network/networkInterfaces/join/action",
          "Microsoft.Network/networkInterfaces/write",
          "Microsoft.Network/networkSecurityGroups/join/action",
          "Microsoft.Network/networkSecurityGroups/read",
          "Microsoft.Network/virtualNetworks/subnets/join/action"
        ],
        "notActions": [],
        "dataActions": [],
        "notDataActions": []
      }
```

```
    ]
  }
```

- Move must have access to the following URLs:
  - https://management.azure.com/
  - https://login.microsoftonline.com/
  - https://graph.windows.net/
  - https://batch.core.windows.net/
  - https://*.blob.core.windows.net/ (Azure Storage blobs)
  - https://nxmove.blob.core.windows.net/
- Move must have access to the shared access signature (SAS) URL https://*.blob.storage.azure.net/ generated by Azure for snapshot exports.

Prerequisites for Linux guest VMs:

- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

Service Accounts

For successful migration of VMs from Azure to AHV, Move requires the following.

- Prism Element UI for the AHV cluster
- (Only for manual preparation of VMs) An administrator for Windows source VMs or root for Linux source VMs to run the source VM preparation scripts.

Registering an App and Applying Privileges (Azure UI)

For migrating the VMs from Azure, first you need to register the Azure app, and then apply the required privileges for the subscription in the Azure UI. Once you have the Subscription ID, Tenant ID, Application ID and the client secret value, you can add the Azure provider in the Move VM.

About this task

To register the app and apply the required privileges, do the following:

Procedure

1. Create a new App registration in Azure to provide necessary access to Move to perform the VM migration.

   a. Select **App registration** > **New registration**.

Figure 30: App registration

   b. Provide the app name, and then click **Register** to register the application with the default selection.

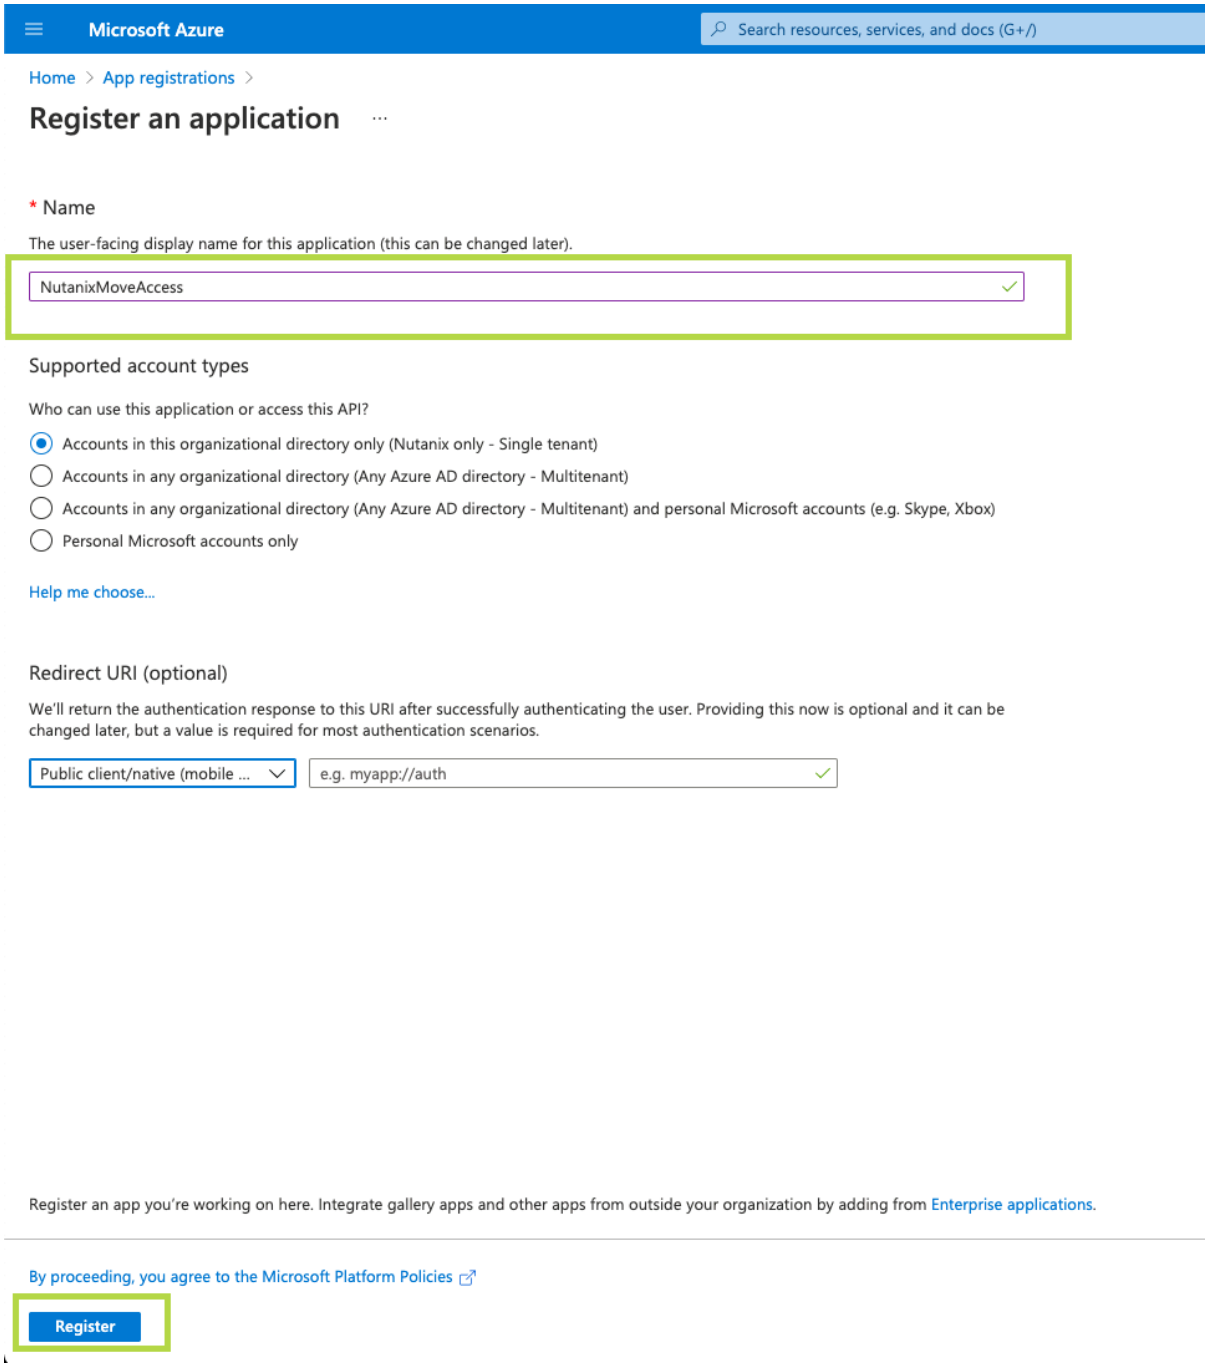

Figure 31: Register the app with default selection

c. Once the application is registered, copy the Application ID and the Tenant ID.

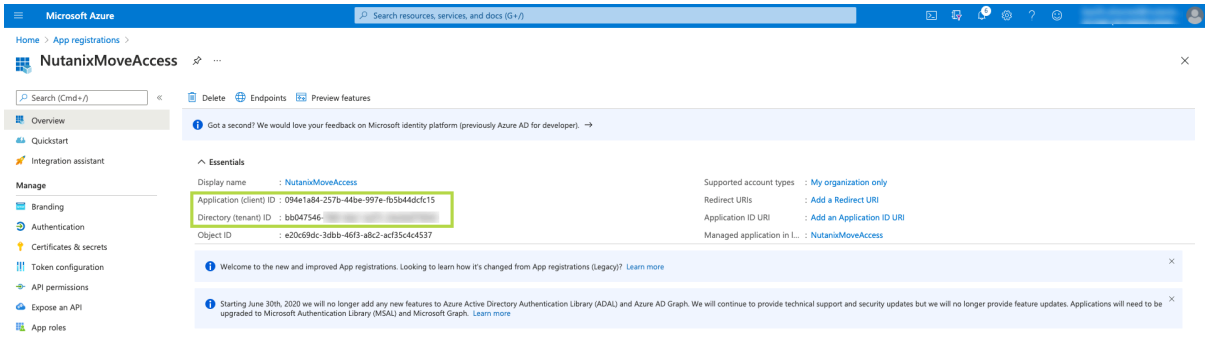

Figure 32: Copy Application ID and Tenant ID

d. Click **Certificates & Secrets** to create a client secret for the registered application. Set the description and the expiry period, and the click **Add**.

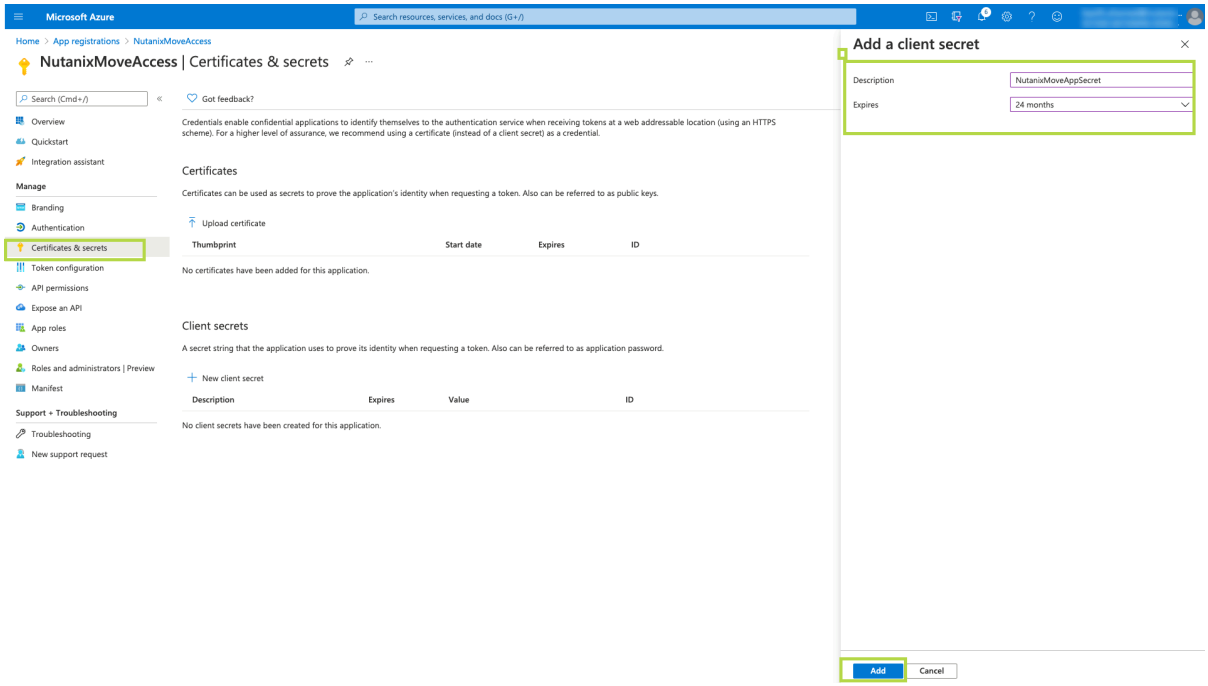

Figure 33: Creating a client secret

e. Copy the client secret value which is required for adding Azure provider in the Move VM.

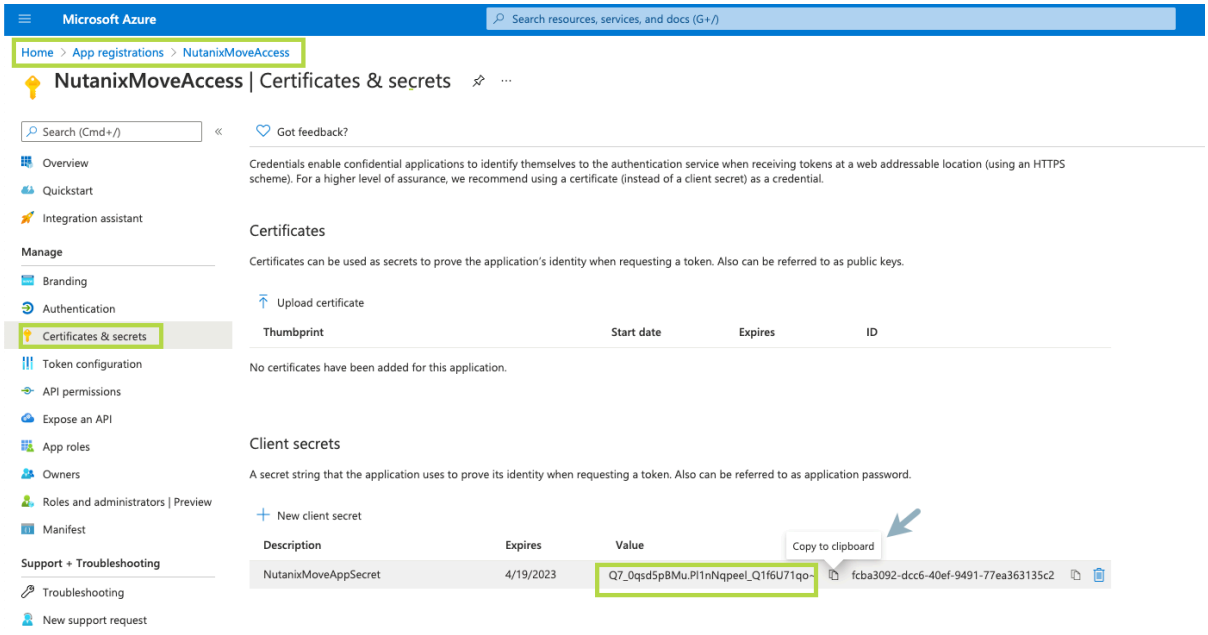

Figure 34: Copying Secret

2. Create a custom role in the subscription and assign that role to the application.

   a. Go to **Subscriptions** > **(Name of your subscription to add to Move)** > **Access control (IAM)**. Click **Add** > **Add custom role**.

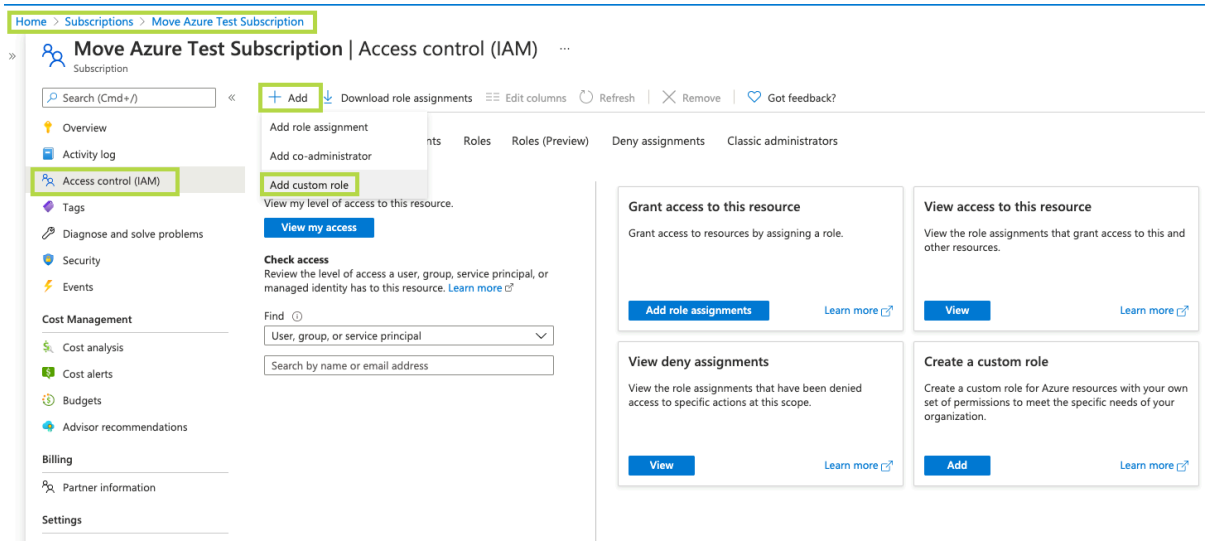

Figure 35: Adding a custom role

   b. Enter a custom role name, and then click the **JSON** tab. Click **Edit**. Replace the `permissions` section in the JSON with the required *set of permissions*, and then click **Save**. Click **Review + Create** to complete the custom role creation.

      To copy the *set of permissions*, refer to Requirements (Azure to AHV) on page 154 or Requirements (Azure to ESXi) on page 177 section.

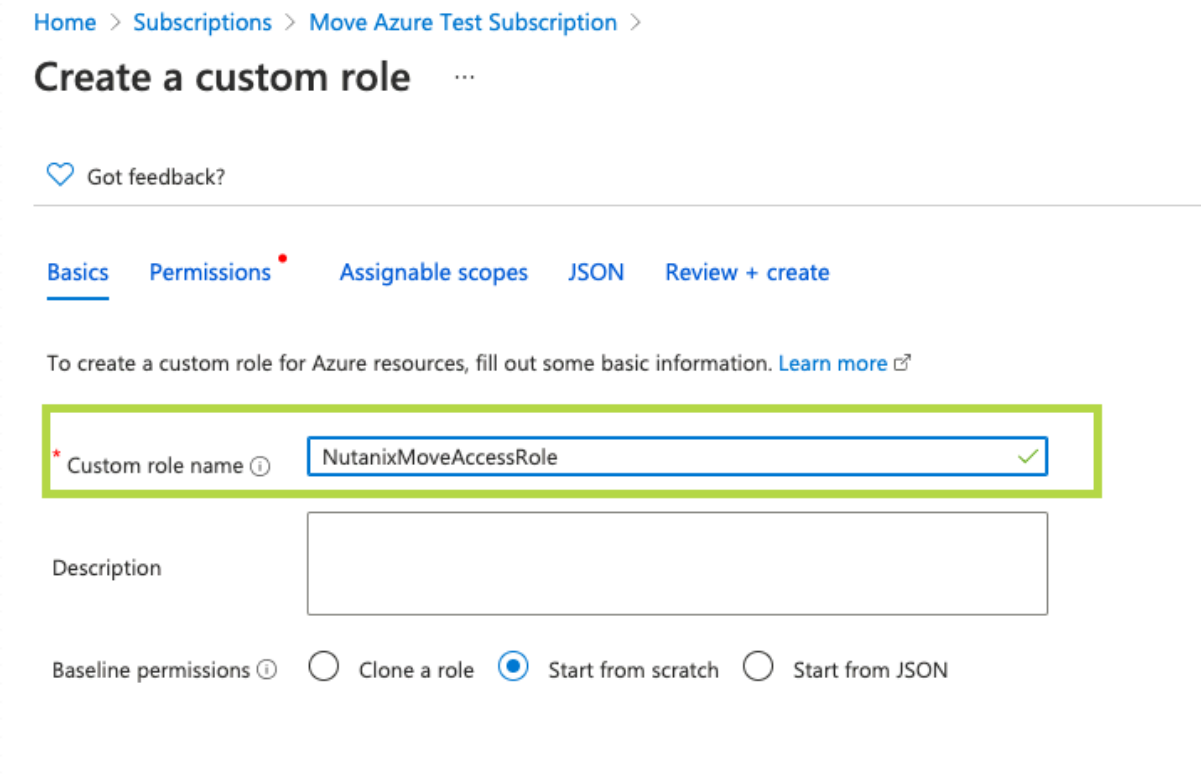

Figure 36: Creating a custom role

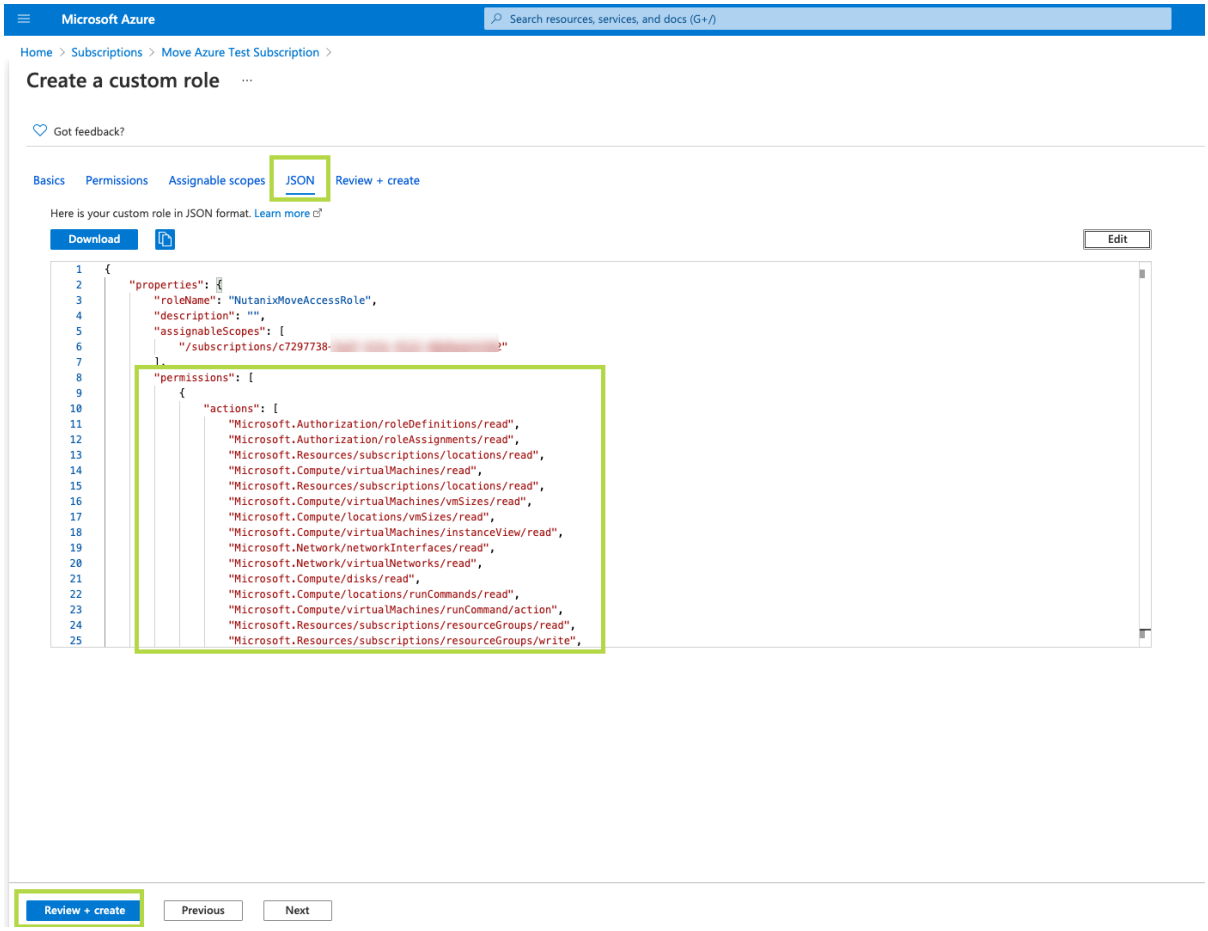

Figure 37: Editing JSON for required set of permissions

c. Go to **Access control (IAM)**. Click **Add** > **Add role assignment**.

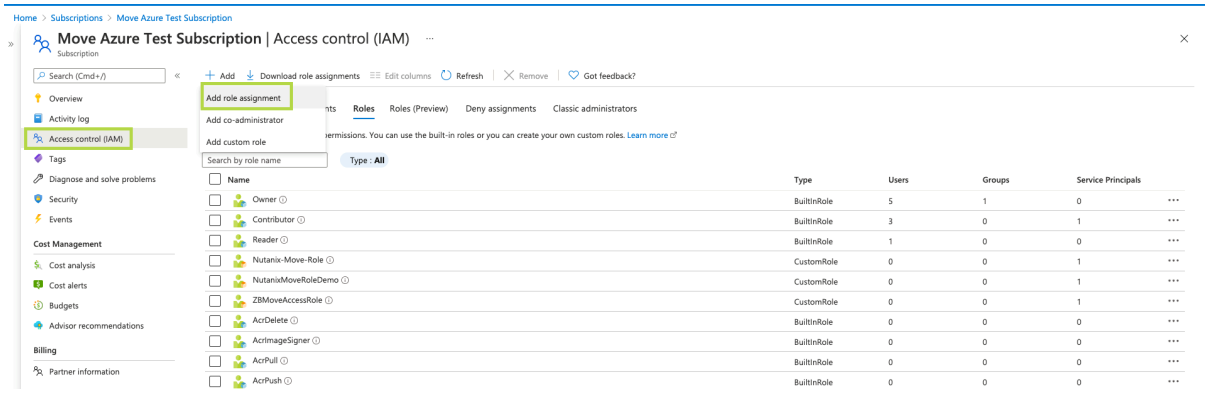

Figure 38: Assigning Permissions

d. Select the created role and the registered application, and then click **Save**.

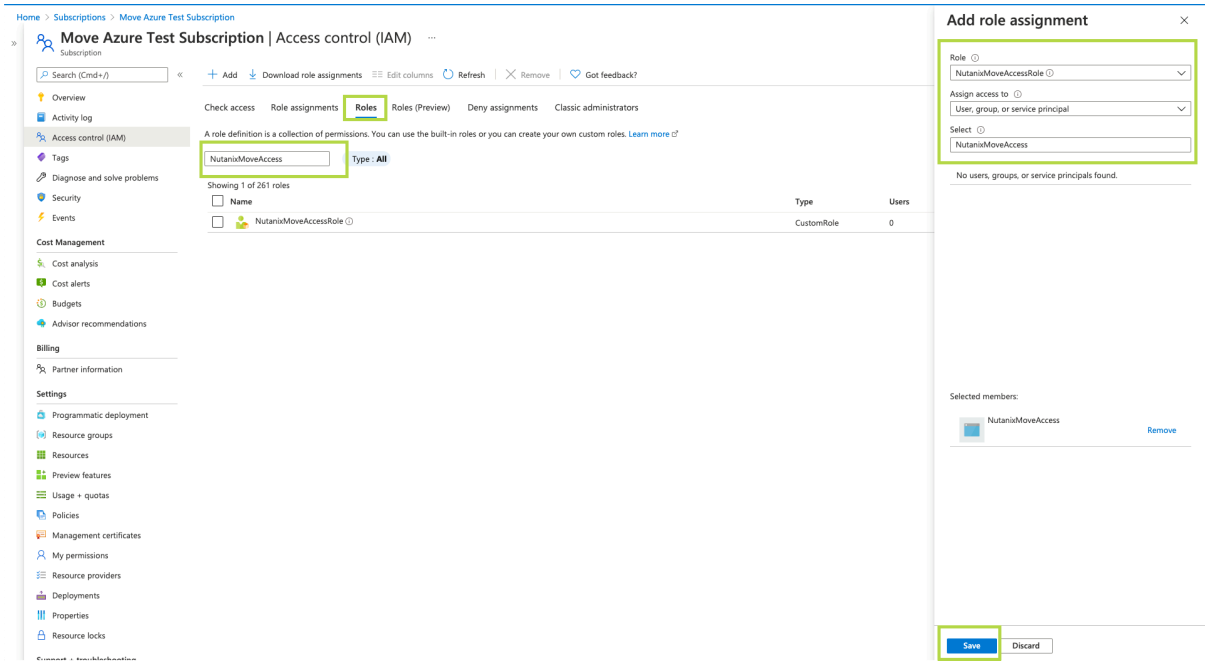

Figure 39: Role assignment

Now you can use this Subscription ID, Tenant ID, Application ID and the client secret value to add the Azure provider in the Move VM.

What to do next

You can now add the Azure environment in the Move UI. Refer to Adding an Azure Environment on page 167. You can also register the app in Azure and assign custom role through Move CLI. Refer to Registering the App in Azure and Assigning Custom Role (Move CLI) on page 164.

Registering the App in Azure and Assigning Custom Role (Move CLI)

For migrating the VMs from Azure, first you need to register the Azure app, and then apply the required privileges for the subscription through Move CLI or Azure UI. Once you have the Subscription ID, Tenant ID, Application ID and the client secret value, you can add the Azure provider in the Move VM.

About this task

To register the app in Azure and assign custom role, do the following:

Procedure

1. SSH to the Move VM as an admin.
   Refer to Accessing Move VM with SSH on page 21.
2. Switch to the root user by entering the password of the Move VM.

```
admin@move on ~ $ rs
[sudo] password for admin:
```

3. To create Azure App, run the following command:

```
root@move on ~ $ create-azure-app
```

Figure 40: Creating Azure App

A link and code is provided to authenticate with Azure.

4. Open the authentication link in a web browser and enter the code for authentication. Select the Azure account.

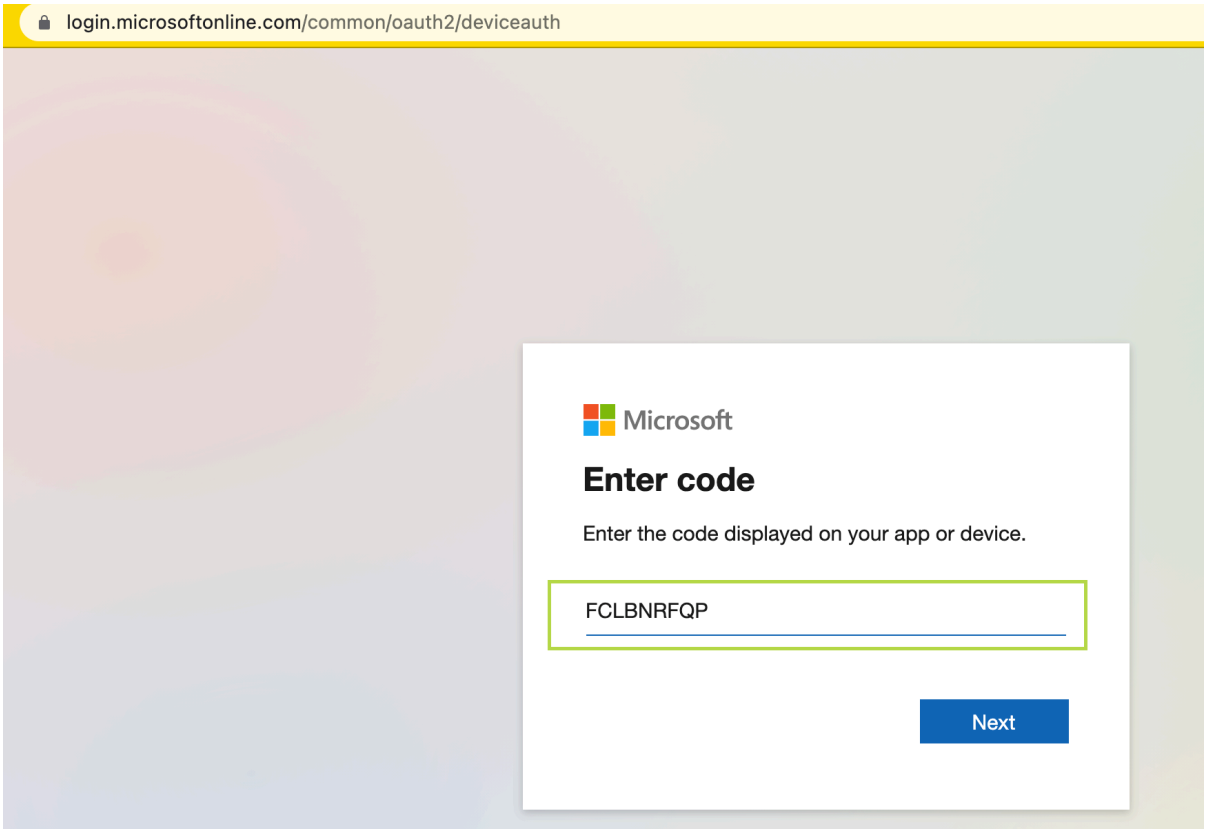

Figure 41: Authenticating with Azure

Authentication will be successful and a list of subscriptions will appear.

5. If your account has multiple subscriptions, provide the subscription ID to be used for migration.

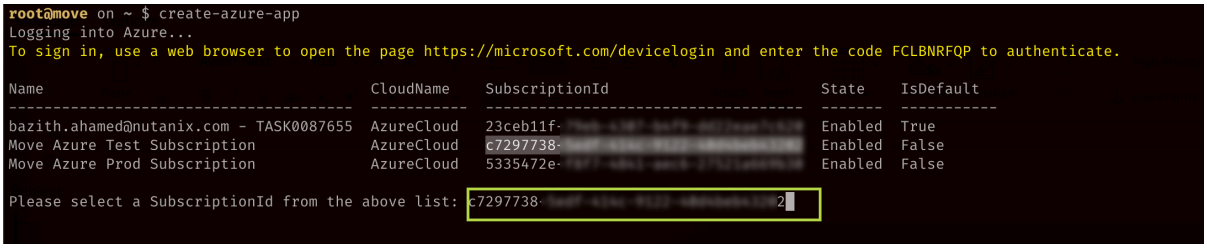

Figure 42: Enter the subscription ID

6. Enter the App name to be registered (default: NutanixMoveApp) and the custom role name which will be created with necessary permission and assigned to the registered app.

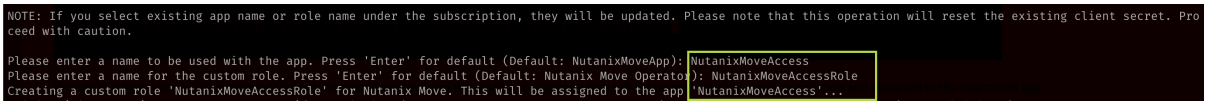

Figure 43: Enter the App and custom role name

Once the App gets registered in Azure and assigned with the custom role, the Subscription ID, Tenant ID, Application ID and the Client Secret is displayed in the output. Use these credentials to add Azure as a provider in the Move VM.

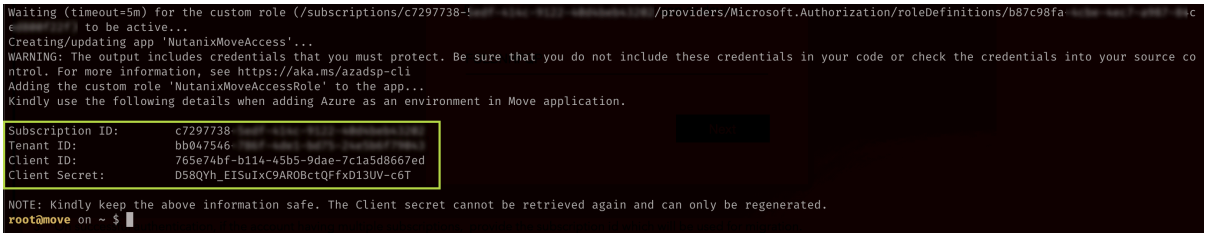

Figure 44: Lists the Subscription ID, Tenant ID, Application ID and the Client Secret

What to do next

You can now add the Azure environment in the Move UI. Refer to Adding an Azure Environment on page 167. You can also register the Azure app, and then apply the required privileges for the subscription from the Azure UI. Refer to Registering an App and Applying Privileges (Azure UI) on page 156.

Qualified Metrics (Azure to AHV)

The following metrics are qualified for migration from Azure.

- Migration of VMs with 2 TB disk.
- Migration of 50 VMs in a single migration plan.
- No restriction regarding any locations.

Unsupported Features (Azure to AHV)

This section lists the unsupported features for migration from Azure.

- Guest operating systems other than the supported operating systems.

  For more information, refer to Supported Guest Operating Systems on page 153
- VMs with unmanaged disks
- VMs with ephemeral disks
- IP address and MAC address retention
- Spot instances gets deleted or shutdown at any time according to the user configuration and have an impact on migrations.
- VDI VMs created with Citrix Machine Creation Service

Limitations (Azure to AHV)

The section lists the limitations for migration from Azure.

Azure to AHV Migration Limitations

The support for migration is constrained by the following.

- VMs with shared disks
- VMs with private Virtual Private Network (VPN)
- VMs with third party backup
- Azure Public Cloud is supported.

  Note: Azure China Cloud, Azure German Cloud, and AzureUSGovernment are not supported.

Adding an Azure Environment

While creating a migration plan, if you need to add an Azure source, you have to add at least one Azure environment for migration.

Before you begin

Ensure to register the Azure app and apply the required privileges for the subscription. You can do it through Azure UI or Move CLI.

About this task

Note:

- This procedure is only applicable for migration from Azure environment.
- Azure accounts with lock at subscription level cannot be added as a provider. If the lock is present at resource group level, VMs under that resource group are marked as unmigratable.

To add an Azure environment, do the following:

Procedure

1. Log on to Move UI.

2. Click **+ Add Environment** under Environments.
   The **Add Environment** appears.

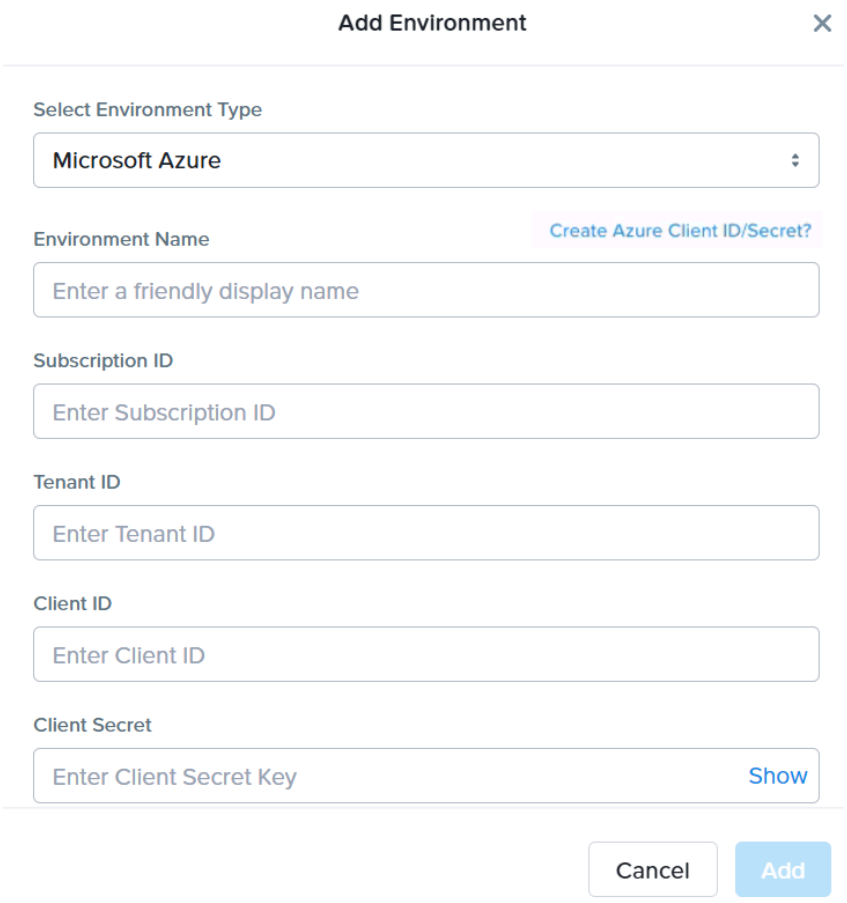

Figure 45: Add Environment Dialog Box

3. Select **Microsoft Azure** as the environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the Azure environment.
   b. **Subscription ID**: Enter your Azure subscription ID.
   c. **Tenant ID**: Enter your Azure tenant ID.
   d. **Client ID**: Enter the client ID of the Azure account.
      If you do not have the Client ID and Secret, click the **Create Azure Client ID/Secret?** link for more details on creating Azure Client ID and Secret.
   e. **Client Secret**: Enter the value of the client secret.

      Note: Client Secret refers to the value of the client secret in Azure and not the secret ID.

   The environment is added to Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

You can add a Nutanix AOS cluster environment. For more information, refer to Adding a Nutanix AOS Cluster Environment on page 39

Adding a Nutanix AOS Cluster Environment

While creating a migration plan, AHV AOS cluster can be added as both source or target. If you want to use ESXi on Nutanix cluster as target, then add the corresponding AOS cluster environment for ESXi.

About this task

> **Note:** When you add a AOS cluster, Move VM IP address with the subnet 255.255.255.255 is added to the global NFS allowlist.
>
> **Caution:** Modifying the NFS allowlist of the destination container disables inheritance of the NFS allowlist at the global level and lead to issues with Move operations.

To add a Nutanix AOS cluster environment, do the following:

Procedure

1. Log on to the Move UI.
2. Click **+ Add Environment** under **Environments**.

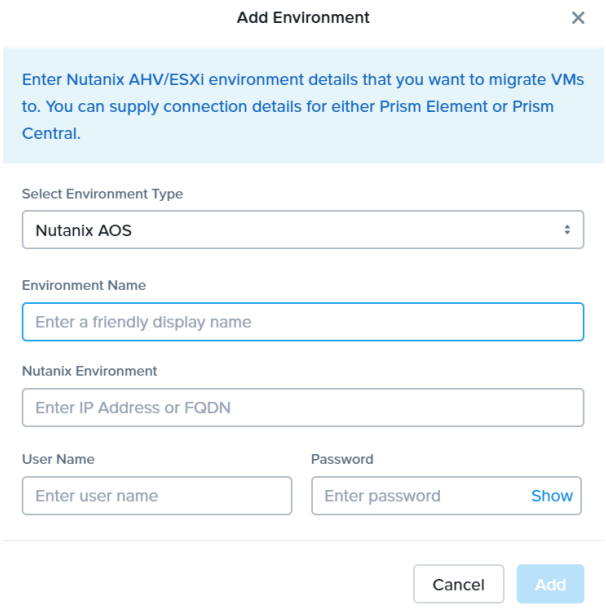

Figure 46: Add AOS Environment Dialog Box

The **Add Environment** window appears.

3. Select **Nutanix AOS** as the target environment type.

4. Complete the indicated fields and click **Add**.

   a. **Environment Name**: Enter a name for the NC2 on AWS and AHV or VMware ESXi on Nutanix environment.

   b. **Nutanix Environment**: Enter the IP address or the FQDN for the target Prism Central or Prism Element.

   c. **User Name**: Enter the username for logging on to the target Nutanix environment.

   d. **Password**: Enter the password for logging on to the target Nutanix environment.

5. (Only for ESXi to ESXi migration) Enter credentials for registered vCenter.

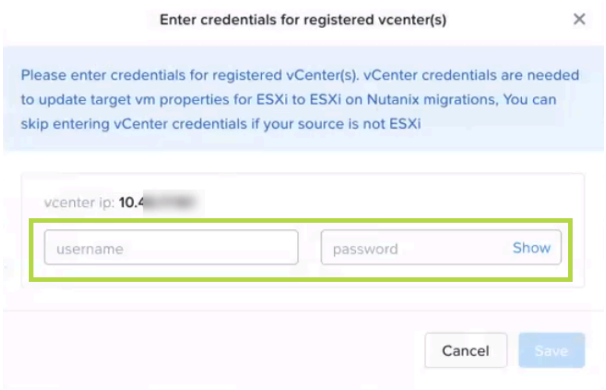

Figure 47: vCenter Credentials Dialog box

vCenter credentials are required to update the target VM properties for ESXi to ESXi on Nutanix migration.

Note:

- You can skip adding vCenter credentials if your source is not ESXi.
- To retain the VM properties on the target VM after migration, be sure to provide the target vCenter credentials. The following properties will be retained:
  - SCSI controller types
  - Network adaptor type
  - MAC address
  - Video card
  - Memory overcommit variables
  - Tool upgrade flag
  - Sync time with host flag
  - Disable acceleration flag
  - Enable logging flag

The AOS environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

Once you have added the environments, you can proceed to create the migration plan. For more information, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan (Hyper-V to ESXi) on page 103 or Creating a Migration Plan (Azure to AHV) on page 171 or Creating a Migration Plan (Azure to ESXi) on page 184

Creating a Migration Plan (Azure to AHV)

You can create a migration plan to seed the data, cutover, and monitor the VMs. You can create the migration plan in Move without initiating the cutover process.

About this task

Note:

- This procedure is only applicable for migration from Azure to AHV.
- On your first logon, you can log on to Move with your defaults credentials.
- If you cancel or discard the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.

  For information on canceling an ongoing migration, see Pausing or Canceling a VM Migration on page 257.

  For information on performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

To create a migration plan, do the following:

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.
3. On the Move dashboard, click **Create a Migration Plan**.

   Note: If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.
4. Enter the new migration plan name, and then click **OK**.

   Note: You can edit the migration plan name by clicking the pencil icon next to the migration plan name.

   The **New Migration Plan** window appears.

   Note: You might at times see a message relating to inventory collection as shown below. During this time, the environments undergoing refresh will not be available for selection. Such environments do not automatically show up for selection once the inventory collection is over. In case you wish to select one of the environments

undergoing refresh (not available for selection), you will need to select **Cancel** and wait for inventory collection to get over, and then create the migration plan again.

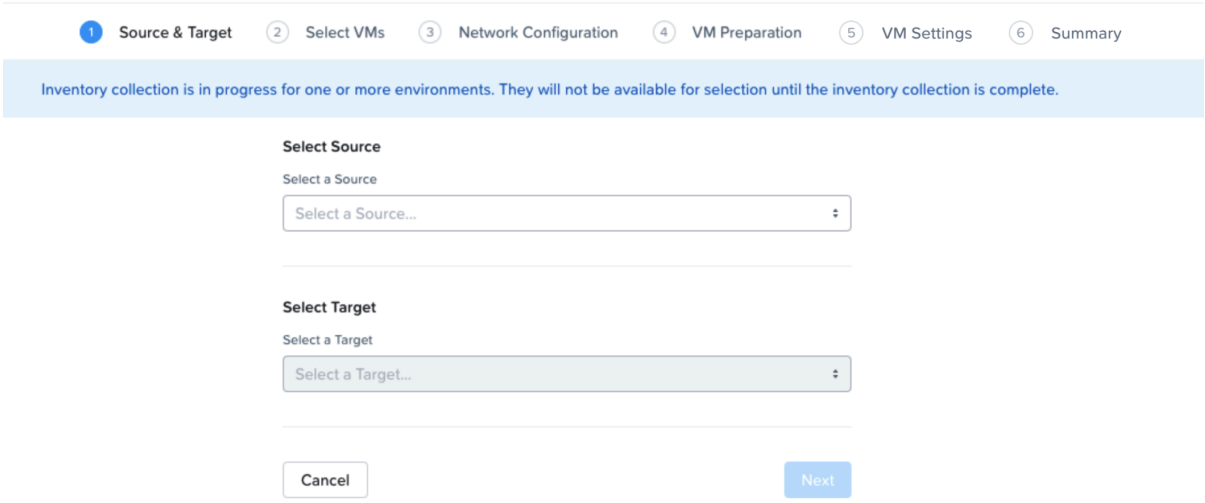

Figure 48: Inventory Collection Message

5. Complete the following fields, and then click **Next**.

   a. **Select a Source**: Select any AWS source for migration.

   b. **Location**: Select a location from where you will migrate the VMs.

      Only locations with VMs are available for selection.

   c. **Select a Target**: Select the target for migrating the VMs.

   d. **Target Project** (optional): Select the project you want as the target.
      This field is available only with Prism Central.

   e. **Target Owners**: Select the owners for the selected target project.
      This field is available only when a target project is selected.

   f. **Target Cluster**: Select the cluster on which you will migrate the VMs.

   g. **Target Container**: Select the container on which you will migrate the VMs.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click **+Add All**, and then click **Next**.

   Note: You cannot add more than 50 VMs in a single migration plan.

   You can filter the VMs by name in the **Filter** bar. You can also sort the VM types by selecting the appropriate column.

   Note: The VMs for which migration has failed are not displayed. To show the entire list of VMs, select **All VMs** from the drop-down list. To show the configuration of VMs, select **Configuration** from the drop-down list. A

question mark icon appears beside an unavailable VM that displays more information about that VM and indicate the reason of a failed VM migration.

Note: Migrated VMs retain their power state on the destination cluster.

The selected VMs are displayed in the sidebar.

7. In the **Network Configuration** screen, select the target network from the drop-down.
   It might take some time to check the capacity on the target cluster for availability of resources for migration.

   - [Applicable only if Prism Central is used] Move provides the option to select VPC-based or VLAN-based target subnets. Based on the selection of a VPC or VLAN ID as the target network, the respective subnets are listed in the target subnet drop-down menu. Select the required subnet from the drop-down menu.

     Note: Overlay subnets which do not have IP address pool(s) associated will be disabled in the subnet drop-down menu.

   Click **Next**.

8. In the **VM Preparation** screen, select one of the following VM preparation modes.

   » **Automatic**. By default, this option is selected. Move automatically runs scripts on the source VMs to prepare them for migration. Ensure that the Azure agent is running on the Azure VMs.

   » **Manual**. Move displays the VM preparation scripts for Windows and Linux VMs. Copy the scripts and manually run them on the respective source VMs, and then click **Next**.

     Note: These scripts prepare the instance by performing the NTNX VirtIO driver installation.

   » **Mixed**. Move allows you to select the VM preparation mode for each VM. This setting allows you to manually prepare one set of VMs and automatically prepare another set of VMs in the same migration plan. If you want to have such a hybrid combination of **Automatic** and **Manual** VM preparation modes, proceed to the next step.

     Note: The preparation mode automatically switches to **Mixed** if you change the mode of preparation for a set of VMs under **Change Settings** and have a mix of **Automatic** and **Manual** VM preparation modes. You cannot manually select the **Custom** option from the **Preparation Mode** drop-down list.

9. In the **Override individual VM Preparation** section, click **Change Settings** to override settings for the individual VMs. You can edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

   Note: If you change the **Mode of Preparation** to **Manual** for a VM, then copy the new generated scripts of that specific VM and run them on the source VM.

   Access the newly generated VM preparation script for that VM by selecting the Copy Scripts icon under the **Actions** column.

   After making the required changes, click **Done**.

   Then, click **Next**.

10. In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

    a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

    b. **Timezone**: Set the timezone as the hardware clock of the VMs in target.

    If set to default, Move configures UTC timezones for Linux VMs and cluster timezones for Windows VMs.

    c. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

    Only those categories which have values are available for selection.

    d. **Skip CDROM addition on target VMs**: Select this check box to skip the CDROM addition on the target VMs.

    e. **Enable Memory Overcommit**: Select this option to enable memory overcommit on the target VM.

    For more information on memory overcommit deployment, refer to AHV Administration Guide.

    f. **Settings for individual VMs**: Click **Change Settings** to configure settings such as timezone, VM priority, and skip CDROM addition for individual VMs. You can also search the VM by typing the name of the VM and change the settings.

    g. **Schedule Data Seeding**: Check this check box to select the date and time for migration.

11. In the **Summary** screen, choose one of the following, and then proceed to review the VM migration summary.

    » **Back**: To edit the information, click this option.

    » **Save**: To save the migration plan, click this option.

    For more information about how to start the migration later, and check the migrated VM status and details, refer to Environments and Migration Plan Management on page 259.

    » **Save and Start**: Click this option to save the migration plan and begin the migration immediately

    The seeding process for migration begins. You can monitor this information by selecting **Status** for the migration plan.

    > **Note:** The seeding process can take several minutes depending on the number of VMs.

What to do next

If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover on page 174

Performing a Migration Cutover

When the migration plan is started and the seeding process is complete, you can cut over the selected VMs in the AHV cluster. You can monitor the VM migration progress by clicking the **Status** link.

About this task

> **Note:**
>
> - You can perform this operation only when the VM status is Ready to Cutover.
> - Recommended that cutover should be performed within one week of initial data seeding.

- If the initial data seeding finishes in less than 10 minutes, Move continues to wait for 10 minutes to take the incremental snapshot; however, you can trigger the cutover immediately.
- The cutover process increments in absolute numbers.

To perform a migration cutover, do the following:

Procedure

1. In the Move UI, click the **Ready to Cutover** status to display the list of available VMs.
2. To perform the cutover, select the VMs or group of VMs.
3. Click **Cutover**.
   The cutover process performs the following VM actions.
   - Shuts down the instance
   - Takes the final snapshots for the instance and copying the final changes to AHV
   - Creates a VM in the target AHV cluster
   - Attaches the replicated drives to the VM
   - Powers on or off the VM (depends on the initial power state)

   The cutover process begins immediately and might take a few minutes. Once cutover is complete, the VM is ready for use in the new AHV cluster.

What to do next

You can manage the migration plan once the migration plan is ready. For more information, refer to Environments and Migration Plan Management on page 259

Performance Matrix for Large Data Migration

Move performs end-to-end migration of large VMs. The scenarios are tested based on the following parameters. The following tables show the performance numbers from the Move lab.

Table 16: Performance Numbers of Large Data Migration (Azure to AHV)

Total migration size	Number of Disks	Location used	Average latency source - target	Data seeding duration	Data transfer rate (Mbps)
2 TiB	1	US West	19 ms	3 hours 5 minutes	1157 Mbps
2 TiB	4*500 GiB	US West	19 ms	4 hours 43 minutes	985 Mbps
2 TiB	1	Central India	239 ms	13 hours 9 minutes	354 Mbps

AZURE TO ESXI

You can prepare and migrate Azure VMs to ESXi by using Move.

> **Note:** ESXi as a target refers to ESXi running on Nutanix appliances.

Migration Considerations

You must consider the supported guest operating systems, requirements, recommendations, unsupported features, and limitations provided in this section before starting the migration process.

Supported Guest Operating Systems

Move supports some common operating systems. Unless otherwise specified, Nutanix has qualified the following 64-bit guest operating system versions.

Generation 1 VMs Support

- Windows 8, 8.1, 10
- Windows Server 2012, 2012 R2, 2016, 2019
- RHEL 6.3 (32-bit and 64-bit supported), 6.5-6.10, 7.0–7.7, and 8.0–8.2

  > **Note:** RHEL 6.3 is supported only with IDE as the disk controller.

- CentOS 6.3 (32-bit and 64-bit supported) to 6.9, 7.0–7.7, 8.0-8.2

  > **Note:** CentOS 6.3 is supported only with IDE as the disk controller.

- Ubuntu Server and Desktop 12.04.5, 14.04.x, 16.04.x, 16.10 (32-bit and 64-bit supported)
- Ubuntu Server 12.0.4, 18.04, 19.04
- SUSE Linux Enterprise Server 12 SP5, 15 SP1

  > **Note:** Kernels which have *vmxnet3*, *mptbase*, *mptsas*, *mptscsih* drivers will work.

- Oracle Linux 7.5-8.0
- Debian 10

  > **Note:** Kernels which have *vmxnet3*, *mptbase*, *mptsas*, *mptscsih* drivers will work.

Generation 2 VMs Support (UEFI Enabled VMs)

- Windows 10 Pro, Windows 10 Enterprise
- Windows Server 2012, 2012 R2, 2016, 2019
- RHEL 7.0, 7.4–7.7, and 8.0–8.1
- CentOS 7.4–7.7, 8.0-8.1
- Ubuntu Server 16.04, 18.04, 19.04, 19.10

- SUSE Linux Enterprise Server 12 SP4, 15 SP1

  Note: Kernels which have *vmxnet3, mptbase, mptsas, mptscsih* drivers will work.

- Oracle Linux 7.7-ci

Requirements (Azure to ESXi)

Before attempting to migrate VMs running on Azure using Move, make sure to conform to the requirements listed here.

General Requirements

Ensure to conform to the following requirements for Azure to ESXi migration.

- Supported browser: Google Chrome
- Ensure that you enable all outbound ports in the source VM.
- Ensure that you have PowerShell version 4.0 or later.
- Ensure TCP 443 connection to Azure endpoint for operations in Azure.
- For Linux source VMs, install the following packages along with their dependencies: *wget, curl, jq, bash, sudo*.
- For automated guest VM preparation, make sure that the Azure Guest Agent is running on the Azure VMs.
- Ensure to register the Azure app and apply the required privileges for the subscription. You can do it through Azure UI, refer to Registering an App and Applying Privileges (Azure UI) on page 156 or Move CLI, refer to Registering the App in Azure and Assigning Custom Role (Move CLI) on page 164.
- Network Security Services (NSS) 3.44 is required for Linux VMs.
- Before you initiate a migration, ensure to disable backups.
- The resource provider `Microsoft.compute` should be registered for the subscription.
- While adding an Azure source, the Azure custom role assigned to the registered app must have the set of permissions as provided in the following JSON to do end-to-end migration of the VMs.

```
{ "permissions": [
            {
                "actions": [
                    "Microsoft.Authorization/roleAssignments/read",
                    "Microsoft.Authorization/roleDefinitions/read",
                    "Microsoft.Authorization/locks/read",
                    "Microsoft.Resources/subscriptions/locations/read",
                    "Microsoft.Compute/virtualMachines/read",
                    "Microsoft.Compute/virtualMachines/write",
                    "Microsoft.Compute/virtualMachines/vmSizes/read",
                    "Microsoft.Compute/locations/vmSizes/read",
                    "Microsoft.Compute/virtualMachines/instanceView/read",
                    "Microsoft.Compute/locations/runCommands/read",
                    "Microsoft.Compute/virtualMachines/runCommand/action",
                    "Microsoft.Network/networkInterfaces/read",
                    "Microsoft.Network/virtualNetworks/read",
                    "Microsoft.Compute/disks/read",
                    "Microsoft.Resources/subscriptions/resourceGroups/read",
                    "Microsoft.Resources/subscriptions/resourceGroups/write",
                    "Microsoft.Resources/subscriptions/resourceGroups/delete",
                    "Microsoft.Network/networkSecurityGroups/securityRules/write",
                    "Microsoft.Network/networkSecurityGroups/securityRules/delete",
                    "Microsoft.Compute/snapshots/write",
```

```
                    "Microsoft.Compute/snapshots/read",
                    "Microsoft.Compute/snapshots/delete",
                    "Microsoft.Compute/snapshots/beginGetAccess/action",
                    "Microsoft.Compute/snapshots/endGetAccess/action",
                    "Microsoft.Compute/virtualMachines/start/action",
                    "Microsoft.Compute/virtualMachines/powerOff/action",
                    "Microsoft.Compute/disks/beginGetAccess/action"
                ],
                "notActions": [],
                "dataActions": [],
                "notDataActions": []
            }
        ]
    }
```

- While adding Azure as both source and target, the Azure custom role assigned to the registered app must have the set of permissions as provided in the following JSON to do end-to-end migration of the VMs.

```
{
    "permissions": [
      {
        "actions": [
          "Microsoft.Authorization/roleAssignments/read",
          "Microsoft.Authorization/roleDefinitions/read",
          "Microsoft.Compute/disks/beginGetAccess/action",
          "Microsoft.Compute/disks/read",
          "Microsoft.Authorization/locks/read",
          "Microsoft.Compute/locations/runCommands/read",
          "Microsoft.Compute/locations/vmSizes/read",
          "Microsoft.Compute/snapshots/beginGetAccess/action",
          "Microsoft.Compute/snapshots/delete",
          "Microsoft.Compute/snapshots/endGetAccess/action",
          "Microsoft.Compute/snapshots/read",
          "Microsoft.Compute/snapshots/write",
          "Microsoft.Compute/virtualMachines/instanceView/read",
          "Microsoft.Compute/virtualMachines/powerOff/action",
          "Microsoft.Compute/virtualMachines/read",
          "Microsoft.Compute/virtualMachines/runCommand/action",
          "Microsoft.Compute/virtualMachines/start/action",
          "Microsoft.Compute/virtualMachines/vmSizes/read",
          "Microsoft.Compute/virtualMachines/write",
          "Microsoft.Network/networkInterfaces/read",
          "Microsoft.Network/networkSecurityGroups/securityRules/delete",
          "Microsoft.Network/networkSecurityGroups/securityRules/write",
          "Microsoft.Network/virtualNetworks/read",
          "Microsoft.Resources/subscriptions/locations/read",
          "Microsoft.Resources/subscriptions/resourceGroups/delete",
          "Microsoft.Resources/subscriptions/resourceGroups/read",
          "Microsoft.Resources/subscriptions/resourceGroups/write",
          "Microsoft.Compute/disks/delete",
          "Microsoft.Compute/disks/endGetAccess/action",
          "Microsoft.Compute/disks/write",
          "Microsoft.Network/networkInterfaces/delete",
          "Microsoft.Network/networkInterfaces/effectiveNetworkSecurityGroups/
action",
          "Microsoft.Network/networkInterfaces/join/action",
          "Microsoft.Network/networkInterfaces/write",
          "Microsoft.Network/networkSecurityGroups/delete",
          "Microsoft.Network/networkSecurityGroups/join/action",
          "Microsoft.Network/networkSecurityGroups/read",
          "Microsoft.Network/networkSecurityGroups/write",
          "Microsoft.Network/publicIPAddresses/delete",
```

```
          "Microsoft.Network/publicIPAddresses/join/action",
          "Microsoft.Network/publicIPAddresses/read",
          "Microsoft.Network/publicIPAddresses/write",
          "Microsoft.Network/virtualNetworks/subnets/join/action"
        ],
        "notActions": [],
        "dataActions": [],
        "notDataActions": []
      }
    ]
  }
```

- Move must have access to the following URLs:
  - https://management.azure.com/
  - https://login.microsoftonline.com/
  - https://graph.windows.net/
  - https://batch.core.windows.net/
  - https://*.blob.core.windows.net/ (Azure Storage blobs)
  - https://nxmove.blob.core.windows.net/
- Move must have access to the shared access signature (SAS) URL https://*.blob.storage.azure.net/ generated by Azure for snapshot exports.

Prerequisites for Linux guest VMs:

- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

Service Accounts

For successful migration of VMs from Azure to ESXi, Move requires the following.

- Prism Element UI for the AHV cluster
- (Only for manual preparation of VMs) An administrator for Windows source VMs or root for Linux source VMs to run the source VM preparation scripts.

Qualified Metrics (Azure to ESXi)

The following metrics are qualified for migration from Azure to ESXi.

- Migration of VMs with 2 TB disk.
- Migration of 50 VMs in a single migration plan.
- No restriction regarding any locations.

Unsupported Features (Azure to ESXi)

This section lists the unsupported features for migration from Azure to ESXi.

- Guest operating systems other than the supported operating systems.

  For more information, refer to Supported Guest Operating Systems on page 153
- VMs with unmanaged disks

- VMs with ephemeral disks
- IP address and MAC address retention
- Spot instances gets deleted or shutdown at any time according to the user configuration and have an impact on migrations.
- VDI VMs created with Citrix Machine Creation Service

Limitations (Azure to ESXi)

The section lists the limitations for migration from Azure to ESXi.

Azure to ESXi Migration Limitations

The support for migration is constrained by the following.

- VMs with shared disks
- VMs with private Virtual Private Network (VPN)
- VMs with third party backup
- The preparation of source VMs is successful only if the SLES and Debian kernels have the *vmxnet3, mptbase, mptsas, mptscsih* drivers in them.
- Azure Public Cloud is supported.

> **Note:** Azure China Cloud, Azure German Cloud, and AzureUSGovernment are not supported.

Adding an Azure Environment

While creating a migration plan, if you need to add an Azure source, you have to add at least one Azure environment for migration.

Before you begin

Ensure to register the Azure app and apply the required privileges for the subscription. You can do it through Azure UI or Move CLI.

About this task

> **Note:**
>
> - This procedure is only applicable for migration from Azure environment.
> - Azure accounts with lock at subscription level cannot be added as a provider. If the lock is present at resource group level, VMs under that resource group are marked as unmigratable.

To add an Azure environment, do the following:

Procedure

1. Log on to Move UI.

2. Click **+ Add Environment** under Environments.
   The **Add Environment** appears.

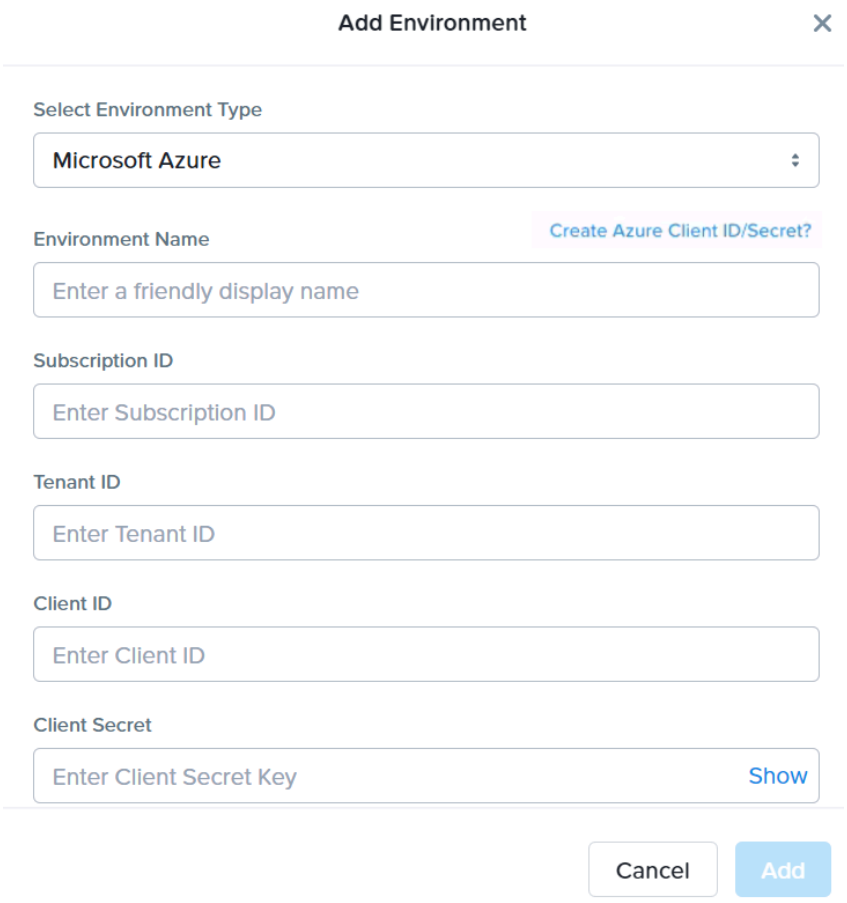

Figure 49: Add Environment Dialog Box

3. Select **Microsoft Azure** as the environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the Azure environment.
   b. **Subscription ID**: Enter your Azure subscription ID.
   c. **Tenant ID**: Enter your Azure tenant ID.
   d. **Client ID**: Enter the client ID of the Azure account.
      If you do not have the Client ID and Secret, click the **Create Azure Client ID/Secret?** link for more details on creating Azure Client ID and Secret.
   e. **Client Secret**: Enter the value of the client secret.

      Note: Client Secret refers to the value of the client secret in Azure and not the secret ID.

   The environment is added to Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

You can add a Nutanix AOS cluster environment. For more information, refer to Adding a Nutanix AOS Cluster Environment on page 39

Adding a Nutanix AOS Cluster Environment

While creating a migration plan, AHV AOS cluster can be added as both source or target. If you want to use ESXi on Nutanix cluster as target, then add the corresponding AOS cluster environment for ESXi.

About this task

> **Note:** When you add a AOS cluster, Move VM IP address with the subnet 255.255.255.255 is added to the global NFS allowlist.
>
> **Caution:** Modifying the NFS allowlist of the destination container disables inheritance of the NFS allowlist at the global level and lead to issues with Move operations.

To add a Nutanix AOS cluster environment, do the following:

Procedure

1. Log on to the Move UI.
2. Click **+ Add Environment** under **Environments**.

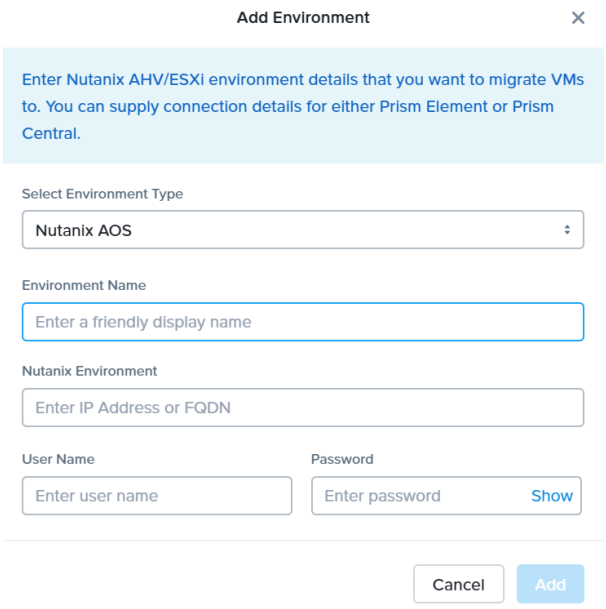

Figure 50: Add AOS Environment Dialog Box

The **Add Environment** window appears.

3. Select **Nutanix AOS** as the target environment type.

4. Complete the indicated fields and click **Add**.

   a. **Environment Name**: Enter a name for the NC2 on AWS and AHV or VMware ESXi on Nutanix environment.

   b. **Nutanix Environment**: Enter the IP address or the FQDN for the target Prism Central or Prism Element.

   c. **User Name**: Enter the username for logging on to the target Nutanix environment.

   d. **Password**: Enter the password for logging on to the target Nutanix environment.

5. (Only for ESXi to ESXi migration) Enter credentials for registered vCenter.

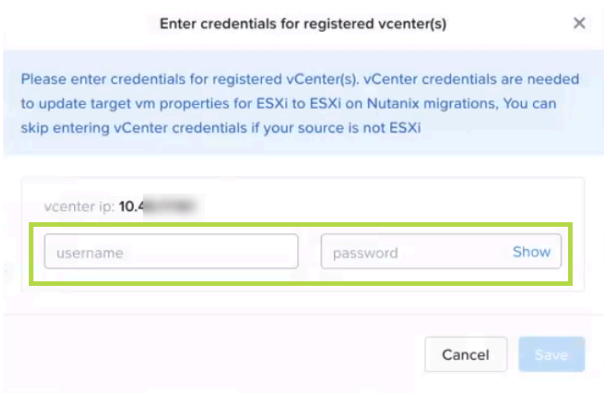

Figure 51: vCenter Credentials Dialog box

vCenter credentials are required to update the target VM properties for ESXi to ESXi on Nutanix migration.

Note:

- You can skip adding vCenter credentials if your source is not ESXi.
- To retain the VM properties on the target VM after migration, be sure to provide the target vCenter credentials. The following properties will be retained:
  - SCSI controller types
  - Network adaptor type
  - MAC address
  - Video card
  - Memory overcommit variables
  - Tool upgrade flag
  - Sync time with host flag
  - Disable acceleration flag
  - Enable logging flag

The AOS environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

Once you have added the environments, you can proceed to create the migration plan. For more information, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan (Hyper-V to ESXi) on page 103 or Creating a Migration Plan (Azure to AHV) on page 171 or Creating a Migration Plan (Azure to ESXi) on page 184

Creating a Migration Plan (Azure to ESXi)

You can create a migration plan to seed the data, cutover, and monitor the VMs. You can create the migration plan in Move without initiating the cutover process.

About this task

Note:

- This procedure is only applicable for migration from Azure to ESXi.
- On your first logon, you can log on to Move with your defaults credentials.
- If you cancel or discard the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.

  For information on canceling an ongoing migration, see Pausing or Canceling a VM Migration on page 257.

  For information on performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

To create a migration plan, do the following:

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.
3. On the Move dashboard, click **Create a Migration Plan**.

   Note: If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.
4. Enter the new migration plan name, and then click **OK**.

   Note: You can edit the migration plan name by clicking the pencil icon next to the migration plan name.

   The **New Migration Plan** window appears.

   Note: You might at times see a message relating to inventory collection as shown below. During this time, the environments undergoing refresh will not be available for selection. Such environments do not automatically show up for selection once the inventory collection is over. In case you wish to select one of the environments

undergoing refresh (not available for selection), you will need to select **Cancel** and wait for inventory collection to get over, and then create the migration plan again.

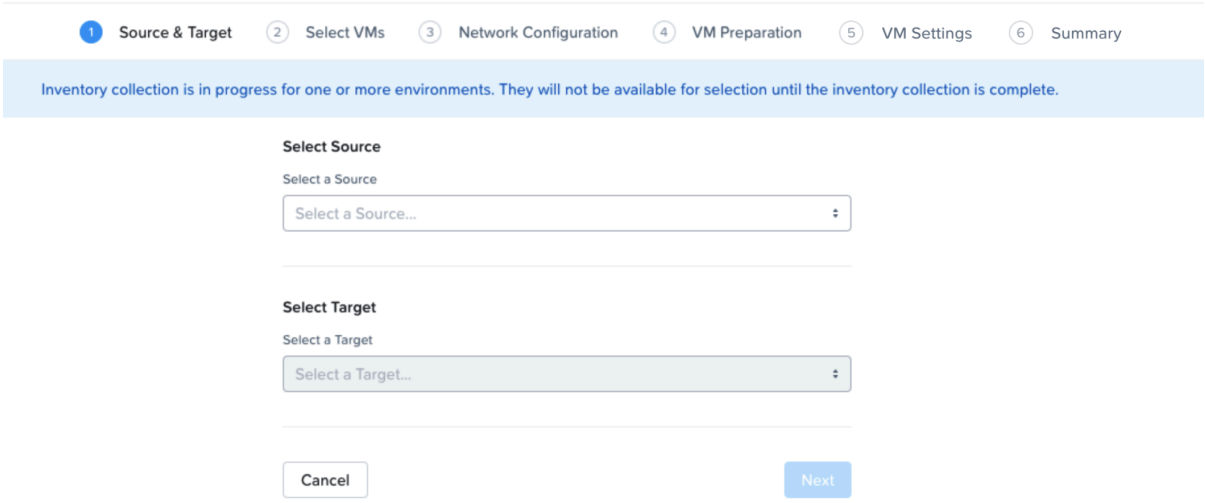

Figure 52: Inventory Collection Message

5. Complete the following fields, and then click **Next**.

   a. **Select a Source**: Select any AWS source for migration.

   b. **Location**: Select a location from where you will migrate the VMs.

      Only locations with VMs are available for selection.

   c. **Select a Target**: Select the target for migrating the VMs.

   d. **Target Container**: Select the container on which you will migrate the VMs.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click **+Add All**, and then click **Next**.

   Note: You cannot add more than 50 VMs in a single migration plan.

   You can filter the VMs by name in the **Filter** bar. You can also sort the VM types by selecting the appropriate column.

   Note: The VMs for which migration has failed are not displayed. To show the entire list of VMs, select **All VMs** from the drop-down list. To show the configuration of VMs, select **Configuration** from the drop-down list. A question mark icon appears beside an unavailable VM that displays more information about that VM and indicate the reason of a failed VM migration.

   Note: Migrated VMs retain their power state on the destination cluster.

   The selected VMs are displayed in the sidebar.

7. In the **Network Configuration** screen, select the target network from the drop-down, and then click **Next**.

   It might take some time to check the capacity on the target cluster for availability of resources for migration.

8. In the **VM Preparation** screen, select one of the following VM preparation modes.

   » **Automatic**. By default, this option is selected. Move automatically runs scripts on the source VMs to prepare them for migration. Ensure that the Azure agent is running on the Azure VMs.

   » **Manual**. Move displays the VM preparation scripts for Windows and Linux VMs. Copy the scripts and manually run them on the respective source VMs, and then click **Next**.

   > **Note:** These scripts prepare the instance by performing the NTNX VirtIO driver installation.

   » **Mixed**. Move allows you to select the VM preparation mode for each VM. This setting allows you to manually prepare one set of VMs and automatically prepare another set of VMs in the same migration plan. If you want to have such a hybrid combination of **Automatic** and **Manual** VM preparation modes, proceed to the next step.

   > **Note:** The preparation mode automatically switches to **Mixed** if you change the mode of preparation for a set of VMs under **Change Settings** and have a mix of **Automatic** and **Manual** VM preparation modes. You cannot manually select the **Custom** option from the **Preparation Mode** drop-down list.

9. In the **Override individual VM Preparation** section, click **Change Settings** to override settings for the individual VMs. You can edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

   > **Note:** If you change the **Mode of Preparation** to **Manual** for a VM, then copy the new generated scripts of that specific VM and run them on the source VM.
   >
   > Access the newly generated VM preparation script for that VM by selecting the Copy Scripts icon under the **Actions** column.

   After making the required changes, click **Done**.

   Then, click **Next**.

10. In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

    a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

    b. **Timezone**: Set the timezone as the hardware clock of the VMs in target.

       If set to default, Move configures UTC timezones for Linux VMs and cluster timezones for Windows VMs.

    c. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

       Only those categories which have values are available for selection.

    d. **Skip CDROM addition on target VMs**: Select this check box to skip the CDROM addition on the target VMs.

    e. **Settings for individual VMs**: Click **Change Settings** to configure settings such as timezone, VM priority, and skip CDROM addition for individual VMs. You can also search the VM by typing the name of the VM and change the settings.

    f. **Schedule Data Seeding**: Check this check box to select the date and time for migration.

11. In the **Summary** screen, choose one of the following, and then proceed to review the VM migration summary.

    » **Back**: To edit the information, click this option.

    » **Save**: To save the migration plan, click this option.

    For more information about how to start the migration later, and check the migrated VM status and details, refer to Environments and Migration Plan Management on page 259.

    » **Save and Start**: Click this option to save the migration plan and begin the migration immediately

    The seeding process for migration begins. You can monitor this information by selecting **Status** for the migration plan.

    > **Note:** The seeding process can take several minutes depending on the number of VMs.

What to do next

If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover on page 174

Performing a Migration Cutover

When the migration plan is started and the seeding process is complete, you can cut over the selected VMs in the AHV cluster. You can monitor the VM migration progress by clicking the **Status** link.

About this task

> **Note:**
>
> - You can perform this operation only when the VM status is Ready to Cutover.
> - Recommended that cutover should be performed within one week of initial data seeding.
> - If the initial data seeding finishes in less than 10 minutes, Move continues to wait for 10 minutes to take the incremental snapshot; however, you can trigger the cutover immediately.
> - The cutover process increments in absolute numbers.

To perform a migration cutover, do the following:

Procedure

1. In the Move UI, click the **Ready to Cutover** status to display the list of available VMs.
2. To perform the cutover, select the VMs or group of VMs.

3. Click **Cutover**.

   The cutover process performs the following VM actions.

   - Shuts down the instance
   - Takes the final snapshots for the instance and copying the final changes to AHV
   - Creates a VM in the target AHV cluster
   - Attaches the replicated drives to the VM
   - Powers on or off the VM (depends on the initial power state)

   The cutover process begins immediately and might take a few minutes. Once cutover is complete, the VM is ready for use in the new AHV cluster.

What to do next

You can manage the migration plan once the migration plan is ready. For more information, refer to Environments and Migration Plan Management on page 259

AHV TO AHV, AHV TO NC2, NC2 TO AHV, AND NC2 (AZURE) TO NC2 (AZURE)

You can perform the following VM migrations using Move:

- AHV to AHV
- AHV to NC2 (on AWS/Azure)
- NC2 (on AWS/Azure) to AHV
- NC2 (on Azure) to NC2 (on Azure)

Note:

- For migration from any source (AHV, ESXi, Hyper-V, and AWS) to AHV target and from any source (ESXi, Hyper-V, and AWS) to Nutanix Cloud Clusters (NC2) on AWS target, Move should be deployed on the same destination target cluster where the VMs need to be migrated. Move appliance is recommended to be deployed on the target cluster (AHV). However, Move can be deployed on the source (AHV) side for either of the following:
  - Your source and target are across geographical locations.
  - The latency between your source and target is more than 200ms.
- For NC2 on AWS, static IP retention is not enabled by default.

Nutanix Guest Tools (NGT) behaviour

The following list details the behaviour of Nutanix Guest Tools (NGT) for migrations:

- NGT will be enabled on the migrated VM only if the source VM has it enabled before migration.
- If the **skip CD-ROM** option is selected during migration, then NGT will not be enabled on the target VM.
- If the source VM does not have a CDROM, then NGT will not be enabled in the target VM.
- If the target cluster does not have virtual IP address configured, then the migrated VM will not have NGT enabled.
- Move will only enable NGT on the migrated VM. It does not change the NGT version on the target VM.
- NGT will not be enabled on the target VM if there is an issue in the target VM network (*ENG-472608*).

Migration Considerations

You must consider the qualified guest operating systems, requirements, recommendations, unsupported features, and limitations provided in this section before starting the migration process.

Qualified Guest Operating Systems

Move supports two ways to migrate a VM from one AHV cluster to another: full migration support and data-only support. Unless otherwise specified, Nutanix has qualified the following 64-bit guest operating system versions.

Full migration support migrates the data, prepares the operating system with the scripts for retaining the IP addresses, and recreates the VM on a target cluster. For full migration support in Windows, the UAC must be disabled.

Caution: During full migration, UAC enabled on a Windows guest breaks the workflow of Move if a built-in local administrator user is not used for migration.

If UAC is enabled or automatic VM preparation fails for certain VMs, you can choose to use manual preparation to prepare such VMs.

Data-Only support migrates the data and recreates the VM on the target. The following lists shows the qualified supported guest operating systems.

For more information about the supported operating systems for the VMs created using UEFI firmware, refer to Compatibility and Interoperability Matrix. It also indicates whether an operating system is community-supported, legacy, or deprecated on AHV.

Note: Either Windows 7 or Windows Server 2008 R2 and earlier versions are not supported with UEFI on AHV.

Qualified Guest Operating Systems

- Ubuntu Server 14.04, 18.04, 20.04
- CentOS 7.5, 8.3
- SUSE 12
- SUSE Linux 15
- RHEL 8.5
- Windows Desktop 10
- Windows Server 2012 R2 DC, 2016 DC, 2019 DC, 2022
- Oracle Linux 7.5, 8.4, 9.1

Note: If you face kernel panic issue on Oracle Linux versions after migration to AHV, then refer and apply the KB article 000004604 for these Oracle Linux VMs.

Data-Only Support

Any guest operating system can be migrated with data-only option. Ensure that it is supported by the target cluster.

Support for UEFI Enabled VMs Migration

Move supports the migration of UEFI-enabled VMs. This topic provides the list of qualified guest operating systems.

Note: Ensure that the target cluster supports the UEFI VMs.

Table 17: Qualified Operating Systems

Operating systems
Windows Server 2016
CentOS 7.4
SUSE 12 SP5

Support for UEFI with Secure Boot Enabled VMs

Move supports UEFI with secure boot enabled VMs.

Table 18: Qualified Guest Operating Systems

Operating systems
Windows Server 2019, 2022
CentOS 7.3, 8.4
RHEL 7.7
Oracle Linux 9.1

Requirements

Before attempting to migrate VMs running on AHV using Move, make sure to conform to the requirements listed here.

General Requirements

Ensure to conform to the following requirements for the migration:

- Supported browser - Google Chrome.
- Ensure you have PowerShell version 4.0 or later.
- For Windows source VMs, ensure to disable UAC for Windows administrator user.
- AHV should support V3 API for migrations.
- Network Security Services (NSS) 3.44 is required for Linux VMs.
- Before you initiate a migration, ensure to disable backups.

For data-only migration, Move does not perform any operation on the operating system. Hence, there is no requirement for the guest operating system.

Prerequisites for Linux guest VMs:

- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

Requirements for Modules or Drivers

Make sure that the following modules and drivers are installed in Windows, Linux, and AHV VMs.

For Windows:

- Enable WinRM for automatic preparation.

  Ensure to enable the following inbound and outbound ports using TCP protocol for the Windows Remote Management (WinRM) feature to work.

  - WinRM-HTTPS: 5986
  - WinRM-HTTP: 5985
  - RDP: 3389 (only for inbound)
  - SSH: 22

- Install Nutanix Guest Tools (NGT) for automatic preparation

For Linux:

- Install Nutanix Guest Tools (NGT) for automatic preparation

- SSH should be enabled with root user privilege for automatic preparation

Service Accounts

For successful migration of VMs, Move requires either of the following to run the source VM preparation scripts:

- An administrator for Windows source VMs
- A root for Linux source VMs

Qualified Metrics

The section lists the qualified metrics for the migration.

The following metrics are qualified for migration:

- Migration of 100 VMs where each VM has 2 disks in a single migration plan.
- 3.5 TB VM migration.
- 100 VMs migration in a single migration plan.

Unsupported Features

This topic lists the unsupported features for the migration.

- Move uses lift-and-shift migration. If the source guest operating system, its features, and its configuration are supported by the target cluster, VM can be migrated by Move. In case of incompatibility between the guest VM and the target cluster, VM may not work after migration to the target.

  For more information about supported guest operating systems on AHV, refer to Compatibility and Interoperability Matrix.
- VMs with volume groups attached.
- GenID and BIOS ID will not be retained.
- Migration of AHV-Native VMs such as File Server VMs and Objects VMs is not supported/qualified in Move.

Limitations

The section lists the limitations for the migration.

Migration Limitations

The support for migration is constrained by the following:

- After migration, Windows VMs with connected NICs only will retain their IP address. Those with disconnected NICs will not retain their IP address.
- Migration of a source VM which is in a VPC-based network is not supported.
- (For AHV to NC2) Retention of static IP address is not supported for VM migrations.

Creating a Migration Plan

You can create a migration plan to seed the data, cutover, and monitor the VMs. You can create the migration plan in Move UI without initiating the cutover process.

Before you begin

Ensure that you have added the source and target AOS environments.

For more information, refer to Adding a Nutanix AOS Cluster Environment on page 39

About this task

Note:

- This procedure is applicable for the following migrations only:
  - AHV to AHV
  - AHV to NC2 (on AWS/Azure)
  - NC2 (on AWS/Azure) to AHV
  - NC2 (on Azure) to NC2 (on Azure)
- If you are logging in for the first time, log on to the Move UI with your default credentials.
- You must have admin user credentials to complete the migration process.
- If you restart the management server, scheduled VM migration does not begin automatically.
- While migrating Prism Central VM, the DHCP IP address of the Prism Central VM is not retained post migration, and you have to reconfigure the IP address. IP address must be same before and after the migration for proper connectivity between the Prism Central and the Prism Element.
- If you cancel or discard the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.

  For information on canceling an ongoing migration, see Pausing or Canceling a VM Migration on page 257.

  For information on performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

To create a migration plan, do the following:

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.
3. On the Move dashboard, click **Create a Migration Plan**.

   Note: If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.
4. Enter the new migration plan name, and then click **OK**.

   Note: You can edit the migration plan name by clicking the pencil icon next to the migration plan name.

   The **New Migration Plan** window appears.

5. Complete the following fields, and then click **Next**.

   a. **Select a Source**: Select any AHV source for migration.

   If you select Prism Central IP address, a new field **Select Cluster** appears to select any cluster of that PC.

   Once you select the source, an appropriate target appears.

> **Note:** You might at times see a message relating to inventory collection as shown below. During this time, the environments undergoing Refresh will not be available for selection. Such environments do not automatically show up for selection once the inventory collection is over. In case you wish to select one of the environments undergoing Refresh (not available for selection), you will need to select **Cancel** and wait for inventory collection to get over and then create the migration plan again.

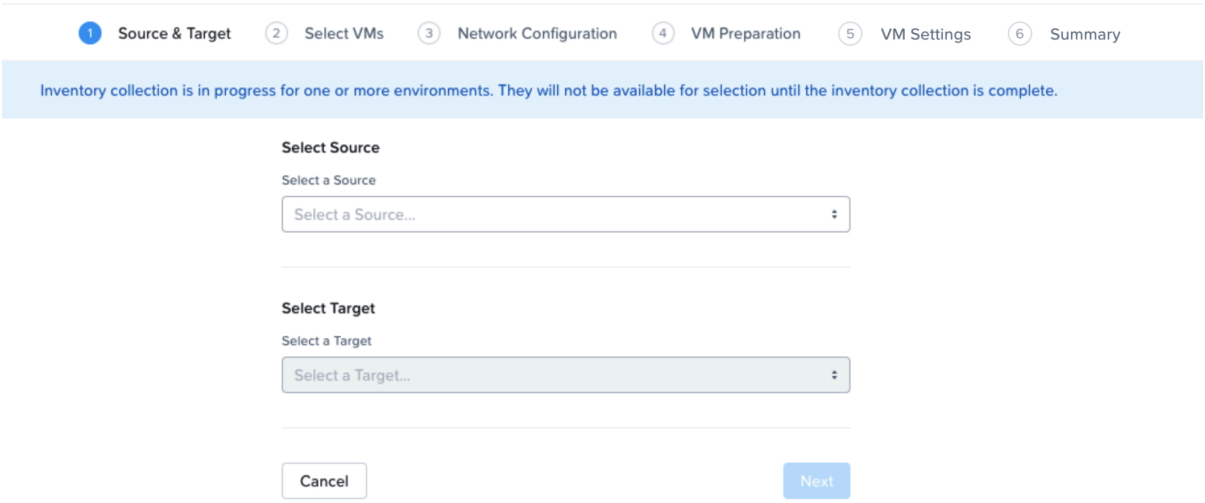

Figure 53: Inventory Collection Message

   b. **Select a Target**: Select the target AHV cluster for the migrating VMs.

   c. **Target Project** (optional): Select the project you want as the target.
   This field is available only with Prism Central.

   d. **Target Owners**: Select the owners for the selected target project.
   This field is available only when a target project is selected.

   e. **Target Cluster**: Select the cluster on which you will migrate the VMs.

   f. **Target Containers**: Select the container to which you are migrating the VMs.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click **+Add All**, and then click **Next**.

   Note:

   - You cannot add more than 50 VMs in a single migration plan.
   - Migration of VMs which have NIC(s) attached to VPCs is not supported.

   You can filter the VMs by name in the **Filter** bar. You can also sort the VM types by selecting the appropriate column.

   Note: The VMs for which migration has failed are not displayed. To show the entire list of VMs, select **All VMs** from the drop-down list. To show the configuration of VMs, select **Configuration** from the drop-down list. A question mark icon appears beside an unavailable VM that displays more information about that VM and indicate the reason of a failed VM migration.

   Note: Migrate VMs retain their power state on the destination cluster.

   The selected VMs are displayed in the sidebar.

7. In the **Network Configuration** screen, select the target network from the drop-down.

   - [Applicable only if Prism Central is used] Move provides the option to select VPC-based or VLAN-based target subnets. Based on the selection of a VPC or VLAN ID as the target network, the respective subnets are listed in the target subnet drop-down menu. Select the required subnet from the drop-down menu.

     Note: Overlay subnets which do not have IP address pool(s) associated will be disabled in the subnet drop-down menu.

   Click **Next**.

8. In the **VM Preparation** screen, select one of the following VM preparation modes.

   Note: Applicable to AHV to NC2 migration only.

   - You can skip the VM preparation options.
   - (For automatic preparation mode) Ignore the warning from Move to enter the login credentials.

   » **Automatic**. Move automatically runs scripts on the source VMs to prepare them for migration. Provide the credentials of the source VMs under **Windows VMs** or **Linux VMs**, depending on the type of the source VM.

   Note:

   - For Windows VMs, Move supports only username-password sign-in option for authentication. It does not support username-PIN sign-in option.
   - Currently, the default location where the preparation scripts are stored is the /tmp folder.

     If the /tmp folder is mounted as noexec, then Move will fallback to the /var/tmp folder. If the /var/tmp folder is also mounted as noexec, then Move will fallback to the /usr/tmp folder.

   For more information, refer to Automatic VM Preparation on page 198.

   » **Manual**. Move displays the VM preparation scripts for Windows and Linux VMs. To manually download and run the migration preparation software, select this option, and then run the scripts provided in the **VM Preparation** screen on the respective source VMs.

   Note: Ensure to run the VM preparation script on all selected VMs. If not, Move will only migrate VM data. All operating system configuration options will be bypassed.

   » **Mixed**. Move allows you to select the VM preparation mode for each VM. This setting allows you to prepare one set of VMs manually and prepare another set of VMs automatically in the same migration plan. If you want to have such a hybrid combination of **Automatic** and **Manual** VM preparation modes, proceed to the step 8.

   Note: The preparation mode automatically switches to **Mixed** if you change the mode of preparation for a set of VMs under **Change Settings** and have a mix of **Automatic** and **Manual** VM preparation modes. You cannot manually select the **Custom** option from the **Preparation Mode** drop-down list.

9. In the **Override individual VM Preparation** section, click **Change Settings** to override the **Guest Operations** settings (configured in the above steps) for the individual VMs. You can also edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

   Note: If you do any of the following for a VM, then copy the new generated scripts of that specific VM and run them on the source VM:

   - Change the **Mode of Preparation** of a VM to **Manual**.
   - Change any of the guest operation settings of a VM with the preparation mode set to **Manual** (an icon appears next to the VM Name prompting to regenerate a new guest script).

   Access the newly generated VM preparation script for that VM by selecting the Copy Scripts icon under the **Actions** column.

   After making the required changes, click **Done**.

10. (Optional) In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

    a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

    b. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

    Only those categories which have values are available for selection.

    c. **VM Migration Type**: Select one of the following VM migration types.

    At the VM level, some of the target VM properties can be customized manually after the migration plan is created. For information on manually customizing the target VM configuration, refer to Customizing the Target VM Configuration on page 255.

    - **Configure Target VM Properties**: The target VM synchronizes with the source VM properties at the time of migration plan creation. Selecting this option allows you to edit the target VM properties at the VM level (during migration). For information on editing the target VM properties, refer to Customizing the Target VM Configuration on page 255.

      Note: At the VM level, only the following properties can be edited:

      - Target VM name
      - Number of vCPUs
      - Number of cores per vCPU
      - Memory
      - Power state

    - **Retain Source VM Properties**: The target VM synchronizes with the source VM properties whenever Move refreshes the source VM configuration details. Only the customizable properties are refreshed on the target. Selecting this option does not allow you to edit the target VM properties at the VM level.

      Note:

      - The source VM properties are refreshed in the following ways.
        - (Manually) When you click the **Refresh Source VM Properties** button.
        - (Automatically) When you start a migration plan.
        - (Automatically) When you initiate a cutover.
      - When you start a migration plan, Move refreshes both source VM and target VM properties by default. However, it will not refresh the target VM properties at the start of a migration plan if you modified the target VM properties after migration plan creation.

    d. **Settings for Individual VMs**: Click **Change Settings** to configure settings such as Instance Type settings and VM priority for individual VMs.

    e. **Schedule Data Seeding**: Select this checkbox to select the date and time for migration.

11. In the **Summary** screen, choose one of the following, and then proceed to review the VM migration summary.

    » **Back**: To edit the information, click this option.

    » **Save**: To save the migration plan, click this option.

    For more information about how to start the migration later, and check the migrated VM status and details, refer to Environments and Migration Plan Management on page 259.

    » **Save and Start**: Click this option to save the migration plan and begin the migration immediately

    Once you save and proceed, the seeding process for migration begins.

    Note: This process takes some time.

    Note: The seeding process can take several minutes depending on the number of VMs.

What to do next

- To customize the target VM configuration at the VM level, refer to Customizing the Target VM Configuration on page 255.
- If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover on page 201.

Automatic VM Preparation

You can automate the guest VM preparation.

About this task

Note:

- Before automatic Windows VM preparation, enable WinRM and enable the ports.
- Automatic VM preparation does not work with the Windows domain account. Use the Windows built-in administrator credentials for the Windows VMs.

To automatically prepare the VMs, do the following:

Procedure

1. In the **Preparation Mode** drop-down, select **Automatic**.
2. In the **Credentials for Source VMs** section, enter the user name and password for the guest VMs to allow Move to install the necessary drivers.
3. (Optional) In the **VM Preparation** screen, do the following, and then click **Next**.
    a. In the **Override individual VM Preparation** section, click **Change Settings** to override the same settings for the individual VMs. You can update the VM preparation credentials of the VMs, remove VMs, or update the VM preparation mode.

4. (Optional) In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

   a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

   b. **Category Settings (Optional)**: Select the categories to which the target VM(s) should be assigned.

      Only those categories which have values are available for selection.

   c. **Settings for Individual VMs**: Click **Change Settings** to configure settings such as Instance Type settings and VM priority for individual VMs.

   d. **Schedule Data Seeding**: Select this checkbox to select the date and time for migration.

   The credentials of VMs are validated. Once the validation is successful, the Guest Tools are downloaded and installed in all the VMs of the migration plan. Then, the VMs are validated for readiness.

   > **Note:** If the validation of credentials or Guest Tools installation fails, you can update the credentials or remove the VM from the migration plan and proceed by clicking **Next**.

What to do next

If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover on page 201

Enabling WinRM

You need to enable WinRM to install the Guest Tools on AHV VMs.

Before you begin

Ensure that the ingress ports 5985 and 5986 are enabled.

About this task

> **Note:** This method is a prerequisite for Automatic VM preparation to work with Windows source VMs.

To enable WinRM.

Procedure

1. Open PowerShell in Windows VM.
2. Run the script to enable WinRM on AHV VMs.

```
> winrm quickconfig -q
winrm set winrm/config/winrs '@{MaxMemoryPerShellMB="300"}'
winrm set winrm/config '@{MaxTimeoutms="1800000"}'
winrm set winrm/config/service '@{AllowUnencrypted="true"}'
winrm set winrm/config/service/auth '@{Basic="true"}'

netsh advfirewall firewall add rule name="WinRM 5985" protocol=TCP dir=in
 localport=5985 action=allow
netsh advfirewall firewall add rule name="WinRM 5986" protocol=TCP dir=in
 localport=5986 action=allow

net stop winrm
cmd /c 'sc config winrm start= auto'
net start winrm
```

3. Run the script to enable Secure Sockets Layer (SSL).

```
> $c = New-SelfSignedCertificate -DnsName "$(hostname)" -CertStoreLocation cert:
\LocalMachine\My
winrm create winrm/config/Listener?Address=*+Transport=HTTPS
 "@{Hostname=`"$(hostname)`";CertificateThumbprint=`"$($c.ThumbPrint)`"}"
```

Performing Data-Only Migration

Move performs data-only migration when you select **Automatic** preparation mode while creating a migration plan, and bypass the guest operations or does not provide the source VM credentials while preparing a migration plan. Or when you select **Manual** preparation mode while creating a migration plan, and do not run the preparation script in the source VMs. In data-only migration, Move skips the source VM guest operating system preparation tasks which includes copying of the scripts to retain the IP address.

About this task

Note: Data-only migration is only supported for the following migrations:

- From ESXi to AHV and ESXi to NC2 on AWS
- From Hyper-V to AHV and Hyper-V to NC2 on AWS
- From Hyper-V to ESXi
- From AHV to AHV

To perform data-only migration, do the following:

Procedure

1. In the **VM Preparation** screen, if you select **Automatic**, then proceed without providing the credentials for the source VMs or select the **Bypass Guest Operations on Source VMs** check box or if you select **Manual**, do not run the preparation script in the source VMs.
   The following message appears when the **Automatic** preparation mode is selected,

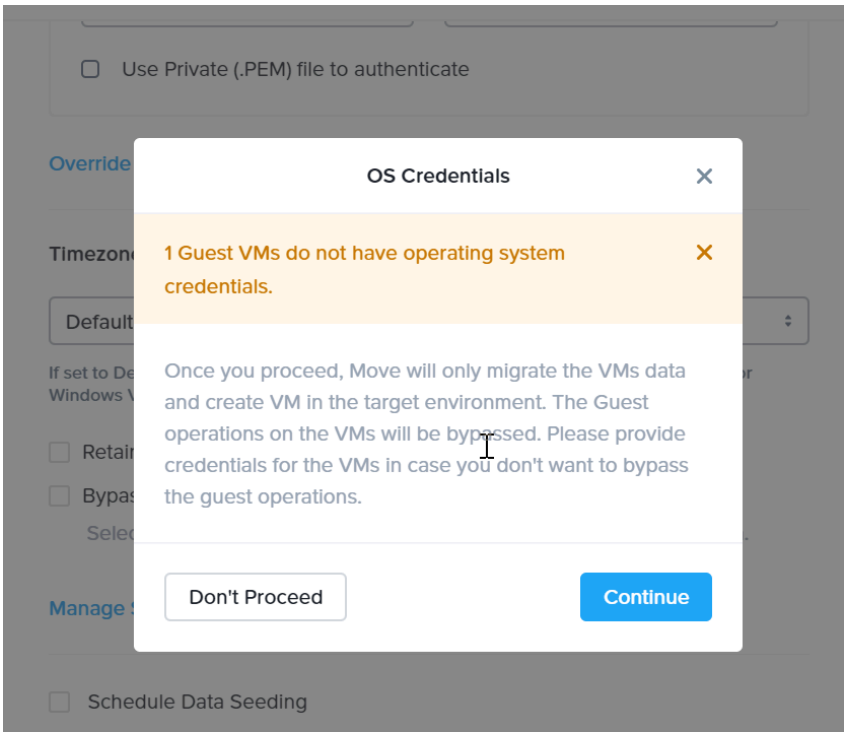

Figure 54: OS Credentials Dialog Box

2. Click **Continue**.
   Move migrates the VMs data and creates a VM in the target, and bypasses the operating system operations.

Performing a Migration Cutover

When the migration plan is started and the seeding process is complete, you can cutover the selected VMs to the AHV cluster. You can monitor the VM migration progress by clicking the **Status** link.

About this task

Note:

- You can perform this operation only when the VM status is Ready to Cutover.
- Recommended that cutover should be performed within one week of initial data seeding.
- If the initial data seeding finishes in less than 10 minutes, the Move VM continues to wait for 10 minutes to take the incremental snapshot; however, you can trigger the cutover immediately.
- The cutover process increments in absolute numbers.

To perform the migration cutover, do the following:

Procedure

1. In the Move UI, click the **Ready to Cutover** status to display the list of available VMs.
2. To perform a cutover, select the VMs or group of VMs.
3. Click **Cutover**.
   The cutover process performs the following VM actions.

   - Shuts down the VM
   - Takes the final snapshots for the VM and copying the final changes to AHV
   - Adds a note in the VM in the AHV cluster
   - Disconnects the source VM network interfaces
   - Creates a VM in the target AHV cluster
   - Attaches the replicated disks to the VM
   - Powers on or off the VM (depends on the initial power state)
   - Runs the scripts to set the static IP address

   The cutover process begins immediately and takes a few minutes. Once cutover is complete, the VM is ready for use in the new AHV cluster.

What to do next

You can manage the migration plan once the migration plan is ready. For more information, refer to Environments and Migration Plan Management on page 259

AHV TO AWS

To provide you the flexibility to migrate some workloads between private and public cloud, Move has introduced AHV to AWS migration. You can now prepare and migrate AHV VMs to AWS by using Move.

Migration Considerations

You must consider the supported guest operating systems, requirements, recommendations, unsupported features, and limitations provided in this section before starting the migration process.

Supported Guest Operating Systems (AHV to AWS)

Move supports some common operating systems. Unless otherwise specified, Nutanix has qualified the following 64-bit guest operating system versions.

Fully Supported

- Windows Server 2012 R2, 2016, 2019, 2022
- RHEL 6.8–6.10, 7.3–7.5
- CentOS 6.8–6.10, 7.3–7.5
- Ubuntu 14.04, 16.04, 18.04

Requirements (AHV to AWS)

Before attempting to migrate VMs running on AHV using Move, make sure to conform to the requirements listed here.

General Requirements

Ensure to conform to the following requirements for AHV to AWS migration.

- Supported browser: Google Chrome
- Ensure you have PowerShell version 4.0 or later.
- Ensure TCP 443 connection to AWS endpoint for operations in AWS.
- For Windows source VMs, ensure to disable UAC for Windows administrator user.
- AHV should support V3 API for AHV to AWS migrations.
- Network Security Services (NSS) 3.44 is required for Linux VMs.
- Before you initiate a migration, ensure to disable backups.

Prerequisites for Linux guest VMs:

- SSH service should be up and running.
- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

Requirements for Modules or Drivers

Make sure that the following modules and drivers are installed in Windows, Linux, and AWS VMs:

For Windows:

- Enable WinRM for automatic preparation.

  Ensure to enable the following inbound and outbound ports using TCP protocol for the Windows Remote Management (WinRM) feature to work.

  - WinRM-HTTPS: 5986
  - WinRM-HTTP: 5985
  - RDP: 3389 (only for inbound)
  - SSH: 22
- Install Nutanix Guest Tools (NGT)

For Linux:

- Install Linux-AWS for Ubuntu 14.04, 16.04, 18.04 (apt-get install linux-aws)
- Install Nutanix Guest Tools (NGT)
- SSH should be enabled with root user privilege for automatic preparation

AWS

- The AWS account provided while adding an AWS target must have the set of permissions as provided in the following JSON to do end-to-end migration.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ec2:*KeyPair*",
        "ec2:RunInstances",
        "ec2:TerminateInstances",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateVolume",
        "ec2:DeleteVolume",
        "ec2:AttachVolume",
        "ec2:DetachVolume",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ec2:ModifyInstanceAttribute",
        "ebs:ListSnapshotBlocks",
        "ebs:StartSnapshot",
        "ebs:PutSnapshotBlock",
        "ebs:CompleteSnapshot",
        "ec2:DeleteSnapshot",
        "ec2:RegisterImage",
        "ec2:DeregisterImage"
      ],
      "Resource": "*"
    }
  ]
}
```

- The AWS account provided while adding AWS as both source and target must have the set of permissions as provided in the following JSON to do end-to-end migration.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
        "ebs:GetSnapshotBlock",
        "ec2:*KeyPair*",
        "ec2:RunInstances",
        "ec2:TerminateInstances",
        "ec2:CreateVolume",
        "ec2:DeleteVolume",
        "ec2:AttachVolume",
        "ec2:DetachVolume",
        "ec2:ModifyInstanceAttribute",
        "ebs:StartSnapshot",
        "ebs:PutSnapshotBlock",
        "ebs:CompleteSnapshot",
        "ec2:RegisterImage",
        "ec2:DeregisterImage"
      ],
      "Resource": "*"
    }
  ]
}
```

- Move should have all the permissions listed above for the following regions.

  Note: There can be restrictions for certain regions. For more information, refer to the next point.

```
"af-south-1"     // Africa (Cape Town).
"ap-east-1"      // Asia Pacific (Hong Kong).
"ap-northeast-1" // Asia Pacific (Tokyo).
"ap-northeast-2" // Asia Pacific (Seoul).
"ap-northeast-3" // Asia Pacific (Osaka).
"ap-south-1"     // Asia Pacific (Mumbai).
"ap-south-2"     // Asia Pacific (Hyderabad).
"ap-southeast-1" // Asia Pacific (Singapore).
"ap-southeast-2" // Asia Pacific (Sydney).
"ap-southeast-3" // Asia Pacific (Jakarta).
"ap-southeast-4" // Asia Pacific (Melbourne).
"ca-central-1"   // Canada (Central).
"eu-central-1"   // Europe (Frankfurt).
"eu-central-2"   // Europe (Zurich).
```

```
"eu-north-1"     // Europe (Stockholm).
"eu-south-1"     // Europe (Milan).
"eu-south-2"     // Europe (Spain).
"eu-west-1"      // Europe (Ireland).
"eu-west-2"      // Europe (London).
"eu-west-3"      // Europe (Paris).
"me-central-1"   // Middle East (UAE).
"me-south-1"     // Middle East (Bahrain).
"sa-east-1"      // South America (Sao Paulo).
"us-east-1"      // US East (N. Virginia).
"us-east-2"      // US East (Ohio).
"us-west-1"      // US West (N. California).
"us-west-2"      // US West (Oregon).
```

- Policies needed for explicit deny by region:

  Move should always have the following permissions for the region `us-east-1`:

  - `iam:GetUser`
  - `iam:SimulatePrincipalPolicy`

  Examples of the JSON are provided below.

  When restricting access to specific regions, the AWS account provided while adding an AWS target must have the set of permissions as provided in the following JSON to do end-to-end migration.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ec2:*KeyPair*",
        "ec2:RunInstances",
        "ec2:TerminateInstances",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateVolume",
        "ec2:DeleteVolume",
        "ec2:AttachVolume",
        "ec2:DetachVolume",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ec2:ModifyInstanceAttribute",
        "ebs:ListSnapshotBlocks",
        "ebs:StartSnapshot",
        "ebs:PutSnapshotBlock",
        "ebs:CompleteSnapshot",
        "ec2:DeleteSnapshot",
        "ec2:RegisterImage",
        "ec2:DeregisterImage"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "<AWS-REGION-1>",
```

```
            "<AWS-REGION-2>",
            ...
          ]
        }
      }
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": [
        "iam:GetUser",
        "iam:SimulatePrincipalPolicy"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "us-east-1",
            "<AWS-REGION-1>",
            "<AWS-REGION-2>",
            ...
          ]
        }
      }
    }
  ]
}
```

When restricting access to specific regions, the AWS account provided while adding AWS as source and target must have the set of permissions as provided in the following JSON to do end-to-end migration.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ec2:CreateSnapshots",
        "ec2:DeleteSnapshot",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
        "ebs:GetSnapshotBlock",
        "ec2:*KeyPair*",
        "ec2:RunInstances",
        "ec2:TerminateInstances",
        "ec2:CreateVolume",
        "ec2:DeleteVolume",
        "ec2:AttachVolume",
        "ec2:DetachVolume",
        "ec2:ModifyInstanceAttribute",
        "ebs:StartSnapshot",
        "ebs:PutSnapshotBlock",
```

```
        "ebs:CompleteSnapshot",
        "ec2:RegisterImage",
        "ec2:DeregisterImage"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "<AWS-REGION-1>",
            "<AWS-REGION-2>",
            ...
          ]
        }
      }
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": [
        "iam:GetUser",
        "iam:SimulatePrincipalPolicy"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "us-east-1",
            "<AWS-REGION-1>",
            "<AWS-REGION-2>",
            ...
          ]
        }
      }
    }
  ]
}
```

- Policies needed for explicit deny by IP address/CIDR:

  When restricting access using IP address/CIDR, the AWS account provided while adding an AWS target must have the set of permissions as provided in the following JSON to do end-to-end migration.

  This following JSON is an example. Update the JSON as necessary based on the security policies in your organization.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AwsTgtPermsWithRestrictedIp",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy",
        "iam:GetUser",
        "ec2:*Describe*",
        "ssm:DescribeInstanceInformation",
        "ec2:*KeyPair*",
        "ec2:RunInstances",
        "ec2:TerminateInstances",
        "ec2:StopInstances",
        "ec2:CreateSecurityGroup",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
```

```
        "ec2:StartInstances",
        "ec2:CreateVolume",
        "ec2:DeleteVolume",
        "ec2:AttachVolume",
        "ec2:DetachVolume",
        "ec2:CreateTags",
        "ec2:DeleteTags",
        "ec2:ModifyInstanceAttribute",
        "ebs:ListSnapshotBlocks",
        "ebs:StartSnapshot",
        "ebs:PutSnapshotBlock",
        "ebs:CompleteSnapshot",
        "ec2:DeleteSnapshot",
        "ec2:RegisterImage",
        "ec2:DeregisterImage"
      ],
      "Resource": "*",
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": [
            "<CIDR-BLOCK-1>",
            "<CIDR-BLOCK-2>",
            ...
          ]
        },
      "Bool": {
        "aws:ViaAWSService": "false"
        }
      }
    }
  ]
}
```

Note:

- A CIDR block is a group of IP addresses that share the same network prefix and have the same number of bits.

  Example: `192.168.x.x/29`

- In the above JSON, replace `<CIDR-BLOCK-x>` with the appropriate CIDR block.

- No explicit deny policies should be defined for the account.

  Explicit deny policies restrict access to AWS based on various parameters such as source IP address, VPC, VPC endpoint, and so on. For example, explicit deny for source IP address denies access to AWS when a request comes from an IP address outside the specified range.

Service Accounts

For successful migration of VMs from AHV to AWS, Move requires the following.

- Prism Element for the AHV cluster
- An administrator for Windows source VMs or a root for Linux source VMs to run the source VM preparation scripts.

Qualified Metrics (AHV to AWS)

The section lists the qualified metrics for migration from AHV to AWS.

The following metrics are qualified for migration from AHV.

- Migration of 10 VMs where each VM has 7 disks in a single migration plan
- 3.5 TB VM migration
- 25 VMs migration in a single migration plan

Unsupported Features (AHV to AWS)

This section lists the unsupported features for migration from AHV to AWS.

- Guest operating systems other than the supported operating systems.

  For more information, refer to Supported Guest Operating Systems for AWS Migration
- AHV clusters with protection domain and stretched containers
- VMs with GPUs (Instance type should be updated for such VMs).
- Clusters with encryption enabled
- Clusters with multi-homing setups
- Increasing the disk size of the source VM
- Migration of non-English VMs, such as Japanese is not qualified for AHV to AWS.
- IP address and MAC address retention

Limitations (AHV to AWS)

The section lists the limitations for migration from AHV to AWS.

AHV to AWS Migration Limitations

The support for migration is constrained by the following.

- Time taken for migration depends on the size of the VM, data churn rate within the VM during migration, and Internet connectivity between on-prem data center and AWS.
- If you have multiple NICs, Move gives an option to select multiple Virtual Private Cloud (VPCs) in the user interface. However, AWS supports only one VPC to a VM.
- Migration of a source VM which is in a VPC-based network is not supported.

Adding a Nutanix AOS Cluster Environment

While creating a migration plan, AHV AOS cluster can be added as both source or target. If you want to use ESXi on Nutanix cluster as target, then add the corresponding AOS cluster environment for ESXi.

About this task

> **Note:** When you add a AOS cluster, Move VM IP address with the subnet 255.255.255.255 is added to the global NFS allowlist.
>
> > **Caution:** Modifying the NFS allowlist of the destination container disables inheritance of the NFS allowlist at the global level and lead to issues with Move operations.

To add a Nutanix AOS cluster environment, do the following:

Procedure

1. Log on to the Move UI.

2. Click **+ Add Environment** under **Environments**.

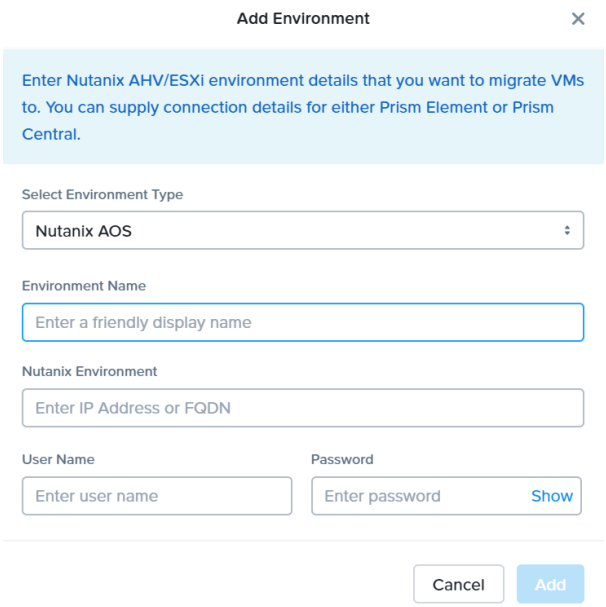

Figure 55: Add AOS Environment Dialog Box

The **Add Environment** window appears.

3. Select **Nutanix AOS** as the target environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the NC2 on AWS and AHV or VMware ESXi on Nutanix environment.
   b. **Nutanix Environment**: Enter the IP address or the FQDN for the target Prism Central or Prism Element.
   c. **User Name**: Enter the username for logging on to the target Nutanix environment.
   d. **Password**: Enter the password for logging on to the target Nutanix environment.

5. (Only for ESXi to ESXi migration) Enter credentials for registered vCenter.

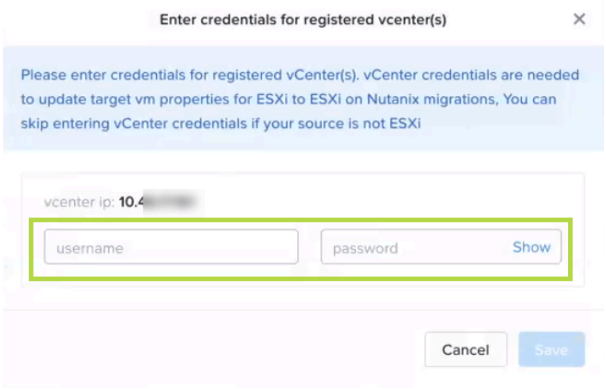

Figure 56: vCenter Credentials Dialog box

vCenter credentials are required to update the target VM properties for ESXi to ESXi on Nutanix migration.

Note:

- You can skip adding vCenter credentials if your source is not ESXi.
- To retain the VM properties on the target VM after migration, be sure to provide the target vCenter credentials. The following properties will be retained:
  - SCSI controller types
  - Network adaptor type
  - MAC address
  - Video card
  - Memory overcommit variables
  - Tool upgrade flag
  - Sync time with host flag
  - Disable acceleration flag
  - Enable logging flag

The AOS environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

Once you have added the environments, you can proceed to create the migration plan. For more information, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan (Hyper-V to ESXi) on page 103 or Creating a Migration Plan (Azure to AHV) on page 171 or Creating a Migration Plan (Azure to ESXi) on page 184

Adding an AWS Environment

While creating a migration plan, if you need to add an AWS source or target, you have to add at least one AWS environment for migration.

About this task

> **Note:** This procedure is only applicable for migration to and from an AWS environment.

To add an AWS environment, do the following:

Procedure

1. Log on to Move UI.
2. Click **+ Add Environment** under Environments.
   The **Add Environment** appears.

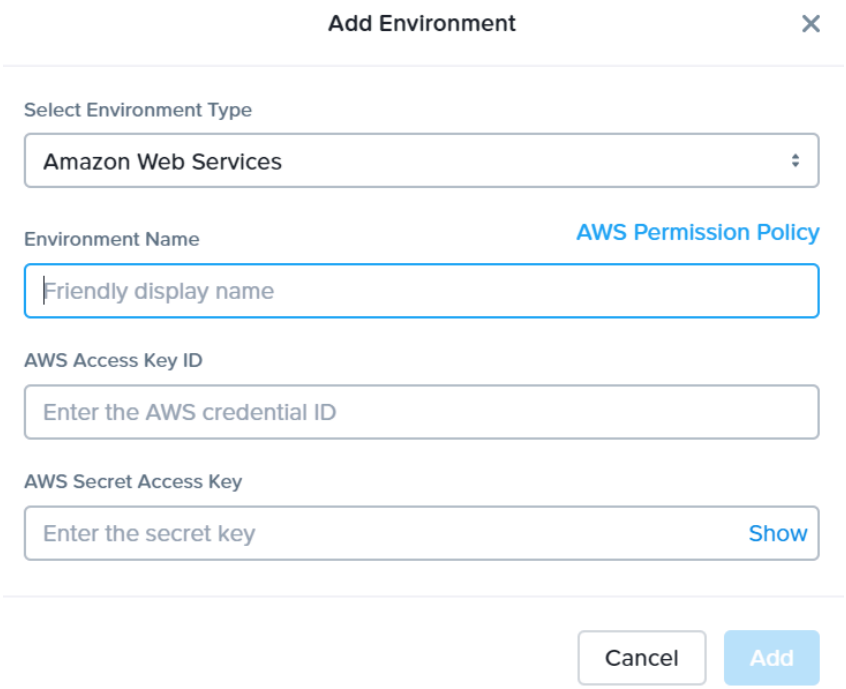

Figure 57: Add AWS Environment Dialog Box

3. Select **Amazon Web Services** as the source environment type.
4. Complete the indicated fields and click **Add**.
   a. **Source Name**: Enter a name for the AWS environment.
   b. **AWS Access Key ID**: Enter the access key ID of the AWS account.
   c. **AWS Secret Access Key**: Enter the key for logging on to AWS account.

   The source environment is added to Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

You can add a Nutanix AOS cluster environment. For more information, refer to Adding a Nutanix AOS Cluster Environment on page 39

Creating a Migration Plan (AHV to AWS)

You can create a migration plan to seed the data, cutover, and monitor the VMs. You can create the migration plan in Move without initiating the cutover process.

About this task

Note:

- This procedure is only applicable for migration from AHV to AWS.
- On your first log on, you can log on to the Move UI with your defaults credentials.
- If you cancel or discard the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.

  For information on canceling an ongoing migration, see Pausing or Canceling a VM Migration on page 257.

  For information on performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

To create a migration plan, do the following:

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.
3. On the Move dashboard, click **Create a Migration Plan**.

   Note: If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.
4. Enter the new migration plan name, and then click **OK**.

   Note: You can edit the migration plan name by clicking the pencil icon next to the migration plan name.

   The **New Migration Plan** window appears.

5. Complete the following fields, and then click **Next**.

   a. **Select a Source**: Select any AHV source for migration.

   If you select Prism Central IP address, a new field **Select Cluster** appears to select any cluster of that PC.

   Once you select the source, an appropriate target appears.

> **Note:** You might at times see a message relating to inventory collection as shown below. During this time, the environments undergoing Refresh will not be available for selection. Such environments do not automatically show up for selection once the inventory collection is over. In case you wish to select one of the environments undergoing Refresh (not available for selection), you will need to select **Cancel** and wait for inventory collection to get over and then create the migration plan again.

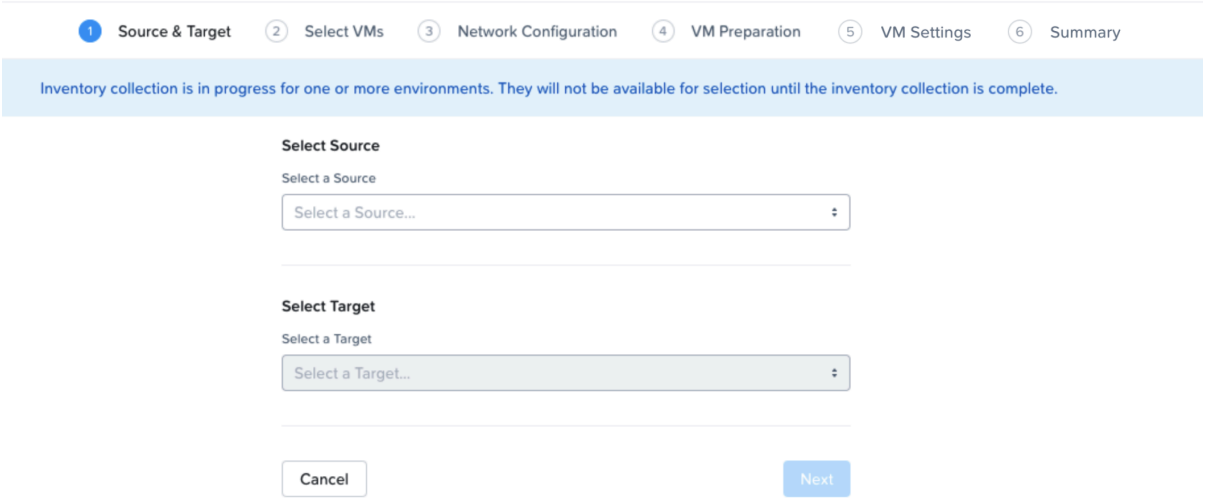

Figure 58: Inventory Collection Message

   b. **Select a Target**: Select the target AWS instance for the migrating VMs.

   c. **Target Region**: To migrate the VMs, select a region.

> **Note:** Only regions with Internet gateway VPCs are available for selection.

   d. **Target Availability Zone**: Select a availability zone to create the target instance. Availability zones are region specific.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click **+Add All**, and then click **Next**.

   Note:

   - You cannot add more than 50 VMs in a single migration plan.
   - Migration of VMs which have NIC(s) attached to VPCs is not supported.

   You can filter the VMs by name in the **Filter** bar. You can also sort the VM types by selecting the appropriate column.

   Note: The VMs for which migration has failed are not displayed. To show the entire list of VMs, select **All VMs** from the drop-down list. To show the configuration of VMs, select **Configuration** from the drop-down list. A question mark icon appears beside an unavailable VM that displays more information about that VM and indicate the reason of a failed VM migration.

   Note: Migrate VMs retain their power state on the destination cluster.

   The selected VMs are displayed in the sidebar.

7. In the **Network Configuration** screen, select the following fields, and then click **Next**.

   a. **Target Network**: Select the network that target instance will use after migration.

   b. **Subnet**: Select the subnet for the target network.

   c. **Security Group**: Select the security group(s) for the target network. A maximum of five security groups can be selected for a target network.

8. In the **VM Preparation** screen, select one of the following VM preparation modes.

   » **Automatic**. Move automatically runs scripts on the source VMs to prepare them for migration. Provide the credentials of the source VMs under **Windows VMs** or **Linux VMs**, depending on the type of the source VM.

   Note: For Windows VMs, Move supports only username-password sign-in option for authentication. It does not support username-PIN sign-in option.

   For more information, refer to Automatic VM Preparation (AHV to AWS) on page 217.

   » **Manual**. Move displays the VM preparation scripts for Windows and Linux VMs. To manually download and run the migration preparation software, select this option, and then run the scripts provided in the **VM Preparation** screen on the respective source EC2 instance.

   These scripts prepare the instance by performing the following installations.

   - For Windows: AWS Paravirtualization (PV) driver, AWS EC2Config, and Elastic Network Adapter (ENA) driver

   - For Linux: ENA driver and Intel Driver

   Note: Ensure to run the VM preparation script on all selected VMs. If not, Move will only migrate VM data and not start the VM in the target environment. Also, all operating system configuration options will be bypassed.

   » **Mixed**. Move allows you to select the VM preparation mode for each VM. This setting allows you to prepare one set of VMs manually and prepare another set of VMs automatically in the same migration plan.

If you want to have such a hybrid combination of **Automatic** and **Manual** VM preparation modes, proceed to the step 8.

> **Note:** The preparation mode automatically switches to **Mixed** if you change the mode of preparation for a set of VMs under **Change Settings** and have a mix of **Automatic** and **Manual** VM preparation modes. You cannot manually select the **Custom** option from the **Preparation Mode** drop-down list.

9. In the **Override individual VM Preparation** section, click **Change Settings** to override settings for the individual VMs. You can edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

   > **Note:** If you change the **Mode of Preparation** to **Manual** for a VM, then copy the new generated scripts of that specific VM and run them on the source VM.
   >
   > Access the newly generated VM preparation script for that VM by selecting the Copy Scripts icon under the **Actions** column.

   After making the required changes, click **Done**.

   Then, click **Next**.

10. In the **VM Settings** screen, do one or more of the settings, and then click **Next**.

    a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

    b. **Create Public IP Address**: Select this option if you want to create public IP address for all the VMs in the migration plan.

       > **Note:** There must be no policies on AWS that block the creation of VM NIC with public IP address. Otherwise, the cutover might fail when creating VMs with public IP address.

    c. **Settings for Individual VMs**: Click **Change Settings** to configure settings such as Instance Type settings and VM priority for individual VMs.

    d. **Schedule Data Seeding**: Check this check box to select the date and time for migration.

11. In the **Summary** screen, choose one of the following, and then proceed review the VM migration summary.

    » **Back**: To edit the information, click this option.

    » **Save**: To save the migration plan, click this option.

      For more information about how to start the migration later, and check the migrated VM status and details, refer to Environments and Migration Plan Management on page 259.

    » **Save and Start**: Click this option to save the migration plan and begin the migration immediately

    Once you save and proceed, the seeding process for migration begins.

    > **Note:** The seeding process can take several minutes depending on the number of VMs.

What to do next

If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover (AHV to AWS) on page 219

Automatic VM Preparation (AHV to AWS)

You can automate the guest VM preparation.

About this task

Note:

- Before automatic Windows VM preparation, enable WinRM and enable the ports.
- Automatic VM preparation does not work with the Windows domain account. Use the Windows built-in administrator credentials for the Windows VMs.

To automatically prepare the VMs, do the following:

Procedure

1. In the **Preparation Mode** drop-down, select **Automatic**.
2. In the **Credentials for Source VMs** section, enter the username and password for the guest VMs to allow Move to install the necessary drivers.
3. (Optional) In the **VM Preparation** screen, do the following, and then click **Next**.
   a. In the **Override individual VM Preparation** section, click **Change Settings** to override the settings for the individual VMs. You can edit the VM preparation credentials, remove VMs, or update the VM preparation mode.
4. (Optional) In the **VM Settings** screen, do one or more of the settings, and then click **Next**.
   a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.
   b. (Optional) **Settings for Individual VMs**: Click **Change Settings** to configure settings such as Instance Type settings and VMs Priority for individual VMs.
   c. **Schedule Data Seeding**: Check this check box to select the date and time for migration.

   The credentials of VMs are validated. Once the validation is successful, the Guest Tools are downloaded and installed in all the VMs of the migration plan. Then, the VMs are validated for readiness.

   Note: If the validation of credentials or Guest Tools installation fails, you can update the credentials or remove the VM from the migration plan and proceed by clicking **Next**.

What to do next

If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover (AHV to AWS) on page 219

Enabling WinRM (AHV-AWS)

You need to enable WinRM to install the Guest Tools on AHV VMs.

Before you begin

Ensure that the ingress ports 5985 and 5986 are enabled.

About this task

Note: This method is a prerequisite for Automatic VM preparation to work with Windows source VMs.

To enable WinRM.

Procedure

1. Open PowerShell in Windows VM.
2. Run the script to enable WinRM on AHV VMs.

```
> winrm quickconfig -q
winrm set winrm/config/winrs '@{MaxMemoryPerShellMB="300"}'
winrm set winrm/config '@{MaxTimeoutms="1800000"}'
winrm set winrm/config/service '@{AllowUnencrypted="true"}'
winrm set winrm/config/service/auth '@{Basic="true"}'

netsh advfirewall firewall add rule name="WinRM 5985" protocol=TCP dir=in
 localport=5985 action=allow
netsh advfirewall firewall add rule name="WinRM 5986" protocol=TCP dir=in
 localport=5986 action=allow

net stop winrm
cmd /c 'sc config winrm start= auto'
net start winrm
```

3. Run the script to enable Secure Sockets Layer (SSL).

```
> $c = New-SelfSignedCertificate -DnsName "$(hostname)" -CertStoreLocation cert:
\LocalMachine\My
winrm create winrm/config/Listener?Address=*+Transport=HTTPS
 "@{Hostname=`"$(hostname)`";CertificateThumbprint=`"$($c.ThumbPrint)`"}"
```

Performing a Migration Cutover (AHV to AWS)

When the seeding process is complete, you can cut over the selected VMs to the AWS target.

About this task

You can monitor the VM migration progress by clicking the status link.

Note:

- You can perform this operation only when the VM status is Ready to Cutover.
- Recommended that cutover should be performed within one week of initial data seeding.
- If the initial data seeding finishes in less than 10 minutes, Move continues to wait for 10 minutes to take the incremental snapshot; however, you can trigger the cutover immediately.
- The cutover process increments in absolute numbers.

To perform a cutover, do the following:

Procedure

1. In the Move UI, click the **Ready to Cutover** status to display the list of available VMs.
2. To cut over, select the VMs or group of VMs.

3. Click **Cutover**.

   The cutover process performs the following VM actions.

   - Shuts down the VM
   - Takes the final snapshots for the VM and copying the final changes to AWS
   - Creates an instance in the AWS target
   - Attaches replicated disks to the VM
   - Powers on or off the instance (depends on the initial power state)

   The cutover process begins immediately and might take a few minutes.

   Once the cutover is complete, the AWS instance is available. By default, the security group attached to the instance has all the ingress traffic blocked. To access the guest VMs, you need to update the security group according to your requirement.

What to do next

You can manage the migration plan once the migration plan is ready. For more information, refer to Environments and Migration Plan Management on page 259

Changing the Default Settings for Commit Data Size For Target Snapshot

You can now change the default settings for commit data size for the target snapshot.

To change the settings, open the srcagent.json and modify the following:

```
{
"AhvProviderConfig" : {
"MaxAWSCommitSize": 5368709120,
         "MaxAWSSnapshotCommitPercent": 20
}
}
```

Maximum of both these parameters are picked, where `MaxAWSCommitSize` is the size of the data commit in bytes and `MaxAWSSnapshotCommitPercent` is the percentage of source disk size.

AHV TO AZURE

To provide you the flexibility to migrate some workloads between private and public cloud, Move has introduced AHV to Azure migration. You can now prepare and migrate AHV VMs to Azure by using Move.

Migration Considerations

You must consider the supported guest operating systems, requirements, recommendations, unsupported features, and limitations provided in this section before starting the migration process.

Supported Guest Operating Systems (AHV to Azure)

Move supports some common operating systems. Unless otherwise specified, Nutanix has qualified the following 64-bit guest operating system versions.

Fully Supported

- Windows Server 2012 R2 DC, 2016, 2019, and 2022
- CentOS 6.8 to 6.10, 7.0 to 7.7, and 8.0 to 8.3
- RHEL 6.8 to 6.10, 7.0 to 7.7, and 8.0 to 8.5
- Ubuntu 14.0.4, 16.0.4, 18.0.4, and 20.0.4
- SLES 11 SP4, 12, and 15
- OEL 7.5 to 8.4

Supported Operating Systems for UEFI-Enabled VMs (AHV to Azure)

Move supports the following operating systems for UEFI-enabled VMs.

Table 19: Supported Operating Systems

Operating systems
Windows 2016
SLES 12 SP5
CentOS 7.4
Windows 2019
CentOS 7.3
RHEL 7.7

Support for UEFI with Secure Boot Enabled VMs (AHV to Azure)

Move supports UEFI with secure boot enabled VMs.

Table 20: Supported Guest Operating Systems

Operating systems
Windows 2019

Operating systems
CentOS 7.3
RHEL 7.7

Requirements (AHV to Azure)

Before attempting to migrate VMs running on AHV using Move, make sure to conform to the requirements listed here.

General Requirements

Following is the list of general requirements for AHV to Azure migration:

- Supported browser: Google Chrome
- Ensure you have PowerShell version 4.0 or later.
- Ensure TCP 443 connection to Azure endpoint for operations in Azure.
- For Windows source VMs, ensure to disable UAC for Windows administrator user.
- AHV should support V3 API for AHV to Azure migrations.
- Ensure to register the Azure app and apply the required privileges for the subscription. You can do it through Azure UI, refer to Registering an App and Applying Privileges (Azure UI) on page 156 or Move CLI, refer to Registering the App in Azure and Assigning Custom Role (Move CLI) on page 164.
- Network Security Services (NSS) 3.44 is required for Linux VMs.
- Move must have access to the following URLs:
  - https://management.azure.com/
  - https://login.microsoftonline.com/
  - https://graph.windows.net/
  - https://batch.core.windows.net/
  - https://*.blob.core.windows.net/ (Azure Storage blobs)
  - https://nxmove.blob.core.windows.net/
- Before you initiate a migration, ensure to disable backups.
- Move must have access to the shared access signature (SAS) URL https://*.blob.storage.azure.net/ generated by Azure for disk exports.

Prerequisites for Linux guest VMs:

- SSH service should be up and running.
- The credentials provided must have root or sudo user permission.
- Guest VM should have curl utility installed.

Requirements for Modules or Drivers

Make sure that the following modules and drivers are installed in Windows, Linux, and Azure VMs:

For Windows:

- Enable WinRM for automatic preparation.

  Ensure to enable the following inbound and outbound ports using TCP protocol for the Windows Remote Management (WinRM) feature to work.

  - WinRM-HTTPS: 5986
  - WinRM-HTTP: 5985
  - RDP: 3389 (only for inbound)
  - SSH: 22
- Install Nutanix Guest Tools (NGT)

For Linux:

- Install Nutanix Guest Tools (NGT)
- SSH should be enabled with root user privilege for automatic preparation

Azure

- The resource provider `Microsoft.compute` should be registered for the subscription.
- The Azure account provided while adding an Azure target must have the set of permissions as provided in the following JSON to do end-to-end migration.

```
{
  "permissions": [
    {
      "actions": [
        "Microsoft.Authorization/locks/read",
        "Microsoft.Authorization/roleAssignments/read",
        "Microsoft.Authorization/roleDefinitions/read",
        "Microsoft.Compute/disks/beginGetAccess/action",
        "Microsoft.Compute/disks/delete",
        "Microsoft.Compute/disks/endGetAccess/action",
        "Microsoft.Compute/disks/read",
        "Microsoft.Compute/disks/write",
        "Microsoft.Compute/locations/vmSizes/read",
        "Microsoft.Compute/virtualMachines/instanceView/read",
        "Microsoft.Compute/virtualMachines/powerOff/action",
        "Microsoft.Compute/virtualMachines/read",
        "Microsoft.Compute/virtualMachines/start/action",
        "Microsoft.Compute/virtualMachines/write",
        "Microsoft.Network/networkInterfaces/delete",
        "Microsoft.Network/networkInterfaces/effectiveNetworkSecurityGroups/action",
        "Microsoft.Network/networkInterfaces/join/action",
        "Microsoft.Network/networkInterfaces/read",
        "Microsoft.Network/networkInterfaces/write",
        "Microsoft.Network/networkSecurityGroups/join/action",
        "Microsoft.Network/networkSecurityGroups/read",
        "Microsoft.Network/virtualNetworks/read",
        "Microsoft.Network/virtualNetworks/subnets/join/action",
        "Microsoft.Resources/subscriptions/locations/read",
        "Microsoft.Resources/subscriptions/resourceGroups/read"
      ],
      "notActions": [],
      "dataActions": [],
      "notDataActions": []
    }
  ]
```

```
}
```

(Optional) If you want to create public IP addresses, then you must also provide the following set of permissions also.

```
        "Microsoft.Network/publicIPAddresses/delete"
        "Microsoft.Network/publicIPAddresses/join/action"
        "Microsoft.Network/publicIPAddresses/read"
        "Microsoft.Network/publicIPAddresses/write"
```

- The Azure account provided while adding Azure as both source and target must have the set of permissions as provided in the following JSON to do end-to-end migration.

```
{
    "permissions": [
      {
        "actions": [
          "Microsoft.Authorization/roleAssignments/read",
          "Microsoft.Authorization/roleDefinitions/read",
          "Microsoft.Compute/disks/beginGetAccess/action",
          "Microsoft.Compute/disks/read",
          "Microsoft.Authorization/locks/read",
          "Microsoft.Compute/locations/runCommands/read",
          "Microsoft.Compute/locations/vmSizes/read",
          "Microsoft.Compute/snapshots/beginGetAccess/action",
          "Microsoft.Compute/snapshots/delete",
          "Microsoft.Compute/snapshots/endGetAccess/action",
          "Microsoft.Compute/snapshots/read",
          "Microsoft.Compute/snapshots/write",
          "Microsoft.Compute/virtualMachines/instanceView/read",
          "Microsoft.Compute/virtualMachines/powerOff/action",
          "Microsoft.Compute/virtualMachines/read",
          "Microsoft.Compute/virtualMachines/runCommand/action",
          "Microsoft.Compute/virtualMachines/start/action",
          "Microsoft.Compute/virtualMachines/vmSizes/read",
          "Microsoft.Compute/virtualMachines/write",
          "Microsoft.Network/networkInterfaces/read",
          "Microsoft.Network/networkSecurityGroups/securityRules/delete",
          "Microsoft.Network/networkSecurityGroups/securityRules/write",
          "Microsoft.Network/virtualNetworks/read",
          "Microsoft.Resources/subscriptions/locations/read",
          "Microsoft.Resources/subscriptions/resourceGroups/delete",
          "Microsoft.Resources/subscriptions/resourceGroups/read",
          "Microsoft.Resources/subscriptions/resourceGroups/write",
          "Microsoft.Compute/disks/delete",
          "Microsoft.Compute/disks/endGetAccess/action",
          "Microsoft.Compute/disks/write",
          "Microsoft.Network/networkInterfaces/delete",
          "Microsoft.Network/networkInterfaces/effectiveNetworkSecurityGroups/
action",
          "Microsoft.Network/networkInterfaces/join/action",
          "Microsoft.Network/networkInterfaces/write",
          "Microsoft.Network/networkSecurityGroups/join/action",
          "Microsoft.Network/networkSecurityGroups/read",
          "Microsoft.Network/virtualNetworks/subnets/join/action"
        ],
        "notActions": [],
        "dataActions": [],
        "notDataActions": []
      }
    ]
```

```
  }
```

(Optional) If you want to create public IP addresses for target VMs on Azure, then you must provide the following set of permissions also.

```
        "Microsoft.Network/publicIPAddresses/delete"
        "Microsoft.Network/publicIPAddresses/join/action"
        "Microsoft.Network/publicIPAddresses/read"
        "Microsoft.Network/publicIPAddresses/write"
```

Service Accounts

For successful migration of VMs from AHV to Azure, Move requires the following.

- Prism Element for the AHV cluster
- An administrator for Windows source VMs or a root for Linux source VMs to run the source VM preparation scripts.

Registering an App and Applying Privileges (Azure UI)

For migrating the VMs from Azure, first you need to register the Azure app, and then apply the required privileges for the subscription in the Azure UI. Once you have the Subscription ID, Tenant ID, Application ID and the client secret value, you can add the Azure provider in the Move VM.

About this task

To register the app and apply the required privileges, do the following:

Procedure

1. Create a new App registration in Azure to provide necessary access to Move to perform the VM migration.

   a. Select **App registration** > **New registration**.

Figure 59: App registration

   b. Provide the app name, and then click **Register** to register the application with the default selection.

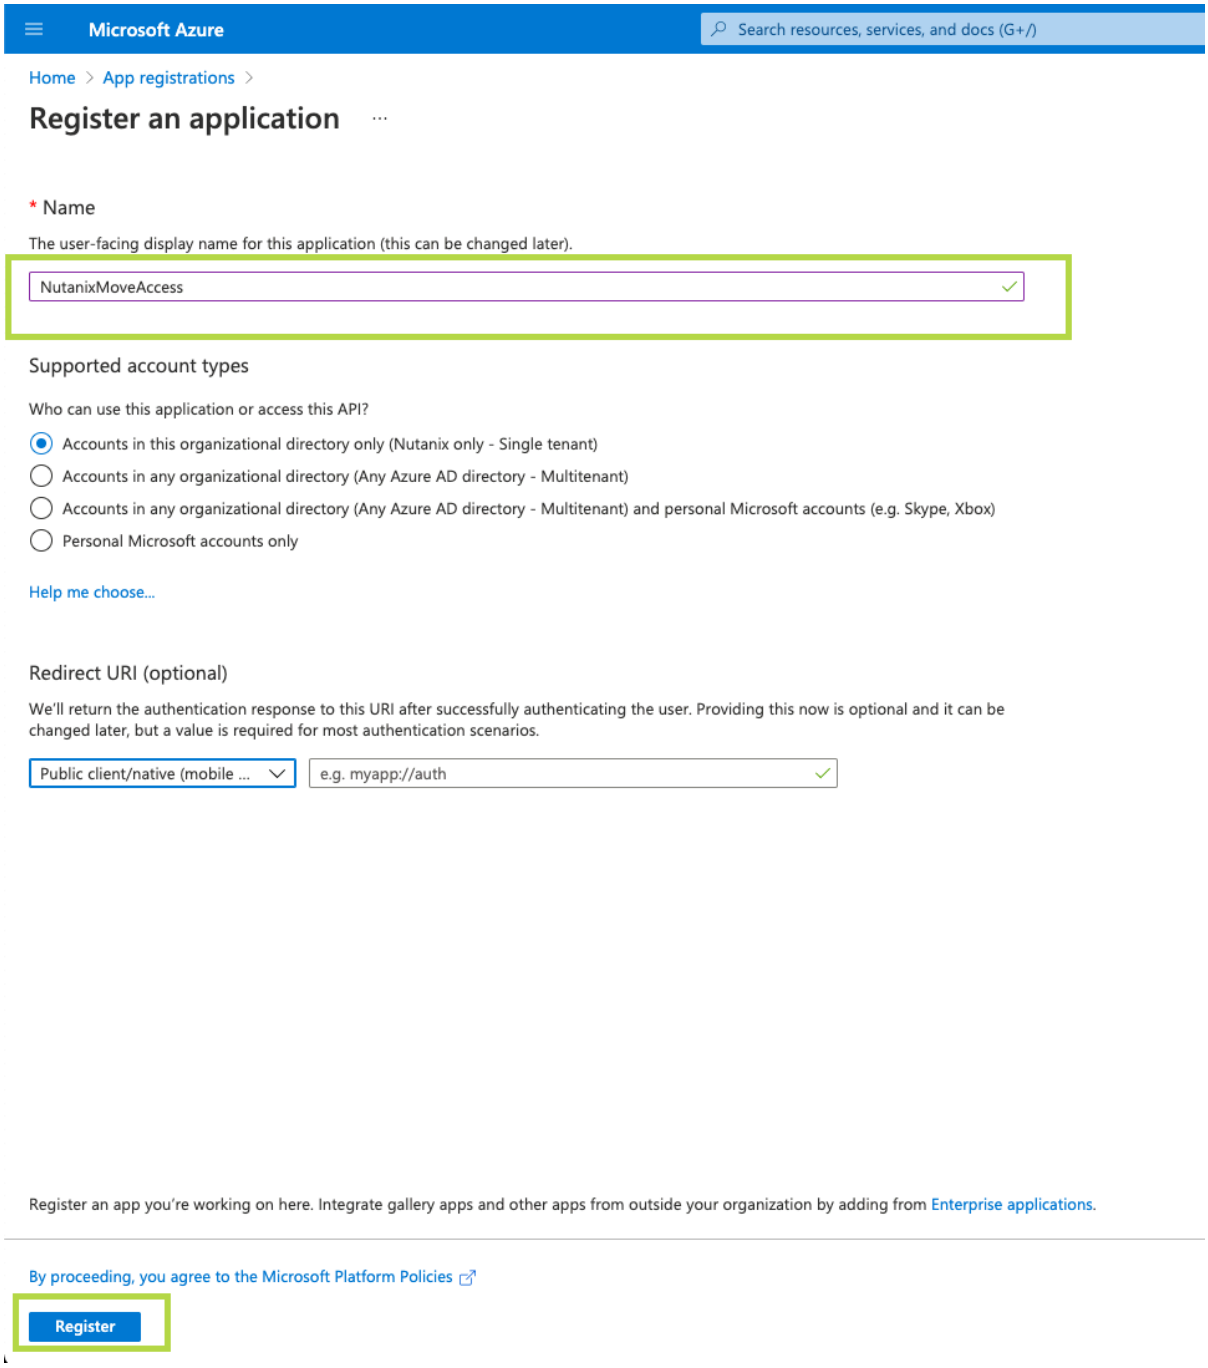

Figure 60: Register the app with default selection

c. Once the application is registered, copy the Application ID and the Tenant ID.

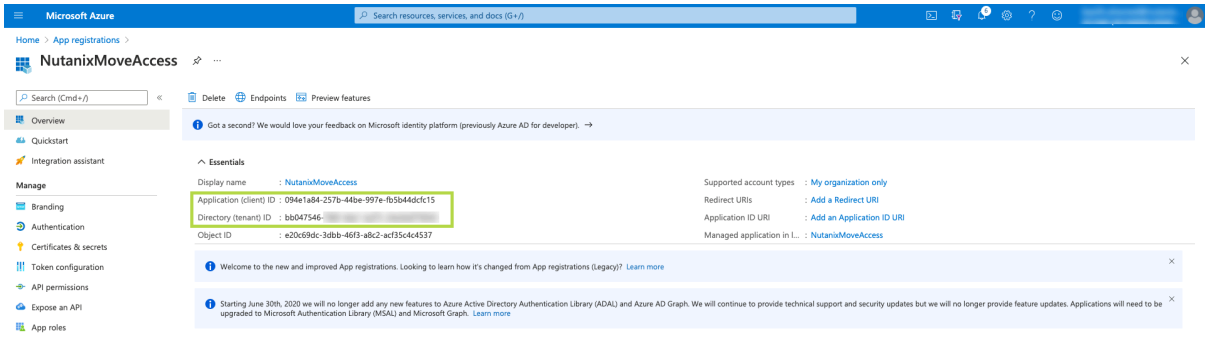

Figure 61: Copy Application ID and Tenant ID

d. Click **Certificates & Secrets** to create a client secret for the registered application. Set the description and the expiry period, and the click **Add**.

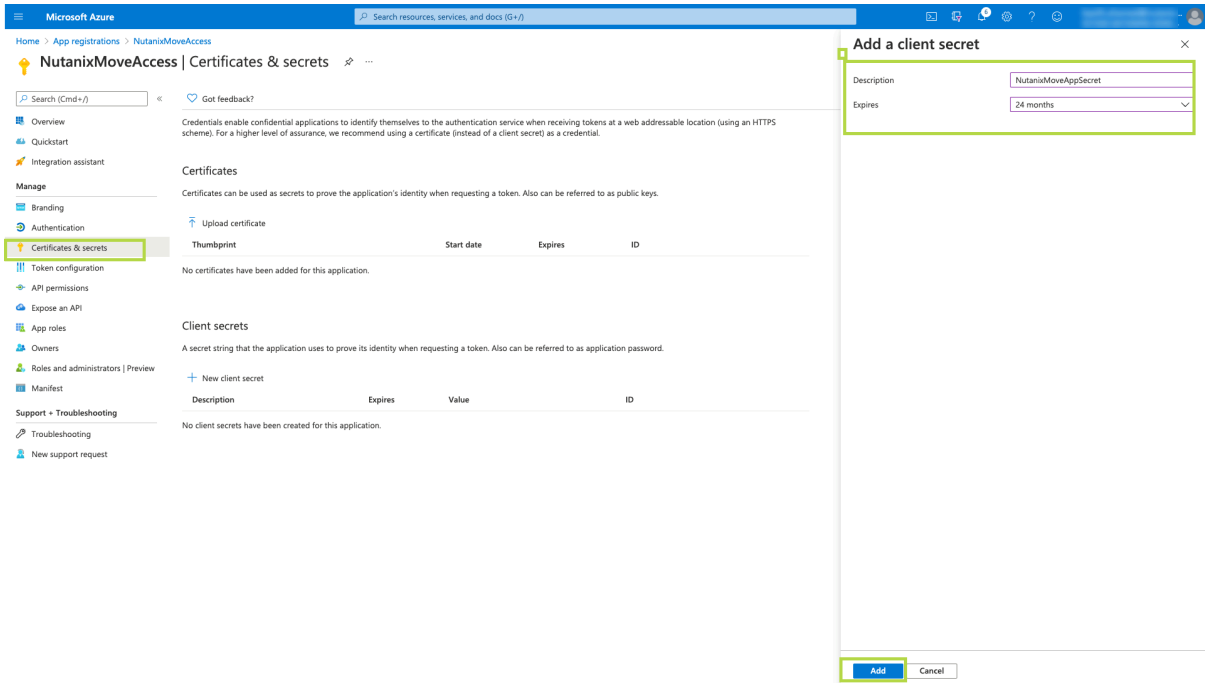

Figure 62: Creating a client secret

e. Copy the client secret value which is required for adding Azure provider in the Move VM.

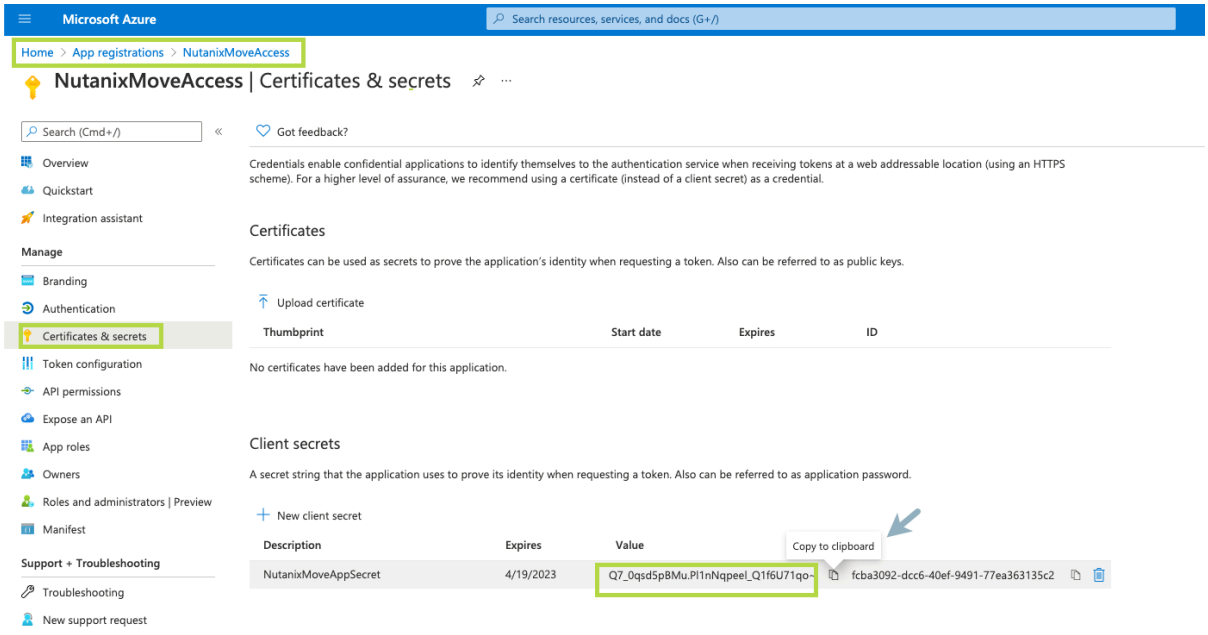

Figure 63: Copying Secret

2. Create a custom role in the subscription and assign that role to the application.

   a. Go to **Subscriptions** > **(Name of your subscription to add to Move)** > **Access control (IAM)**. Click **Add** > **Add custom role**.

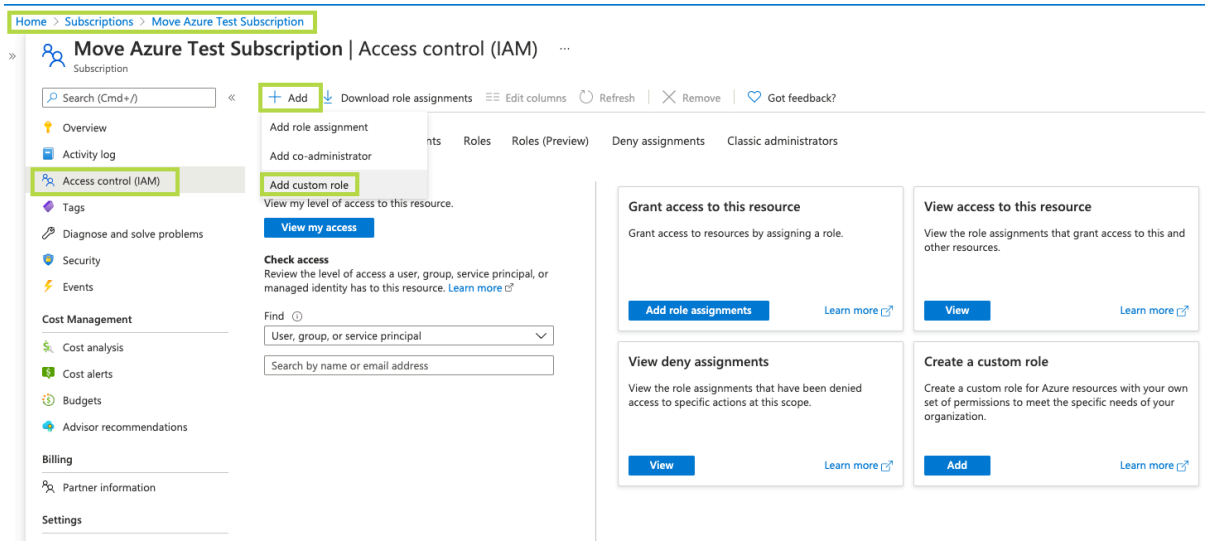

Figure 64: Adding a custom role

   b. Enter a custom role name, and then click the **JSON** tab. Click **Edit**. Replace the `permissions` section in the JSON with the required *set of permissions*, and then click **Save**. Click **Review + Create** to complete the custom role creation.

      To copy the *set of permissions*, refer to Requirements (Azure to AHV) on page 154 or Requirements (Azure to ESXi) on page 177 section.

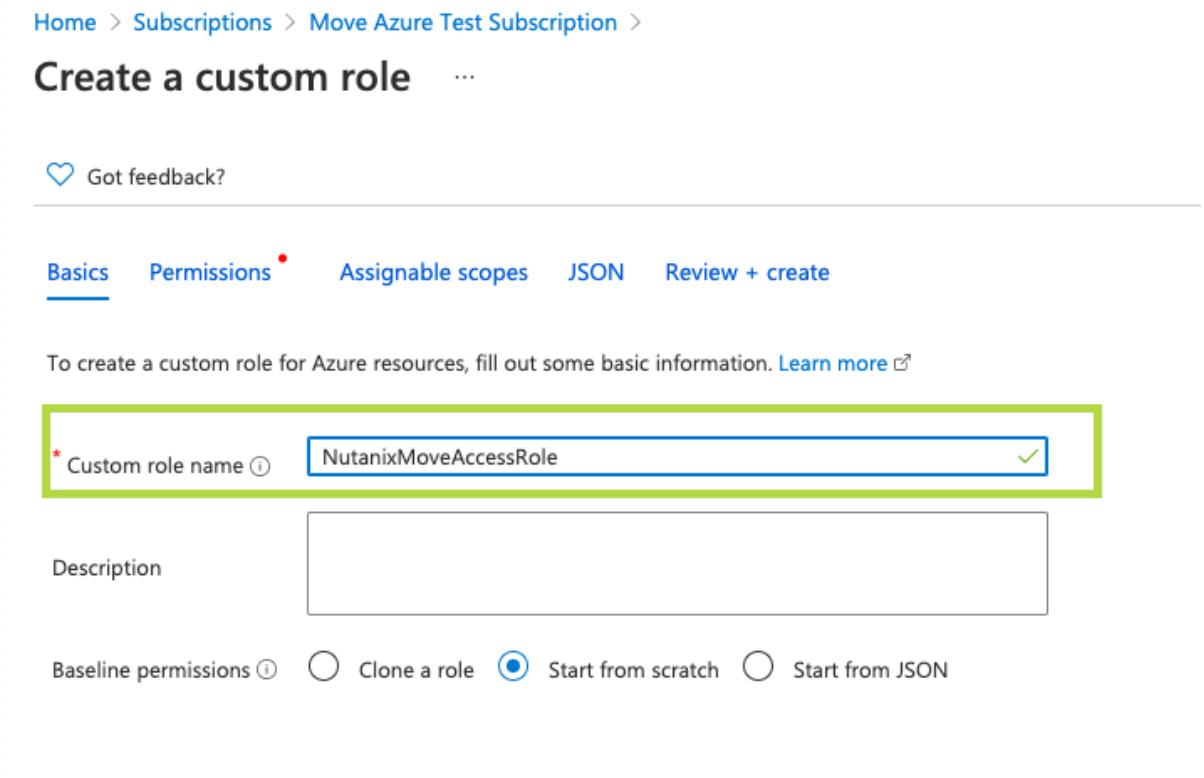

Figure 65: Creating a custom role

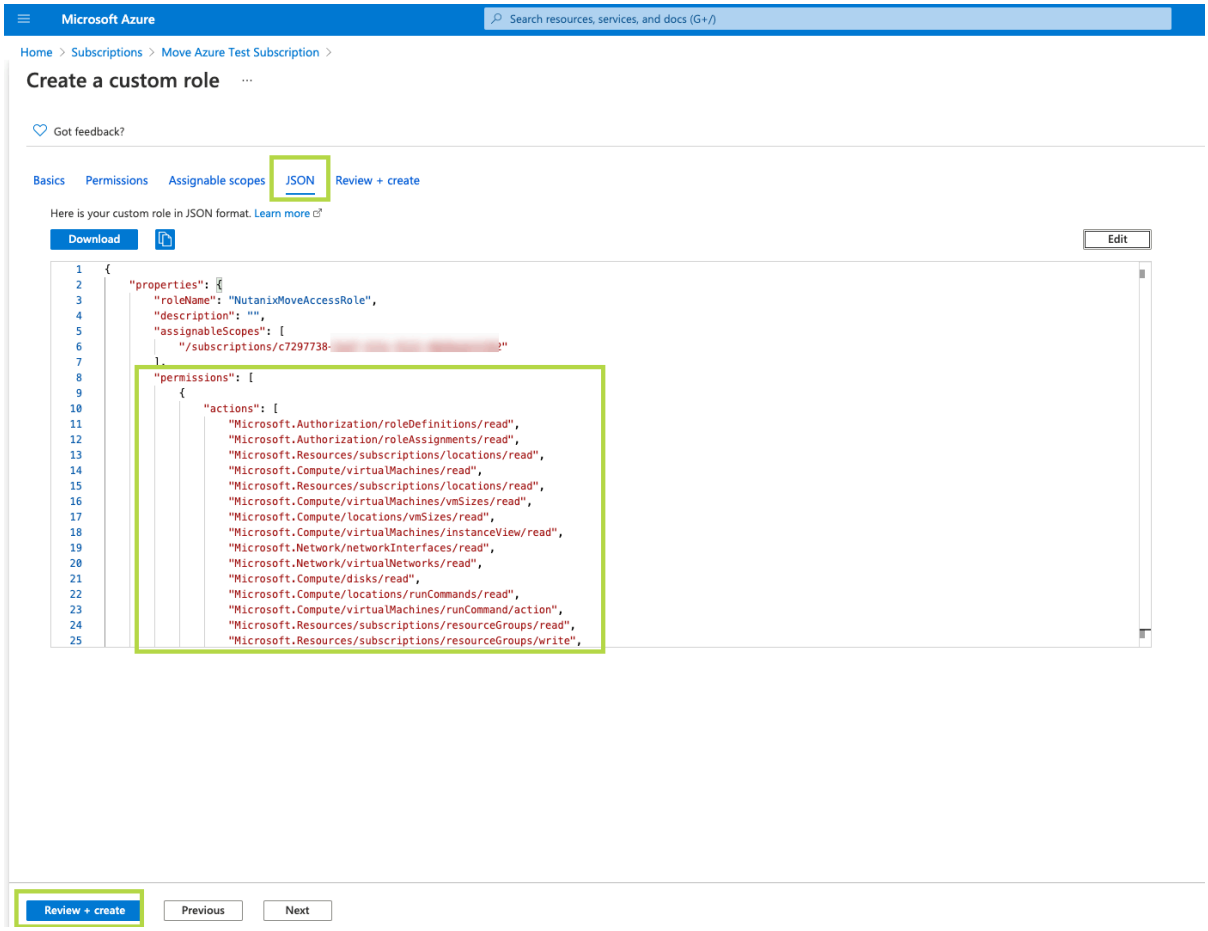

Figure 66: Editing JSON for required set of permissions

c. Go to **Access control (IAM)**. Click **Add** > **Add role assignment**.

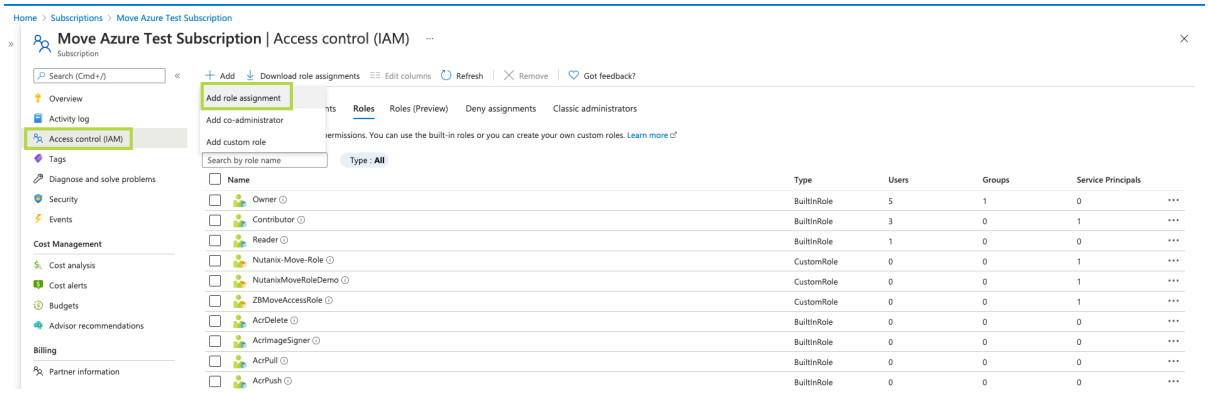

Figure 67: Assigning Permissions

d. Select the created role and the registered application, and then click **Save**.

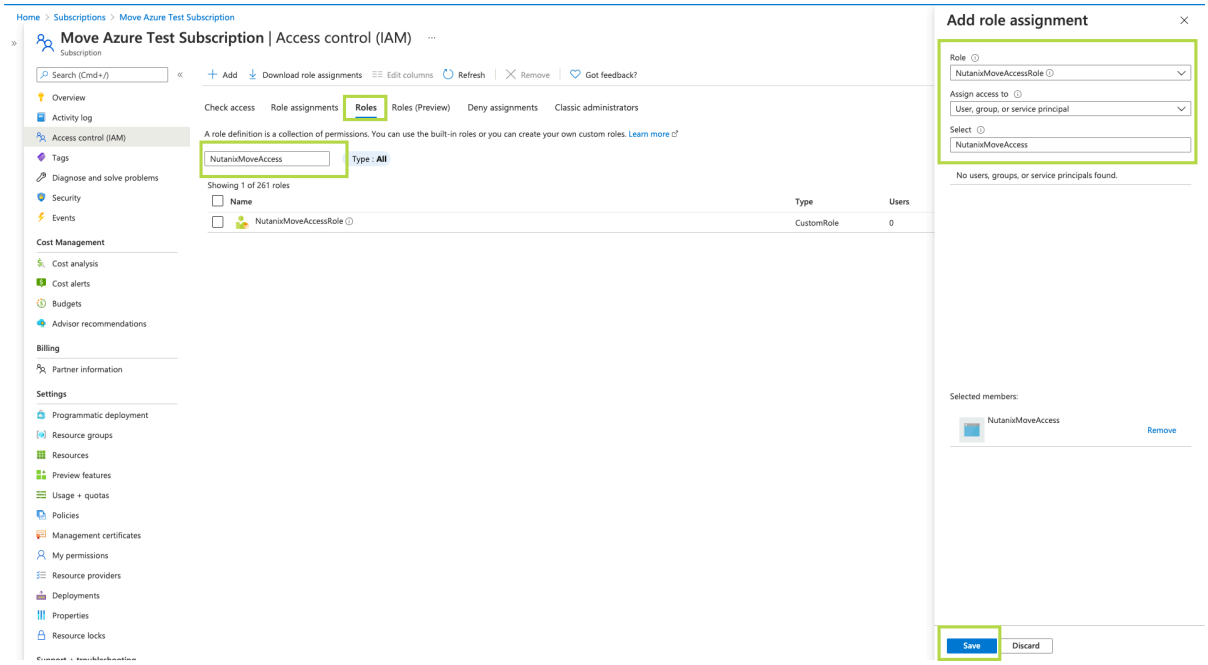

Figure 68: Role assignment

Now you can use this Subscription ID, Tenant ID, Application ID and the client secret value to add the Azure provider in the Move VM.

What to do next

You can now add the Azure environment in the Move UI. Refer to Adding an Azure Environment on page 167. You can also register the app in Azure and assign custom role through Move CLI. Refer to Registering the App in Azure and Assigning Custom Role (Move CLI) on page 164.

Registering the App in Azure and Assigning Custom Role (Move CLI)

For migrating the VMs from Azure, first you need to register the Azure app, and then apply the required privileges for the subscription through Move CLI or Azure UI. Once you have the Subscription ID, Tenant ID, Application ID and the client secret value, you can add the Azure provider in the Move VM.

About this task

To register the app in Azure and assign custom role, do the following:

Procedure

1. SSH to the Move VM as an admin.

   Refer to Accessing Move VM with SSH on page 21.

2. Switch to the root user by entering the password of the Move VM.

```
admin@move on ~ $ rs
[sudo] password for admin:
```

3. To create Azure App, run the following command:

```
root@move on ~ $ create-azure-app
```

Figure 69: Creating Azure App

A link and code is provided to authenticate with Azure.

4. Open the authentication link in a web browser and enter the code for authentication. Select the Azure account.

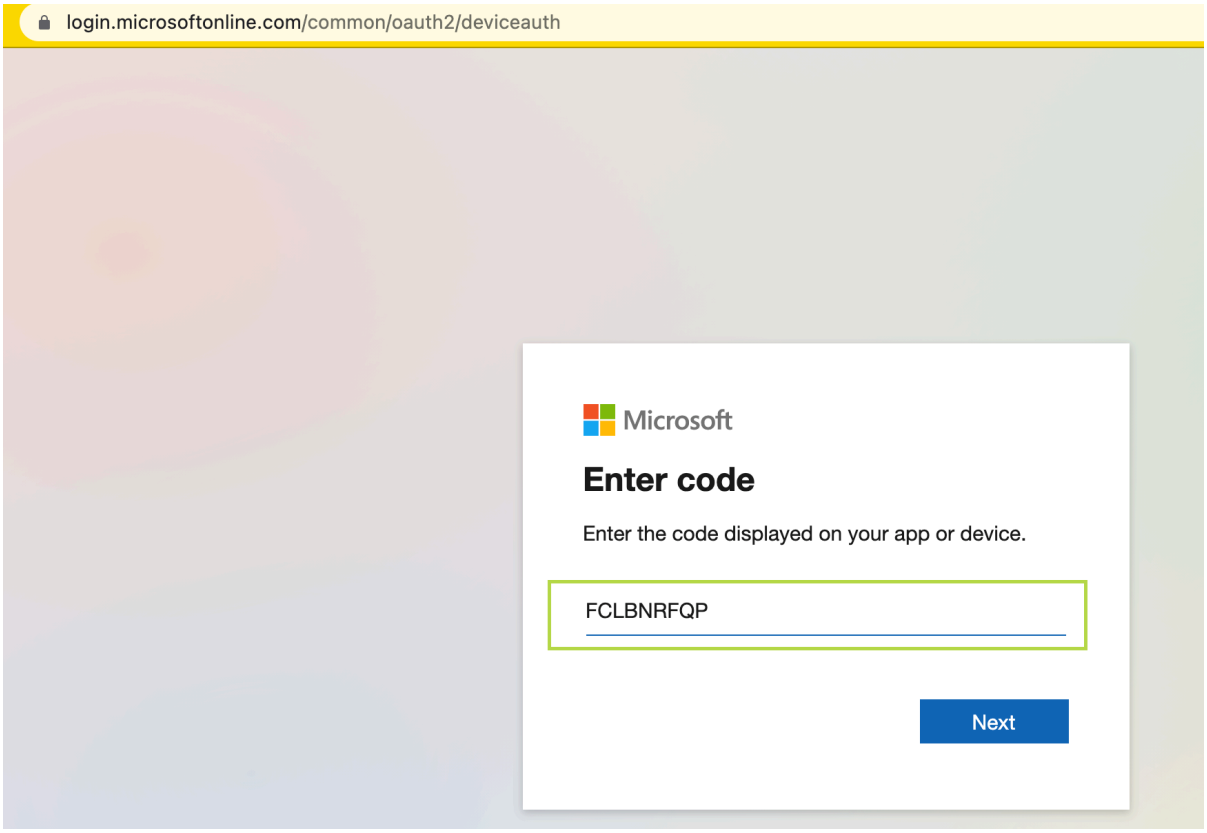

Figure 70: Authenticating with Azure

Authentication will be successful and a list of subscriptions will appear.

5. If your account has multiple subscriptions, provide the subscription ID to be used for migration.

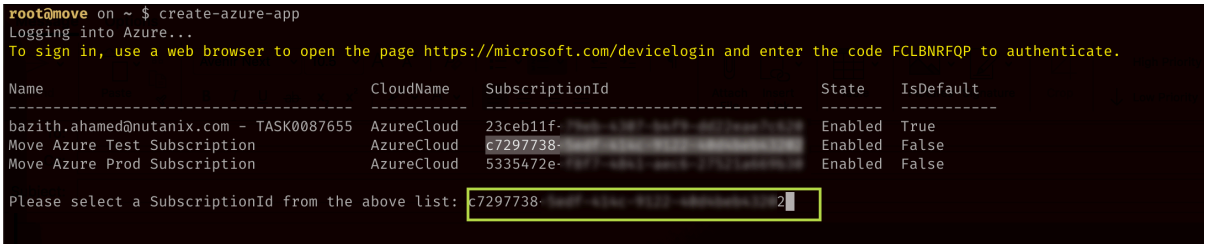

Figure 71: Enter the subscription ID

6. Enter the App name to be registered (default: NutanixMoveApp) and the custom role name which will be created with necessary permission and assigned to the registered app.

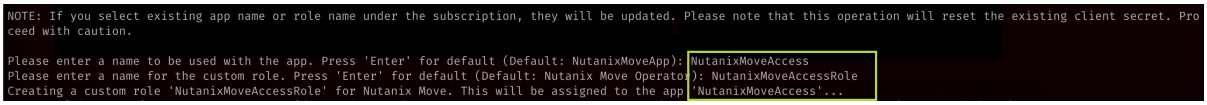

Figure 72: Enter the App and custom role name

Once the App gets registered in Azure and assigned with the custom role, the Subscription ID, Tenant ID, Application ID and the Client Secret is displayed in the output. Use these credentials to add Azure as a provider in the Move VM.

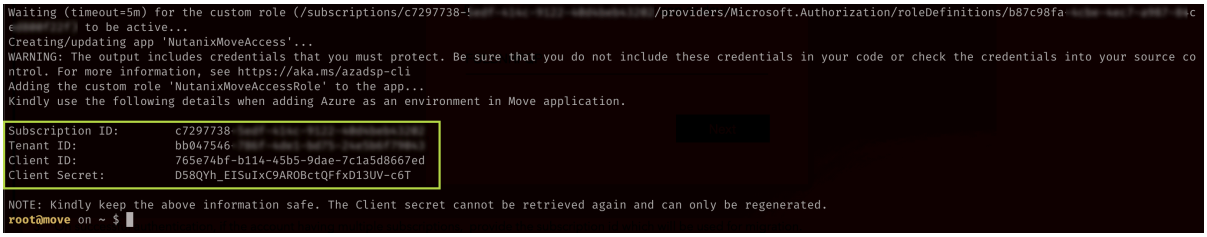

Figure 73: Lists the Subscription ID, Tenant ID, Application ID and the Client Secret

What to do next

You can now add the Azure environment in the Move UI. Refer to Adding an Azure Environment on page 167. You can also register the Azure app, and then apply the required privileges for the subscription from the Azure UI. Refer to Registering an App and Applying Privileges (Azure UI) on page 156.

Unsupported Features (AHV to Azure)

This section lists the unsupported features for migration from AHV to Azure.

- Guest operating systems other than the supported operating systems.

  For more information, refer to Supported Guest Operating Systems for AWS Migration
- AHV clusters with protection domain and stretched containers.
- VMs with GPUs.
- Clusters with encryption enabled.
- Clusters with multi-homing setups.

- Increasing the disk size of the source VM.
- IP address and MAC address retention.
- VMs having disks in volume groups.

Limitations (AHV to Azure)

The section lists the limitations for migration from AHV to Azure.

AHV to Azure Migration Limitations

The support for migration is constrained by the following.

- Windows 32-bit is not supported.
- Only Azure public cloud is supported.
- The following are not supported:
  - PC migration
  - Files migration
  - Objects migration
- Migration of a source VM which is in a VPC-based network is not supported.

Adding a Nutanix AOS Cluster Environment

While creating a migration plan, AHV AOS cluster can be added as both source or target. If you want to use ESXi on Nutanix cluster as target, then add the corresponding AOS cluster environment for ESXi.

About this task

> **Note:** When you add a AOS cluster, Move VM IP address with the subnet 255.255.255.255 is added to the global NFS allowlist.
>
> > **Caution:** Modifying the NFS allowlist of the destination container disables inheritance of the NFS allowlist at the global level and lead to issues with Move operations.

To add a Nutanix AOS cluster environment, do the following:

Procedure

1. Log on to the Move UI.

2. Click **+ Add Environment** under **Environments**.

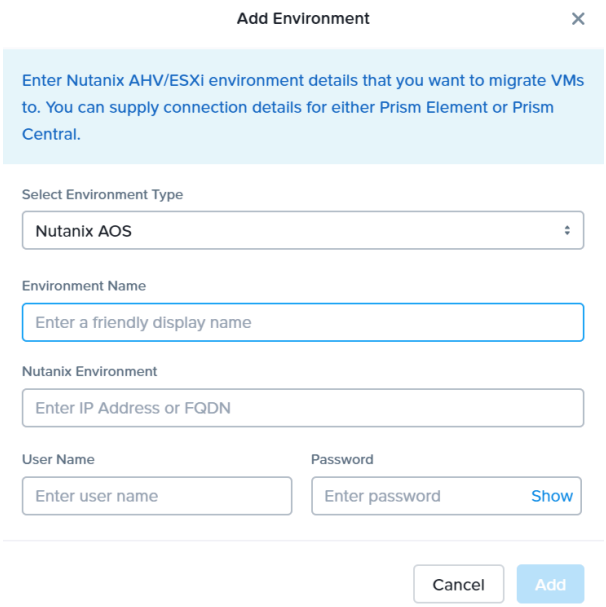

Figure 74: Add AOS Environment Dialog Box

The **Add Environment** window appears.

3. Select **Nutanix AOS** as the target environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the NC2 on AWS and AHV or VMware ESXi on Nutanix environment.
   b. **Nutanix Environment**: Enter the IP address or the FQDN for the target Prism Central or Prism Element.
   c. **User Name**: Enter the username for logging on to the target Nutanix environment.
   d. **Password**: Enter the password for logging on to the target Nutanix environment.

5. (Only for ESXi to ESXi migration) Enter credentials for registered vCenter.

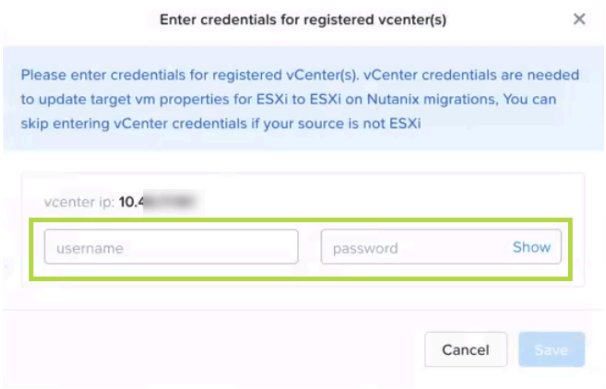

Figure 75: vCenter Credentials Dialog box

vCenter credentials are required to update the target VM properties for ESXi to ESXi on Nutanix migration.

Note:

- You can skip adding vCenter credentials if your source is not ESXi.
- To retain the VM properties on the target VM after migration, be sure to provide the target vCenter credentials. The following properties will be retained:
  - SCSI controller types
  - Network adaptor type
  - MAC address
  - Video card
  - Memory overcommit variables
  - Tool upgrade flag
  - Sync time with host flag
  - Disable acceleration flag
  - Enable logging flag

The AOS environment is added to the Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

Once you have added the environments, you can proceed to create the migration plan. For more information, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan (Hyper-V to ESXi) on page 103 or Creating a Migration Plan (Azure to AHV) on page 171 or Creating a Migration Plan (Azure to ESXi) on page 184

Adding an Azure Environment

While creating a migration plan, if you need to add an Azure source, you have to add at least one Azure environment for migration.

Before you begin

Ensure to register the Azure app and apply the required privileges for the subscription. You can do it through Azure UI or Move CLI.

About this task

Note:

- This procedure is only applicable for migration from Azure environment.
- Azure accounts with lock at subscription level cannot be added as a provider. If the lock is present at resource group level, VMs under that resource group are marked as unmigratable.

To add an Azure environment, do the following:

Procedure

1. Log on to Move UI.

2. Click **+ Add Environment** under Environments.
   The **Add Environment** appears.

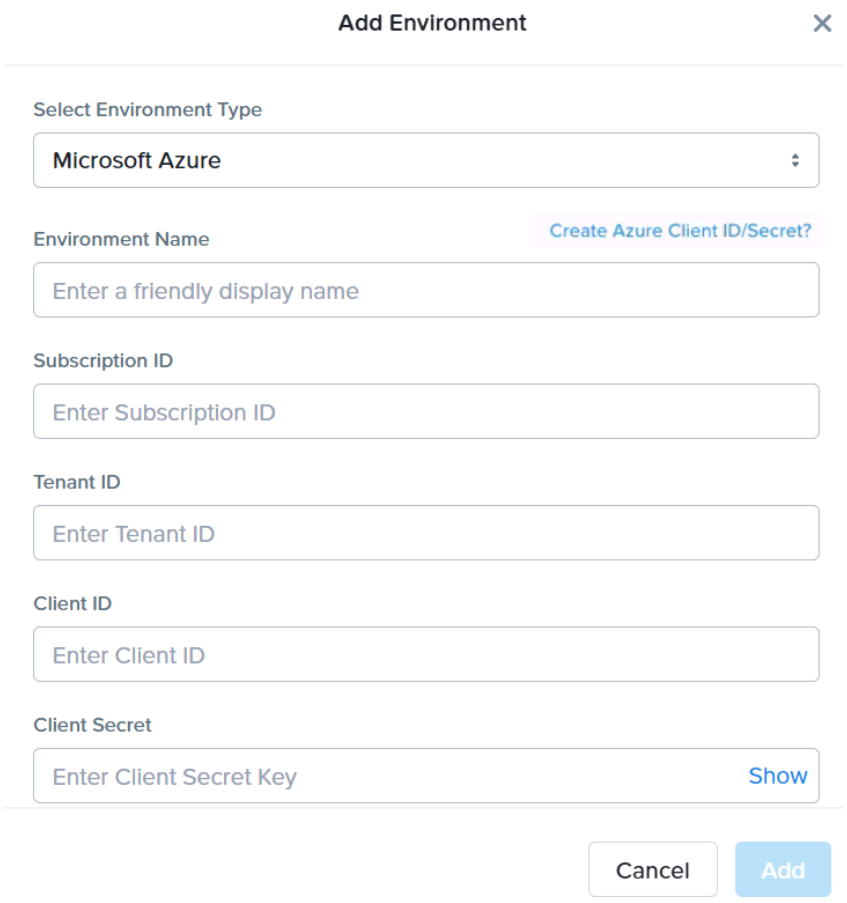

Figure 76: Add Environment Dialog Box

3. Select **Microsoft Azure** as the environment type.
4. Complete the indicated fields and click **Add**.
   a. **Environment Name**: Enter a name for the Azure environment.
   b. **Subscription ID**: Enter your Azure subscription ID.
   c. **Tenant ID**: Enter your Azure tenant ID.
   d. **Client ID**: Enter the client ID of the Azure account.
      If you do not have the Client ID and Secret, click the **Create Azure Client ID/Secret?** link for more details on creating Azure Client ID and Secret.
   e. **Client Secret**: Enter the value of the client secret.

      Note: Client Secret refers to the value of the client secret in Azure and not the secret ID.

   The environment is added to Move UI and can be viewed in the **Environments** list in the left pane of the Move dashboard.

What to do next

You can add a Nutanix AOS cluster environment. For more information, refer to Adding a Nutanix AOS Cluster Environment on page 39

Creating a Migration Plan (AHV to Azure)

You can create a migration plan to seed the data, cutover, and monitor the VMs. You can create the migration plan in Move without initiating the cutover process.

About this task

Note:

- This procedure is only applicable for migration from AHV to Azure.
- On your first log on, you can log on to the Move UI with your defaults credentials.
- If you cancel or discard the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.

  For information on canceling an ongoing migration, see Pausing or Canceling a VM Migration on page 257.

  For information on performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

To create a migration plan, do the following:

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.
3. On the Move dashboard, click **Create a Migration Plan**.

   Note: If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.
4. Enter the new migration plan name, and then click **OK**.

   Note: You can edit the migration plan name by clicking the pencil icon next to the migration plan name.

   The **New Migration Plan** window appears.

5. Complete the following fields, and then click **Next**.

   a. **Select a Source**: Select any AHV source for migration.

   If you select Prism Central IP address, a new field **Select Cluster** appears to select any cluster of that PC.

   Once you select the source, an appropriate target appears.

> **Note:** You might at times see a message relating to inventory collection as shown below. During this time, the environments undergoing Refresh will not be available for selection. Such environments do not automatically show up for selection once the inventory collection is over. In case you wish to select one of the environments undergoing Refresh (not available for selection), you will need to select **Cancel** and wait for inventory collection to get over and then create the migration plan again.

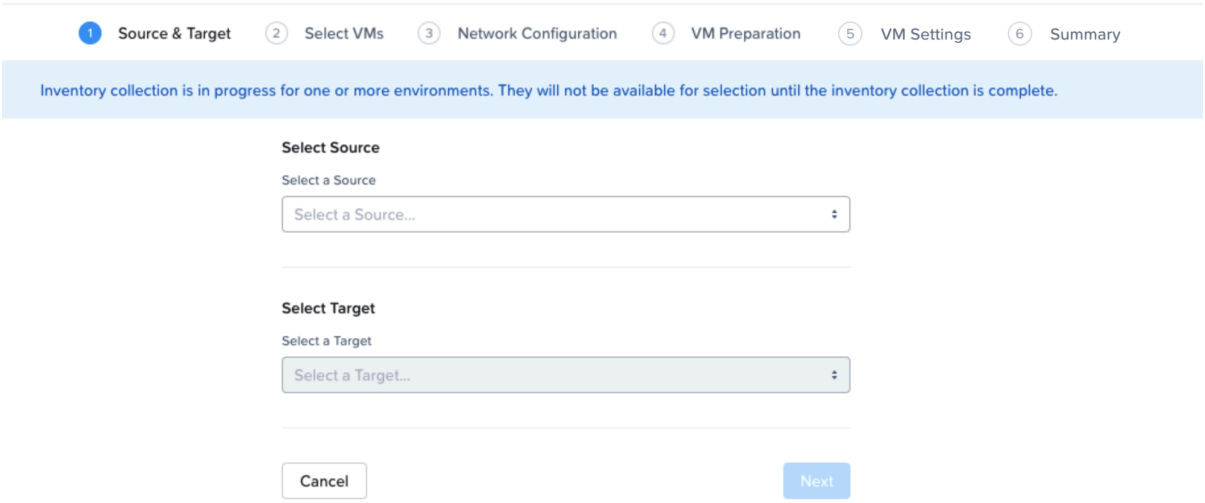

Figure 77: Inventory Collection Message

   b. **Select a Target**: Select the target Azure instance for the migrating VMs.

   c. **Target Location**: To migrate the VMs, select a target location.

> **Note:** Only locations with Internet gateway VPCs are available for selection.

   d. **Target Resource Group**: Select a resource group to create the target instance. Resource groups are location specific.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click **+Add All**, and then click **Next**.

   Note:

   - You cannot add more than 50 VMs in a single migration plan.
   - Migration of VMs which have NIC(s) attached to VPCs is not supported.

   You can filter the VMs by name in the **Filter** bar. You can also sort the VM types by selecting the appropriate column.

   Note: The VMs for which migration has failed are not displayed. To show the entire list of VMs, select **All VMs** from the drop-down list. To show the configuration of VMs, select **Configuration** from the drop-down list. A question mark icon appears beside an unavailable VM that displays more information about that VM and indicate the reason of a failed VM migration.

   Note: Migrate VMs retain their power state on the destination cluster.

   The selected VMs are displayed in the sidebar.

7. In the **Network Configuration** screen, select the following fields, and then click **Next**.

   a. **Target Network**: Select the network that target instance will use after migration.

   b. **Subnet**: Select the subnet for the target network.

   c. **Security Group**: Select the security group for the target network.

8. In the **VM Preparation** screen, select one of the following VM preparation modes.

   » **Automatic**. Move automatically runs scripts on the source VMs to prepare them for migration. Provide the credentials of the source VMs under **Windows VMs** or **Linux VMs**, depending on the type of the source VM.

   Note: For Windows VMs, Move supports only username-password sign-in option for authentication. It does not support username-PIN sign-in option.

   For more information, refer to Automatic VM Preparation (AHV to Azure) on page 245.

   » **Manual**. Move displays the VM preparation scripts for Windows and Linux VMs. To manually download and run the migration preparation software, select this option, and then run the scripts provided in the **VM Preparation** screen on the respective source VMs.

   These scripts prepare the instance by performing the following installations.

   For Windows: Move does not install any drivers.

   For Linux: Move installs hv_vmbus, hv_netvsc, and hv_storvsc (Hyper-V drivers).

   Note: Ensure to run the VM preparation script on all selected VMs. If not, Move will only migrate VM data and not start the VM in the target environment. Also, all operating system configuration options will be bypassed.

   » **Mixed**. Move allows you to select the VM preparation mode for each VM. This setting allows you to prepare one set of VMs manually and prepare another set of VMs automatically in the same migration plan.

If you want to have such a hybrid combination of **Automatic** and **Manual** VM preparation modes, proceed to the step 8.

> **Note:** The preparation mode automatically switches to **Mixed** if you change the mode of preparation for a set of VMs under **Change Settings** and have a mix of **Automatic** and **Manual** VM preparation modes. You cannot manually select the **Custom** option from the **Preparation Mode** drop-down list.

9. In the **Override individual VM Preparation** section, click **Change Settings** to override settings for the individual VMs. You can edit the VM preparation credentials, remove VMs, or update the VM preparation mode.

   > **Note:** If you change the **Mode of Preparation** to **Manual** for a VM, then copy the new generated scripts of that specific VM and run them on the source VM.
   >
   > Access the newly generated VM preparation script for that VM by selecting the Copy Scripts icon under the **Actions** column.

   After making the required changes, click **Done**.

   Then, click **Next**.

10. In the **VM Settings** screen, do one or both of the settings, and then click **Next**.

    a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.

    b. **Create Public IP Address**: Select this option if you want to create public IP address for all the VMs in the migration plan.

       > **Note:** The Azure app must have the following permissions to create target VMs with public IP address. Otherwise, the target VMs will be created without public IP address.
       >
       > ```
       > "Microsoft.Network/publicIPAddresses/delete"
       > "Microsoft.Network/publicIPAddresses/join/action"
       > "Microsoft.Network/publicIPAddresses/read"
       > "Microsoft.Network/publicIPAddresses/write"
       > ```

    c. **Settings for Individual VMs**: Click **Change Settings** to configure settings (such as VM priority) for individual VMs.

    d. **Schedule Data Seeding**: Select this checkbox to select the date and time for migration.

11. In the **Summary** screen, choose one of the following, and then proceed review the VM migration summary.

    » **Back**: To edit the information, click this option.

    » **Save**: To save the migration plan, click this option.

    For more information about how to start the migration later, and check the migrated VM status and details, refer to Environments and Migration Plan Management on page 259.

    » **Save and Start**: Click this option to save the migration plan and begin the migration immediately

    Once you save and proceed, the seeding process for migration begins.

    > **Note:** This process takes some time.

    > **Note:** The seeding process can take several minutes depending on the number of VMs.

What to do next

If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover (AHV to Azure) on page 246

Automatic VM Preparation (AHV to Azure)

You can automate the guest VM preparation.

About this task

Note:

- Before automatic Windows VM preparation, enable WinRM and enable the ports.
- Automatic VM preparation does not work with the Windows domain account. Use the Windows built-in administrator credentials for the Windows VMs.

To automatically prepare the VMs, do the following:

Procedure

1. In the **Preparation Mode** drop-down, select **Automatic**.
2. In the **Credentials for Source VMs** section, enter the username and password for the guest VMs to allow Move to install the necessary drivers.
3. In the **Override Individual VM Settings** section, do the following, and then click **Next**.
   a. Click **Change settings** to override the settings for the individual VMs. You can update the VM preparation credentials of the VMs, remove VMs, or update the VM preparation mode.

      Note: If you select **Manual**, you must copy the displayed scripts and run them on the source VMs.
4. In the **VM Settings** screen, do the following, and then click **Next**.
   a. **VMs Priority**: The scheduling priority of the VMs migration (at the migration-plan level) is set to **Medium** by default. Modify the scheduling priority as required. For more information about VM priority, refer to Virtual Machine (VM) Priority on page 288.
   b. **Settings for Individual VMs**: Click **Change Settings** to configure settings (such as VM Priority) for individual VMs.
   c. **Schedule Data Seeding**: Select this checkbox to select the date and time for migration.

   The credentials of VMs are validated. Once the validation is successful, the Guest Tools are downloaded and installed in all the VMs of the migration plan. Then, the VMs are validated for readiness.

   Note: If the validation of credentials or Guest Tools installation fails, you can update the credentials or remove the VM from the migration plan and proceed by clicking **Next**.

What to do next

If you have started the migration, the next step is to perform the cutover. For more information, refer to Performing a Migration Cutover (AHV to Azure) on page 246

Enabling WinRM (AHV-Azure)

You need to enable WinRM to install the Guest Tools on AHV VMs.

Before you begin

Ensure that the ingress ports 5985 and 5986 are enabled.

About this task

> **Note:** This method is a prerequisite for Automatic VM preparation to work with Windows source VMs.

To enable WinRM.

Procedure

1. Open PowerShell in Windows VM.
2. Run the script to enable WinRM on AHV VMs.

```
> winrm quickconfig -q
winrm set winrm/config/winrs '@{MaxMemoryPerShellMB="300"}'
winrm set winrm/config '@{MaxTimeoutms="1800000"}'
winrm set winrm/config/service '@{AllowUnencrypted="true"}'
winrm set winrm/config/service/auth '@{Basic="true"}'

netsh advfirewall firewall add rule name="WinRM 5985" protocol=TCP dir=in
 localport=5985 action=allow
netsh advfirewall firewall add rule name="WinRM 5986" protocol=TCP dir=in
 localport=5986 action=allow

net stop winrm
cmd /c 'sc config winrm start= auto'
net start winrm
```

3. Run the script to enable Secure Sockets Layer (SSL).

```
> $c = New-SelfSignedCertificate -DnsName "$(hostname)" -CertStoreLocation cert:
\LocalMachine\My
winrm create winrm/config/Listener?Address=*+Transport=HTTPS
 "@{Hostname=`"$(hostname)`";CertificateThumbprint=`"$($c.ThumbPrint)`"}"
```

Performing a Migration Cutover (AHV to Azure)

When the seeding process is complete, you can cut over the selected VMs to the Azure target.

About this task

You can monitor the VM migration progress by clicking the status link.

> **Note:**
>
> - You can perform this operation only when the VM status is Ready to Cutover.
> - Recommended that cutover should be performed within one week of initial data seeding.
> - If the initial data seeding finishes in less than 10 minutes, Move continues to wait for 10 minutes to take the incremental snapshot; however, you can trigger the cutover immediately.
> - The cutover process increments in absolute numbers.

To perform a cutover, do the following:

Procedure

1. In the Move UI, click the **Ready to Cutover** status to display the list of available VMs.
2. To cut over, select the VMs or group of VMs.
3. Click **Cutover**.
   The cutover process performs the following VM actions.

   - Shuts down the VM
   - Takes the final snapshots for the VM and copying the final changes to Azure
   - Creates an instance in the Azure target
   - Attaches replicated disks to the VM
   - Powers on or off the instance (depends on the initial power state)

   The cutover process begins immediately and might take a few minutes.

   Once the cutover is complete, the Azure instance is available. By default, the security group attached to the instance has all the ingress traffic blocked. To access the guest VMs, you need to update the security group according to your requirement.

What to do next

You can manage the migration plan once the migration plan is ready. For more information, refer to Environments and Migration Plan Management on page 259

Performance Matrix for Large Data Migration

Move performs end-to-end migration of large VMs. The scenarios are tested based on the following parameters. The following tables show the performance numbers from the Move lab.

Table 21: Performance Numbers of Large Data Migration (AHV to Azure)

Total migration size	Number of Disks	Location used	Average latency source - target	Data seeding duration	Cutover duration
2 TiB	2	US West	19 ms	3 hours 40 minutes	9 minutes

CREATING A TEST CAPABLE VM MIGRATION PLAN

Creating a test capable migration plan is an optional feature to test the VMs on the target environment prior to the final cutover. During the final cutover, the source VMs are powered off which disrupts the underlying data replication. You can test multiple VMs with this feature without disrupting the source VMs. After the test, you can continue with the final cutover in your production environment. This feature supports all provider combinations except AHV to AWS.

Before you begin

Ensure that you have added the source and target environments before creating the migration plan.

Note:

- ESXi to AHV and ESXi to NC2 on AWS VM migration is used as a sample test migration plan.
- When using the *Test Capable Migration* feature for Azure VMs having multiple disks, there is possibility of inconsistency with disk snapshots which will be used to create target test VMs.
- Test migration is supported for all provider combinations except AHV to AWS.
- Refer to specific migration plans in this guide for information on how to add the required environments.

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.
3. On the Move dashboard, click **Create a Migration Plan**.

   Note: If you have existing migration plans, click the **+ New Migration Plan** link to create a plan.

   The **Enter Migration Plan Name** window appears.

4. Enter the new migration plan name, and then click **OK**.

   The **New Migration Plan** window appears.

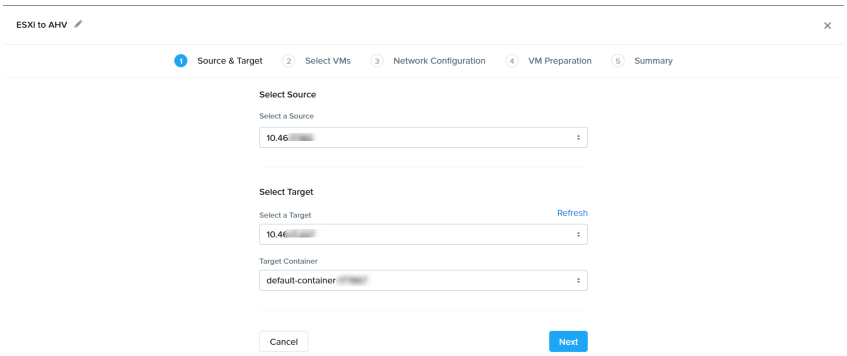

Figure 78: Migration Plan Window

5. Complete the following fields, and then click **Next**.

   a. **Select a Source**: Select any vCenter Server or standalone ESXi host as source for migration. Once you select the source, an appropriate target appears.

   b. **Select a Target**: Select any NC2 on AWS and AHV target for the migrating VMs.

   c. **Target Containers**: Select the container on which you will migrate the VMs.

6. In the **Select VMs** screen, select one or more VMs from the list. To add all the VMs, click **+Add All**, and then click **Next**.

   > **Note:** You cannot add more than 50 VMs in a single migration plan.

   The **Network Configuration** window appears.

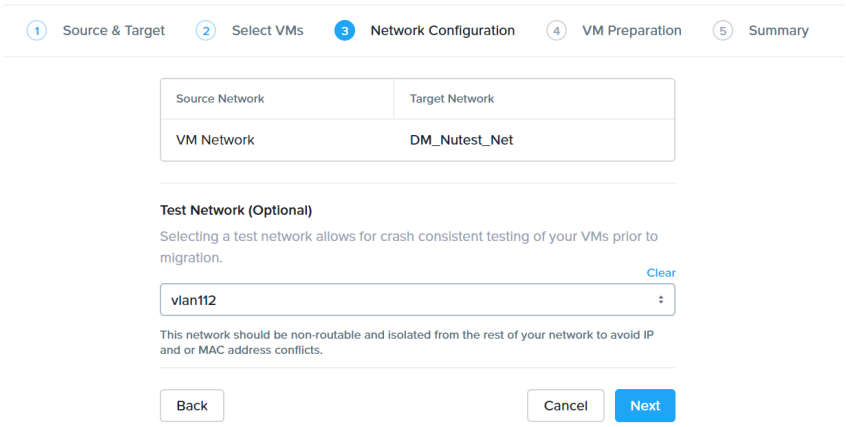

Figure 79: Network Configuration Window

7. On the **Network Configuration** screen, select the target network from the **Target Network** drop-down.

   - [Applicable only if Prism Central is used] Move provides the option to select VPC-based or VLAN-based target subnets. Based on the selection of a VPC or VLAN ID as the target network, the respective subnets are listed in the target subnet drop-down menu. Select the required subnet from the drop-down menu.

     Note: Overlay subnets which do not have IP address pool(s) associated will be disabled in the subnet drop-down menu.

   Test Network (Optional) is enabled.

8. Select the test network from the **Test Network (Optional)** drop-down.

   - For migrations to NC2 Azure, there is an additional field to select a test subnet for the selected test network. Based on the selection of a VPC or VLAN ID as the test network, the respective subnets are listed in the test subnet drop-down menu. Select the required test subnet from the drop-down menu.

     Note:

     - **Test Network (Optional)** is enabled only after selecting the following:
       - The target network.
       - The subnet corresponding to the target network (for migrations to NC2 Azure only).
     - To avoid conflicts, the **Test Network (Optional)** drop-down lists all other networks except the one selected by the target network.
     - **Test Network (Optional)** is an optional feature and you can skip it to go directly to the production target.

   Click **Next**.

9. In the **VM Preparation** screen, select **Automatic** from the **Preparation Mode** drop-down.

   In this example, **Automatic** preparation mode is selected. Refer to specific migration plans for **Manual** and **Mixed** preparation mode. You can choose based on your requirement.

10. In **Credentials for Source VMs** , provide the credentials for your Windows or Linux source VMs. Then, click **Next**.

> **Note:** For Windows VMs, Move supports only username-password sign-in option for authentication. It does not support username-PIN sign-in option.

The **Migration Plan Summary** page appears.

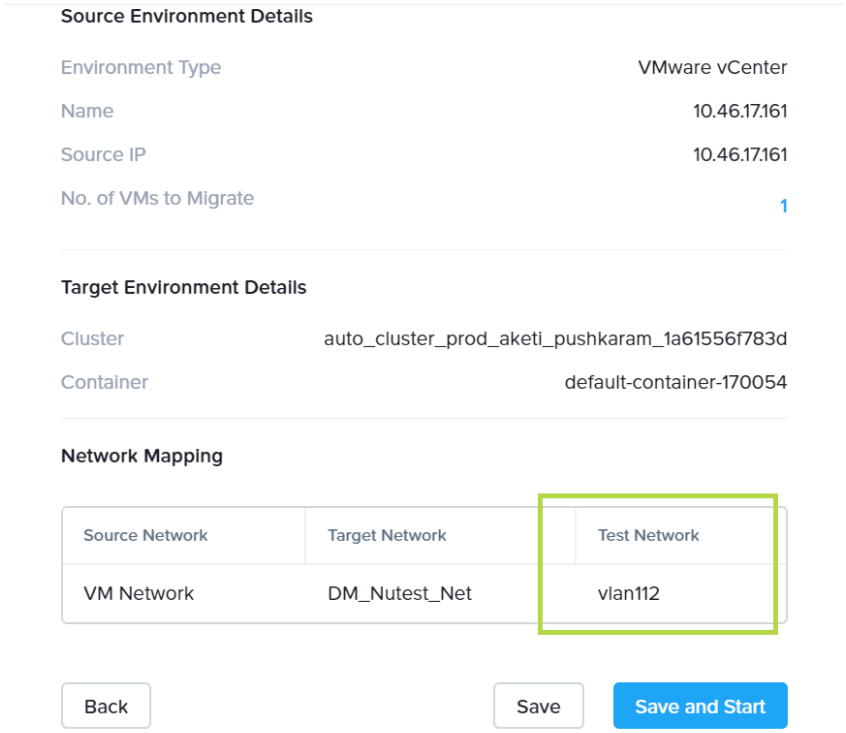

Figure 80: Migration Summary Page

11. On the **Summary** screen, review the migration summary details and click **Save and Start**.

The **Move Dashboard** appears with the existing migration plans and the seeding process begins.

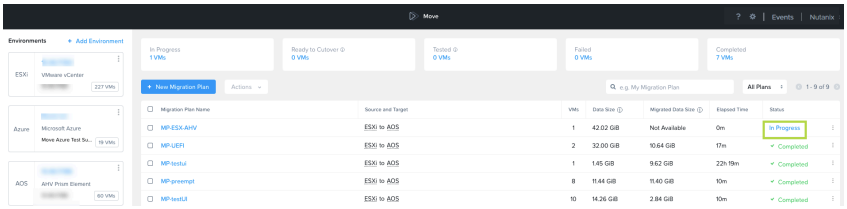

Figure 81: Move Dashboard with Existing Migration Plans

12. Click **In Progress** to monitor the progress of the migration.

    The **Migration Plan Details** window appears.

Figure 82: Migration Plan Details Window

> **Note:**
>
> - All tabs are disabled until you chose a VM and enabled at different stages as needed.
> - During the migration, if you select the VM, only the **Cancel** action is enabled for canceling the migration plan.
> - If there is any failure, the **Retry** and the **Discard** actions are enabled.

13. After the migration status changes to **Ready to Cutover**, the following actions are enabled:

    » **Test Actions**: Click to continue with testing the VMs on the target environment.

    » **Cutover** : Click to continue with normal migration process in the production environment.

14. Select **Create Test VM** from the **Test Actions** drop-down. The source VMs remain powered on and a test VM is created in the target environment.

    This process takes some time.

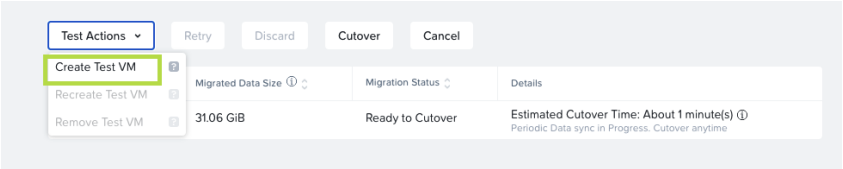

Figure 83: Migration Plan Details Window

Click **Continue** on the dialog box.

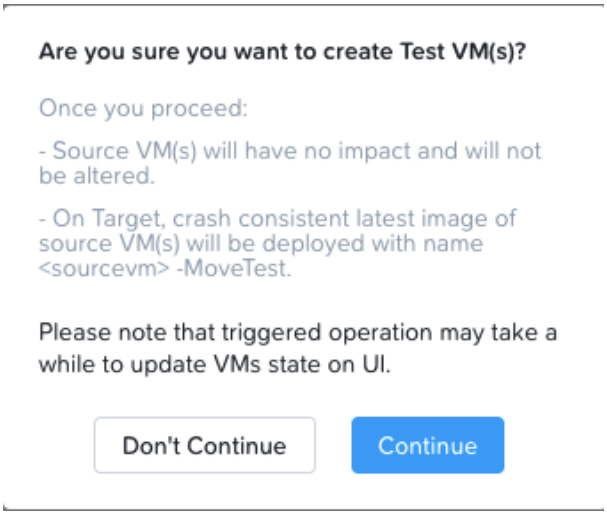

Figure 84: Confirmation Dialog Box

15. Click **View Test VM** option created under the **Migration Status** tab under **Ready to Cutover**.

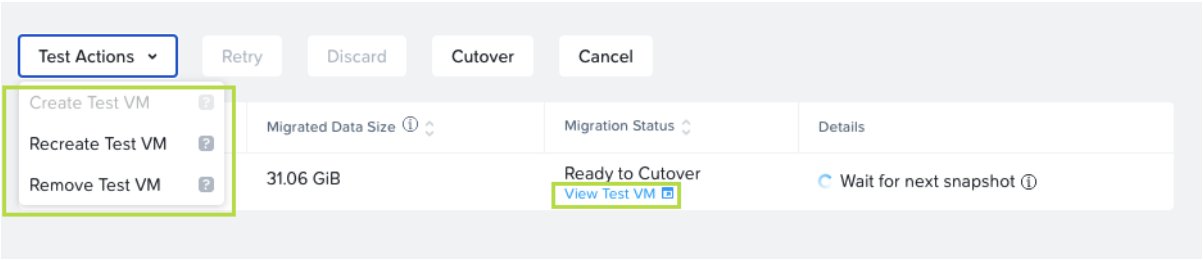

Figure 85: Migration Plan Details Window

The following options are enabled after you click **Test Actions**:

- **Recreate Test VM**: Click to recreate a test VM.
- **Remove Test VM**: Click to remove the deployed test VMs from the target and changes the VM status back to **Ready to Cutover**.

16. Click **View Test VM**.
    A new window for the target network opens up.

17. Enter the credentials of the source VM to log on.

18. Look for the test VM and perform test operations.

    > **Note:** Test VMs are suffixed with `-MoveTest` in the target network.

19. Once you test the VMs in the target environment, come back to the **Move Dashboard** and select **Cutover** to continue with normal migration to the production environment.

    You can click **Remove Test VM** to clean up the target environment and click **Recreate Test VM** to perform the test again.

What to do next

You can perform test migrations for all types of migration supported by Move except AHV to AWS migration. For more information about creating migration plan and performing test migration for supported environments, refer to Creating a Migration Plan on page 42 or Creating a Migration Plan (ESXi to ESXi) on page 66 or Creating a Migration Plan on page 84 or Creating a Migration Plan on page 127 or Creating a Migration Plan (AWS to ESXi) on page 147 Creating a Migration Plan (AHV to AWS) on page 214Creating a Migration Plan (Azure to AHV) on page 171Creating a Migration Plan (Azure to ESXi) on page 184.

CUSTOMIZING THE TARGET VM CONFIGURATION

Nutanix Move provides the option to customize the target VM configuration for migrations. You can do it at the migration-plan level and at the VM level. This topic refers to the customization that can be done at the VM level.

Before you begin

Ensure that at least one migration plan is available in Move.

> **Note:** Customizing the target VM configuration is supported for the following migrations only:
>
> - ESXi to AHV
> - ESXi to NC2
> - ESXi to ESXi
> - Hyper-V to AHV
> - Hyper-V to ESXi
> - AHV to AHV

About this task

The procedure details the steps to customize the target VM configuration for migrations at the VM level.

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **VM** and click **Continue**.
   Move dashboard for VM migration appears.
3. On the Move dashboard, click the **Status** link of the migration plan where you want to customize the target VM configuration.
   The resulting screen displays all the VMs in that migration plan, along with the details of the migration of each VM.

   > **Note:** To view the VMs that are part of all the migration plans, click any of the status fields at the top of the dashboard.

4. Identify the VM whose target VM configuration you want to customize and click **VM Configurations** of that VM.
   VM Configurations window appears, and displays the properties of the source and target VMs, along with options to customize the target VM configuration.

   > **Note:** The **Target VM configuration** selected by default will be the one that was selected while:
   >
   > - creating this migration plan, or
   > - customizing this target VM configuration the last time.

5. Select one of the following options (if necessary) from the **VM Migration Type** drop-down menu.

   - **Configure Target VM Properties**: The target VM synchronizes with the source VM properties at the time of migration plan creation. Selecting this option allows you to edit the target VM properties during migration.

     Note:

     - Customization is disabled during some of the phases of migration such as VM preparation, clean-up phases, test migration, and after cutover initiation.
     - When the migration is in progress, the target VM properties will not reflect any changes to the source VM properties.

   - **Retain Source VM Properties**: The target VM synchronizes with the source VM properties whenever Move refreshes the source VM configuration details. Only the customizable properties are refreshed on the target. Selecting this option does not allow you to edit the target VM properties at the VM level.

     Note:

     - The source VM properties are refreshed in the following ways.
       - (Manually) When you click the **Refresh Source VM Properties** button.
       - (Automatically) When you start a migration plan.
       - (Automatically) When you initiate a cutover.
     - When you start a migration plan, Move refreshes both source VM and target VM properties by default. However, it will not refresh the target VM properties at the start of a migration plan if you modified the target VM properties after migration plan creation.

6. (Applicable if you select **Configure Target VM Properties**) Edit the properties of the target VM configuration as necessary.

   To edit the properties, under the **On Target** column, click the **edit** icon next to the property values and edit the values.

   Note:

   - Only the following properties can be edited:
     - Target VM name
     - Power state
     - Number of vCPUs
     - Number of cores per vCPU
     - Memory
   - **Name** can have a maximum of 80 characters.

7. Click **Apply**, and then click **Close**.

PAUSING OR CANCELING A VM MIGRATION

Move provides the option to pause or cancel VM migrations which are in progress. You can pause or cancel the migration of VMs at the migration-plan level or at the VM level.

Before you begin

1. Ensure that a VM migration plan is created.
2. Migration of the VMs in that migration plan is in progress.

About this task

The procedure details the steps to pause or cancel the migration at the migration-plan level; that is, the migration of all the VMs in a migration plan.

If you want to pause or cancel the migration at the VM level, that is, the migration of specific VMs in a migration plan, then see Pausing or Canceling Migration of Specific VMs in a Migration Plan on page 257.

> **Note:** If you cancel the migration till cutover state, the source VMs will be automatically cleaned up if the **Automatic** preparation mode is selected. If the preparation mode selected is **Manual**, no cleanup is performed by Move. However, you can perform manual cleanup.
>
> For more information about performing manual cleanup, see Manual Cleanup for VM Migrations on page 295.

Procedure

1. Go to the Move dashboard.
2. Identify the migration plan whose VM migration you want to pause or cancel.
3. Click the corresponding vertical ellipsis icon (under the **Status** column).
   A context menu appears with the options **Pause** and **Cancel**.
4. To pause or cancel the migration, select the appropriate option from the context menu.

Pausing or Canceling Migration of Specific VMs in a Migration Plan

Before you begin

1. Ensure that a migration plan is created.
2. Migration of the VMs in that migration plan is in progress.

About this task

The procedure details the steps to pause or cancel the migration at the VM level; that is, the migration of specific VMs in a migration plan.

Procedure

1. Go to the Move dashboard.
2. Identify the migration plan whose VM migration you want to pause or cancel.

3. Click the **In Progress** status of the migration plan (under the **Status** column).
   The resulting interface displays the summary of the ongoing migration.

4. Identify the VM(s) for which you want to pause or cancel the migration, and select the corresponding checkbox.
   Pause and **Cancel** buttons are now enabled.

5. To pause or cancel the migration of the selected VM(s), click the appropriate button.

ENVIRONMENTS AND MIGRATION PLAN MANAGEMENT

You can manage the existing environments and migration plans from the Move dashboard. You can **Refresh**, **Edit**, or **Remove** an environment and **Start**, **Pause**, **Resume**, **Cancel**, **Edit**, or **Delete** the migration plans. Move also provides the option to view the configurations of the VMs in a migration plan.

Basic Actions for Existing Environments

The **ellipses** icon in the left column of the dashboard includes the following actions that can be performed on any existing environment:

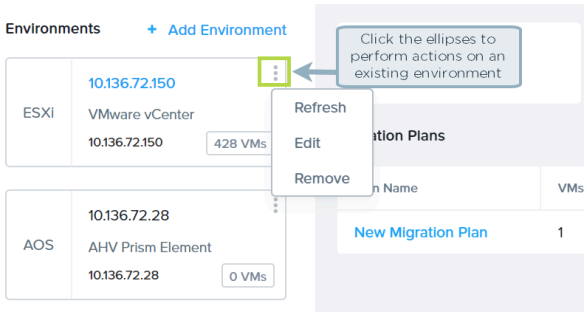

Figure 86: Actions Button for Environments Management

- **Refresh**. Refreshes to the latest changes made to the environment.
- **Edit**. Updates the existing environment.

  For editing the environment, refer to *Adding Environment* section for the respective environment.
- **Remove**. Removes the existing environment.

Basic Actions for Migration Plans

You can click the name of a migration plan and get an insight of the migrated size, status, and details of the migrated VMs.

For more information on limits for data sync during migration, refer to Limits for Data Sync during Migration on page 260.

The **Action** drop-down in the **Status** column of the **Migration Plans** page includes the following actions:

- **Start**. Starts the already created migration plan. The status changes to Validating Plan, and then to Migration in Progress. Once the migration is complete, the status changes to Migration Completed.

  > **Note:** You cannot edit or delete a migration plan once the migration is started.
- **Pause**. Pauses the migration that is in progress. The status changes to Migration Paused. This option is only available once you start a migration plan. When a migration is paused, you can either resume, cancel, or delete the migration.
- **Resume**. Resumes the paused migration. This option is only available once you pause a migration plan.

- **Cancel**. Cancels the migration. The status changes to Canceling Migration, and then to Migration Cancelled. Once the migration is canceled, you can only delete the migration plan, and cannot perform any other actions on that migration plan.
- **Edit**. Updates the existing migration plan.

  For editing the migration plan, refer to *Creating a Migration Plan* section for the respective source and target, and follow the process from Step 4.
- **Delete**. Deletes the existing migration plan.

View VM Configurations

Move provides the option to view the configurations of the source and target VMs in a migration plan.

To view the VM configurations of a migration plan, do the following:

1. In the Move dashboard, click the status of a migration plan.

   The resulting screen displays all the VMs which are part of the selected migration plan along with the configuration details of each VM.
2. Click **View VM Configs**.

   VM details window appears displaying the properties of the source and target VMs.

Limits for Data Sync during Migration

Move supports concurrent migrations and allows users to create and start the migration of multiple migration plans in parallel. To ensure the source providers are not overloaded and avoid throttling errors, internally Move maintains rules to limit the number of concurrent disks (migrations) which can be in data sync phase during migration.

Move allocates its resources to the initial set of disks and once data transfer is completed for one of the disks, it will pick up the next disk in the queue.

Following are different limits for data transfer and if any of these limits are reached, subsequent transfers of disks are not started till some of in-progress data transfers complete.

Source Type	Limit
Number of Disks per Move Instance	32
Number of Disk per ESXi Host	8
Number of Disk per Provider (ESXi/Hyper-V)	32
Number of Disk per Provider (AWS)	8
Number of Disk per Provider (Azure)	8

> **Note:** If ongoing migration disk is of size greater than 2 TB, then all other migrations of the same ESXi host will be in queue.

For example,

- If 8 disks from a given ESXi host are in data transfer phase, any new VM migration from that specific host will be queued till one of the disks completes data transfer.
- If 8 disks from 4 ESXi hosts are in data transfer phase (all 32 disk slots are in use), any new VM migration will be queued till one of the disks completes data transfer.

When the VMs are in queue for data sync and waiting for resources, Move UI shows the VM status as In Queue.

In the following image, migrations started for multiple single disk VMs in parallel. Move started the migration of 8 disks (VMs) in parallel and the rest of VMs are in state In Queue for Resources. For example, VM-10.

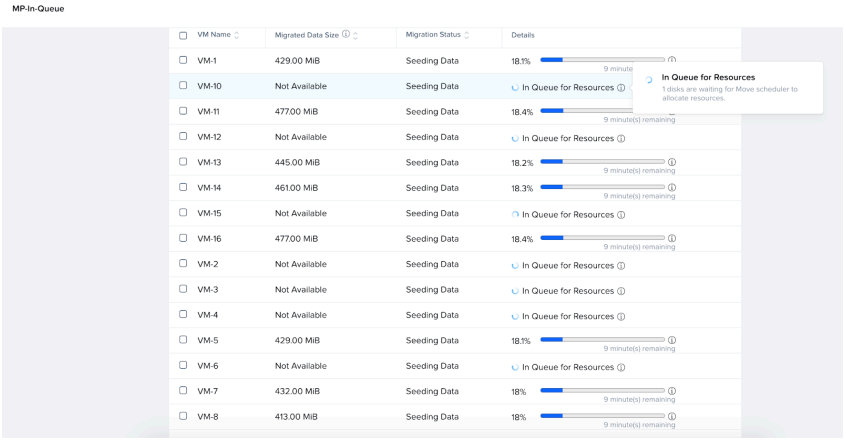

Figure 87: Parallel Migration of Single Disk VMs

Move picked up VM-2, VM-3, VM-15, and VM-10 and started data seeding. Once data seeding was completed for these four VMs, rest of the VMs were in queue. For example, VM-12 is in queue and will wait for the next available slot.

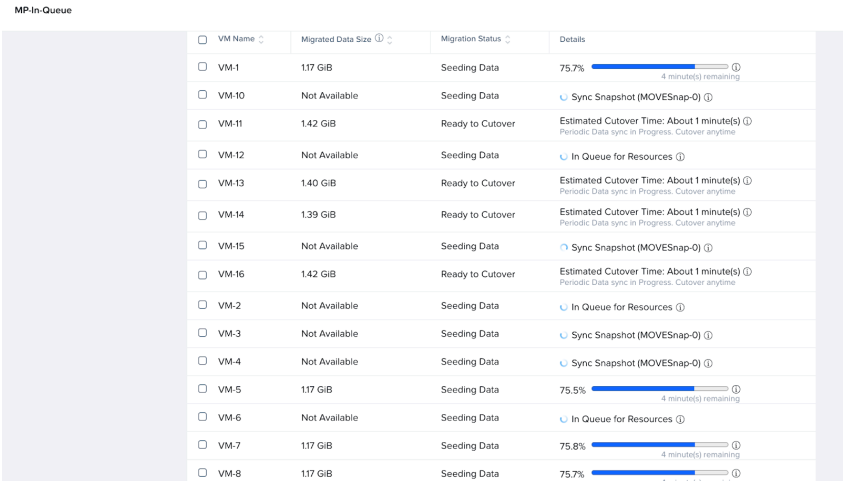

Figure 88: Status Change of VMs during Parallel Migration

FILES MIGRATION

Nutanix Move supports the migration of files from external file servers to Nutanix file servers. You can create a migration plan to seed data, perform cutover, and monitor the progress of the migration of files.

Move manages data migrations by running iterations. It runs the first iteration of a share when a migration plan (that is associated with the share) is initiated. After the first iteration is complete, changes at the source are copied to the target through subsequent iterations. These subsequent iterations are triggered automatically once every 24 hours as long as the migration plan is active.

The following topics discuss the steps to create a files migration plan, initiate the plan, and the operations that can be performed while the plan is active.

Note:

- Before upgrading a file server VM (FSVM), ensure to complete all the migrations.
- You can perform shares migration using Nutanix Files also. For more information on how to migrate shares using Nutanix Files, refer to Share Migration in *Nutanix Files User Guide*.

Requirements

This section lists the requirements for files migration.

- Supported browser: Google Chrome.
- NFS shares at the source server must have the *System Authentication* configured.
- Source NFS shares must be accessible and discoverable using the external IP address of the file server VM (FSVM).
- SMB file servers must be added with a user who has backup operator role.
- For SMB shares migration, both the source and target file servers must be present in the same domain.
- The IP address of the Move appliance must not be within the same subnet as the private network of the FSVMs.

Recommendations

This section lists the recommendations for files migration.

- Do not migrate more than 100 shares in a Move appliance.
- Perform a maximum of five share migrations concurrently.

Limitations

This section lists the limitations for files migration.

- Name of a files migration plan can have a maximum of 60 characters.
- Fully qualified domain name (FQDN) can have a maximum of 70 characters.
- Move will not perform the cleanup of the source and target servers after migration.

Unsupported Features

This section lists the unsupported features for files migration.

- Migration to a target share that has a storage limit (*MaxSize*) configured.
- NFS Shares with authentication type *None* and *Kerberos*.
- NFS Shares with *Root_Squash*.
- Shares with:
  - Protection policy
  - Tier
  - Smart Disaster Recovery (DR)
- Addition of a source file server with a User Principal Name (UPN) format username is not supported.

  Example: john.doe@nutanix.com
- Microsoft Windows file servers with deduplicated files/volumes.

Creating a Files Migration Plan

The topic discusses how to create a files migration plan using Move.

About this task

To create a files migration plan, do the following.

Procedure

1. Log on to the Move VM.
   Select Migration Type window appears with the options - **VM** and **Files**.
2. Select **Files** and click **Continue**.
   Move dashboard appears with the **Migration Plans** page selected by default. If any Files migration plans were already created in the Move VM, then they appear in this page.

   Note:

   - The **Shares** page displays details of the shares that have been migrated.
   - The **File Servers** page displays the list of file servers that have been added as source or target in the migration plans.
   - You can switch between VM migration dashboard and files migration dashboard by selecting the corresponding menu name in the menu bar.

   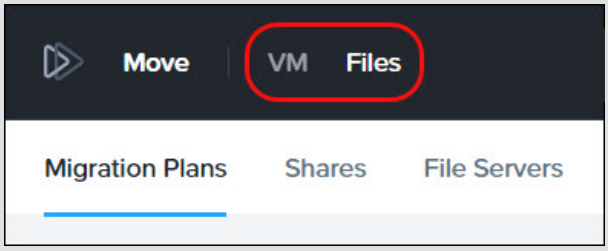

   Figure 89: Option to switch between migration dashboards
3. On the Move dashboard, click **+ New Migration Plan**.
   Enter Migration Plan Name window appears.

4. Enter a name for the new migration plan and click **Proceed**.
New Migration Plan screen appears with the **Select Source and Target** screen selected by default.

   If there are any pre-existing Files migration plans in the Move VM, then a dropdown menu appears under both the fields **Target Nutanix File Server** and **Source File Server**. The dropdown menus list the file servers that were previously saved to the Move VM. If there are no pre-existing Files migration plans in the Move VM, then the dropdown menus do not appear.

   > **Note:** You can edit the migration plan name by clicking the pencil icon next to the name.

5. In the **Target Nutanix File Server** field, do the following.

   If you want to select a file server from the previously saved target file servers, then select one from the dropdown menu (applicable if there are any pre-existing Files migration plans in the Move VM).

   If you want to add a new target file server, then perform the following steps. Else, go to step 6 on page 264.

   a. Click **+ Add New Target**. The **New Target Nutanix File Server** window appears.

   b. Under **Name**, enter a name for the target file server.

   c. Under **FQDN Details/Address**, enter the Fully Qualified Domain Name (FQDN) details or the IP address of the target file server.

   d. Under **User Name** and **Password**, enter the login credentials of the target file server in the respective fields.

      The credentials must be of a REST API user. To create a REST API user, see Managing REST API Roles topic in *Nutanix Files User Guide*.

   e. Click **Add**.

6. In the **Source File Server** field, do the following.

   If you want to select a file server from the previously saved source file servers, then select one from the dropdown menu (applicable if there are any pre-existing migration plans in the Move VM).

   > **Note:** When a source file server and a target file server are added to a Files migration plan, they are tightly coupled. If you want to add a source server that is already coupled with another target server, then the source will not appear under the dropdown menu options under **Source File Server**. You must add it as a new source file server under a different name.

   If you want to add a new source file server, then perform the following steps. Else, go to step 7 on page 265.

   a. Click **+ Add New Source**. The **New Source File Server** window appears.

   b. Under **Name**, enter a name for the source file server.

   c. Under **FQDN Details/Address**, enter the Fully Qualified Domain Name (FQDN) details or the IP address of the source file server.

   d. Under **User Name** and **Password**, enter the login credentials of the SMB File Server in the respective fields.

      > **Note:** If you have only NFS shares, then do not enter the login credentials. This will result in failure to add the source file server as Move uses these credentials to discover SMB shares.

   e. Click **Add**.

7. For **Replication Type**, select one of the following options:

   - **Sync**: Replicates files from the source to the target. It replaces the files in the target with the updated ones from the source. It also replicates any files in the source that are missing in the target.

     It does not delete any files in the target.

   - **Mirror**: Also replicates files from the source to the target. Similar to *Sync*, it replaces the files in the target with the updated ones from the source. It also replicates any files in the source that are missing in the target.

     However, it deletes the files that are present only in the target and not in the source.

   Click **Next**.

   Select Shares screen appears.

8. In the **Select Shares** screen, do the following:

   a. Click **+ Add Share**.

   b. Clicking on the fields, **Source** and **Target**, displays the discovered shares in the source and target file server respectively. Select the required source and target shares under the respective fields.

      Note:

      - Migration is supported between primary protocols only.

        Hence, for a multi-protocol source share, the protocol of the associated target share must be the same as the primary protocol of the source share. Otherwise, the migration will fail.
      - Move lists only the primary protocol of the target share.
      - If you want to enter a custom path for these fields, then you can do so manually. Both the fields are editable. Use forward slash as the directory separator in the path.
      - If you enter a custom source path under **Source**, ensure that the path is present in the source share.
      - If a custom target path entered under **Target** is missing in the target share, then Move will create it in the target share.

   c. From the **Replication Type** dropdown menu, you can change the replication type.

   d. Under **Actions**, click the check mark to save the Share mapping entry.

   e. To add another Share for migration, go to step 8.a on page 265.

      Note: You can edit or delete a Share mapping entry by clicking the pencil icon or the bin icon respectively under **Actions**.

   Click **Next**.

   Move validates the Share mapping entries. If the validation is successful, **Migration Settings** screen appears.

   If Move encounters any issues during validation, it warns that the validation has failed and indicates which Share mapping entries have issues.

9. In the **Migration Settings** screen, do the following if you want to schedule the migration. If you do not want to schedule the migration, click **Next** and go to step 10 on page 266.

   a. Select **Schedule Migration**. Fields to select the date and time appear.

   b. Select the date and time under the respective fields based on when you want to schedule the migration.

   c. Click **Next**.

   Summary screen appears.

10. In the **Summary** screen, review the plan migration summary, and choose one of the following.

    > **Note:** The field **No. of Shares in Migration Plan** indicates the number of shares in this migration plan. Clicking the shares-count displays the summary of these shares.

    » **Back**: Click this option to go back to the previous screens.

    » **Save**: Click this option to save the migration plan. Move dashboard appears with the **Migration Plans** page selected by default. The saved migration plan is displayed in a table.

    » **Save and Start**: Click this option to save the migration plan and start the saved plan. Move starts the migration plan. The seeding process for the migration begins.

What to do next

- If you have created a migration plan, then the next step is to start the migration plan. For more information, refer to Starting a Files Migration plan on page 266.

  Refer to the same topic for the next steps even if you have started the migration plan.

Starting a Files Migration plan

The topic discusses how to start a files migration plan in Move. It also includes information on pausing, resuming, and canceling the migration of a share.

Before you begin

- Ensure that you have a files migration plan available for migration. For information on creating a files migration plan, refer to Creating a Files Migration Plan on page 263.
- If you have already started the migration plan while saving it, then you can skip the first two steps in the following procedure.

About this task

To start a migration plan, do the following:

Procedure

1. In the Move dashboard (for files migration), identify the migration plan you want to start.
2. Click the vertical ellipsis icon next to **Status** of the identified migration plan. From the dropdown menu, select **Start** to start the migration plan.
   Move starts validating the migration plan. Once it is validated, Move starts seeding data. During data seeding, it starts running the iteration to copy files from the source server to the target server.

3. To monitor the progress of the migration plan, click **Status** of the plan.
   The resulting **Migration Plan** screen provides information about the migration status of the shares in the plan.

   Note:

   - The seeding process can take several minutes depending on the volume of data being migrated.
   - To pause or cancel the migration of any shares, select those shares from the table, click **Actions**, and select the corresponding menu item from the dropdown menu.
   - To resume the migration of shares that were paused, select those shares, click **Actions**, and select **Resume** from the dropdown menu.

4. Wait for the iteration to complete. Once the iteration completes, the **Migration Plan** screen indicates the iteration completion status.

   If all the files have been migrated successfully, then the status fields display the following messages.

   - **Last Sync Status**: Completed

   - **Status**: Ready to Cutover

   If some files have failed to migrate successfully during the iteration, then the status fields display the following messages.

   - **Last Sync Status**: Completed with Failed Files

   - **Status**: Ready to Cutover with Failed Files

5. (Optional) This step is applicable only if the migration plan has completed the iteration with failed files.

   Perform the following steps to view the list of files that failed to migrate in the iteration.

   a. Click the source share/path that completed migration with failed files.
      Sync iterations window appears with details of all the iterations of the source share/path. The **Status** field of the iteration that completed with failed files displays the message Completed with Failed Files.

   b. Click **Download Failed Files** link under **Status** of the iteration that completed with failed files.
      A CSV file is downloaded. Refer to this file to view the list of files that failed to migrate during the iteration.

   c. Click **Close** to close the **Sync iterations** window.

What to do next

- After the migration plan has completed the iteration, you can do one of the following:
  - Perform cutover of the shares in the migration plan.
  - Keep the migration plan active.
- To perform a cutover of a share, refer to Performing a Migration Cutover on page 268.
- If you keep the migration plan active, then Move will run iterations periodically to ensure that changes at the source server are copied to the target server. The iterations are triggered automatically once every 24 hours.
- You can manually initiate a sync iteration (for a share). This option is not available if an iteration is already in progress, that is, if the share is in the *incremental data seeding* state.

  To initiate a sync iteration for a share, select the share, click **Actions**, and select **Sync** from the dropdown menu.

- You can also cancel an ongoing iteration for a share. This option is available only when an iteration is in progress, that is, if the share is in the *incremental data seeding* state.

  To cancel an ongoing iteration for a share, select the share, click **Actions**, and select **Cancel Current Iteration** from the dropdown menu.

  > **Note:** This cancels only the current iteration of the share. The migration plan remains active. The next automatic iteration of the share will run as scheduled.

Performing a Migration Cutover

After the files migration plan is started and the iteration is complete, you can perform the cutover of shares in the plan.

Before you begin

Ensure that you have a files migration plan that has completed its iteration. For more information on how to start a files migration plan and run it to completion, refer to Starting a Files Migration plan on page 266.

About this task

To perform the cutover of shares in a migration plan, do the following:

Procedure

1. In the Move dashboard (for files migration), identify the migration plan that has the shares you want to cutover. Click **Status** of that plan.
   The **Migration Plan** screen appears.
2. In the **Migration Plan** screen, select the shares that you want to cutover, click **Actions**, and select **Cutover** from the dropdown menu.
   The cutover process begins. It will trigger the last sync iteration and start seeding data. The status fields reflect the progress.

   > **Note:** To view the information of the current iteration and the sync iteration history of a share, click on the share. A pop-up window will populate this information.
3. Wait for the cutover to complete. You can monitor the cutover progress of the shares in the **Migration Plan** screen.
   Once the cutover completes, the status fields reflect the completion status.

   If all the files have been migrated successfully to the target server, then the status fields display the message Completed. Go to 9 on page 269.

   If some of the files failed to migrate to the target server, then the status fields display the message Completed with Failed Files. Move provides the option to retry migrating the files that failed to migrate during cutover.

   If you want to retry migration of the failed files, then go to 4 on page 269. If you do not want to retry migration and want to mark the share migration as complete, go to 7 on page 269.

4. Select the shares that have completed cutover with failed files, click **Actions**, and select **Retry Failed Files** from the dropdown menu.

   Note:

   - The option to retry migration of files is available only to those shares that have failed files after the cutover.
   - Move provides the option to retry the migration of failed files only once per share.

   A pop-up window appears prompting for confirmation.

5. Click **Confirm**.
   Move reattempts the migration of the failed files. The status fields will reflect the status of the migration progress.

6. Wait for the migration to complete. You can monitor the migration progress of the files in the **Migration Plan** screen.
   Once the migration completes, the status fields reflect the completion status of the migration.

   If all the files have been migrated successfully to the target server, then the status fields display the message Completed. Go to 9 on page 269.

   If some of the files failed to migrate to the target server again, then the status fields display the message Completed with Failed Files. You cannot reattempt the migration of the failed files again.

7. Select the shares that completed with failed files, click **Actions**, and select **Complete** from the dropdown menu.
   A pop-up window appears prompting for confirmation.

8. Click **Confirm**.
   The selected share is now marked as Completed.

9. Click the vertical ellipsis icon next to **Status** of the share that completed migration and select **Delete** from the dropdown menu.
   A pop-up window appears prompting for confirmation.

10. Click **Continue**.
    The share is deleted from the migration plan.

Performance Matrix for Large Shares Migration

This section provides the performance matrix of Move during the migration of large shares.

The following table provides the performance numbers from the Move lab. The performance was tested considering the parameters listed in the table.

Table 22: Performance Numbers for Files Migration

Total Migration Size	Number of Files	Time Taken for Data Seeding	Time Taken for Cutover	Source I/O	Data Churn	FSVM details
1 TB	18 million	17 hours	1 hour, 30 minutes	N/A	N/A	Nodes: 3 CPU: 8 Memory: 12 GiB
1 TB	450	6 hours	40 seconds	N/A	N/A	Nodes: 3 CPU: 8 Memory: 12 GiB

MOVE ADMINISTRATION

You can upgrade, uninstall, and modify the Move configurations.

Move Upgrade Management

You can upgrade to a new version of Move to use the latest available features. The dashboard displays the current version and an option to upgrade the Move VM to the latest version when a new version is available for upgrade.

Note:

- You can now upgrade to the latest Move version only from the last three major releases along with the minor releases during that interval; that is, for Move 4.8.0, upgrade is supported from 4.5.0, 4.5.1, 4.5.2, 4.6.0, and 4.7.0.

  Until Move 4.7.0, you could upgrade to a newer Move version only from the last two major releases along with the minor releases during that interval.
- To upgrade to the latest Move version from a version that is older than the last three major releases, you need to perform a chain upgrade; that is, continue upgrading Move to the highest version possible until you have the required Move version installed.

  For example, to upgrade Move 3.3.0 to Move 4.8.0, continue upgrading Move to the highest version possible until you have Move 4.8.0 installed.
- From Move 4.0.0 onward, Move requires 8 GiB memory. If you are using older versions of Move, upgrade the memory to 8 GiB before upgrading Move version to 4.0.0.
- It is rarely observed that upgrade from 3.2.0 proceeds even though it is not allowed and restricted. Note that the preceding upgrade flow is not supported.

You can upgrade Move in the following two ways:

- Online Upgrade - If the Move VM is already connected to the Internet, the system automatically detects the latest version and you can perform the one-click upgrade operation.

  For more information, refer to Upgrading Move Online on page 271.
- Offline Upgrade - If the Move VM is not connected to the Internet, you can upgrade the server in offline mode by uploading the binary.

  For more information, refer to Upgrading Move Offline on page 273.

Note:

- Both online and offline upgrade require some application downtime to ensure that Move loses no data during the upgrade.
- To reset the upgrade state of the updater container, run the following command:

  ```
  updater-op -t -r -y
  ```

Upgrading Move Online

You can upgrade Move to the latest available version by using the one-click upgrade method.

Before you begin

If the Move VM is deployed behind the firewall, online update will not work and you will also not receive the update notifications. To perform an online upgrade and to receive the update notifications, add the firewall rules to accept the IP addresses of the following URLs. You can also use the nslookup command to get the list of IP addresses.

- auth.docker.io
- registry-1.docker.io
- index.docker.io

About this task

To upgrade Move, do the following:

Procedure

1. Log on to the Move UI by using a supported web browser.

2. In the gear icon drop-down list, click **Upgrade Software**.

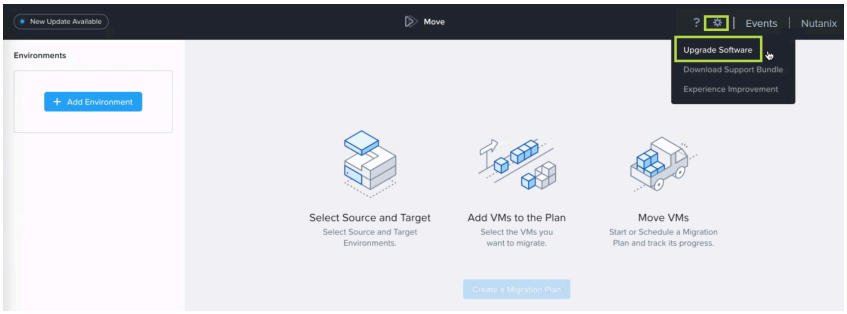

Figure 90: Settings Menu: Upgrade Software

The **Upgrade** dialog box displays the latest available version of Move. Also, a notification New Update Available appears.

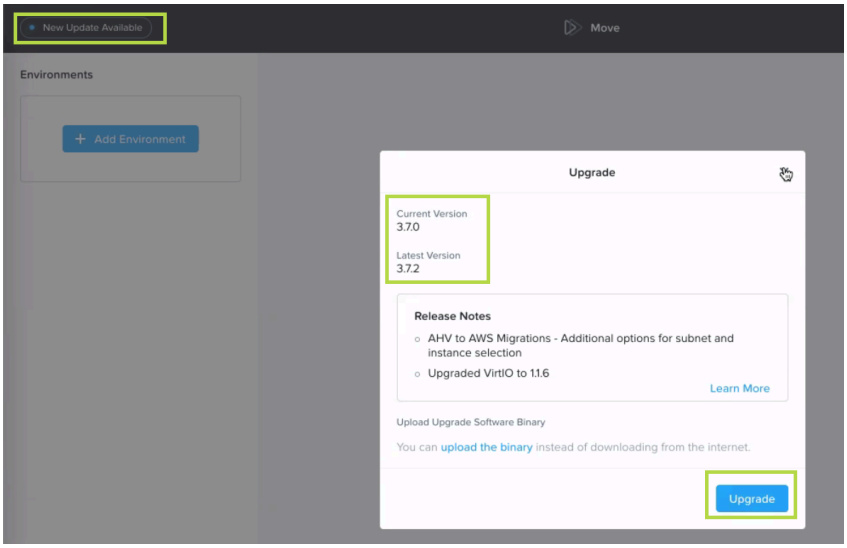

Figure 91: Upgrade Dialog Box

3. Click **Upgrade**.
   A confirmation message is displayed after the upgrade is complete.

4. Log on again to access the updated version of Move.

Upgrading Move Offline

You can upgrade Move VM to the latest available version without connecting to the Internet by uploading the binary.

Before you begin

When you are connected to the Internet, you must download the Move offline upgrade package from Nutanix Support Portal, and unzip it to keep the move-offline-upgrade-x.x.x.x.tar.gz binary file and the update_info.json metadata file. Here, x.x.x.x is the version number of the file.

> **Note:** You can perform offline upgrades from Move 3.1.0 version onward. Offline upgrades from 3.0.3 and earlier versions are not supported.

About this task

To upgrade Move, do the following:

Procedure

1. Log on to the Move UI by using a supported web browser.
2. In the gear icon drop-down list, click **Upgrade Software**.

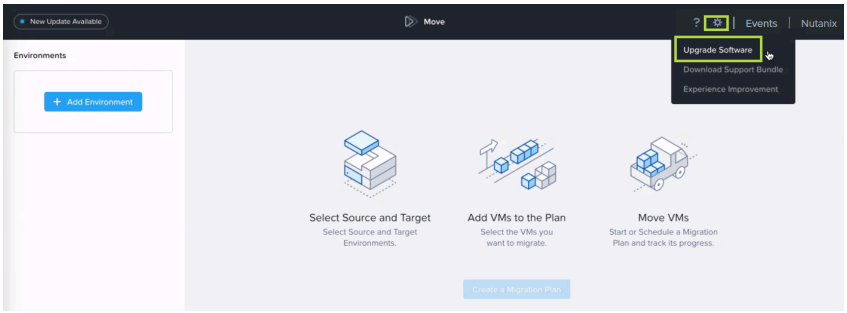

Figure 92: Settings Menu: Upgrade Software

The **Upgrade** dialog box displays the latest available version of Move and a section to perform an offline upgrade by uploading the binary file. Also, a notification New Update Available appears.

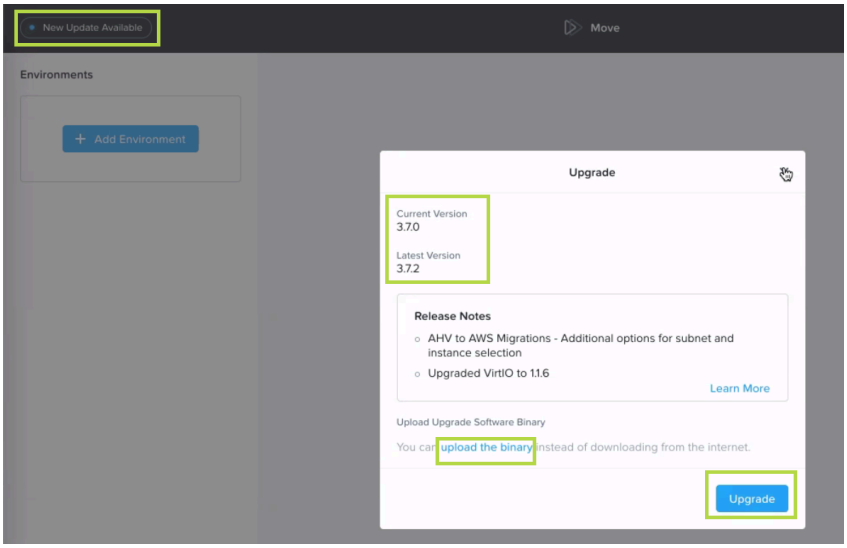

Figure 93: Upgrade Dialog Box

3. Under **Upload Upgrade Software Binary** section, click **Upload the Binary**.

   An **Upgrade** page appears to choose the binary file and metadata file

4. Click **Choose** to select the downloaded binary and metadata file.

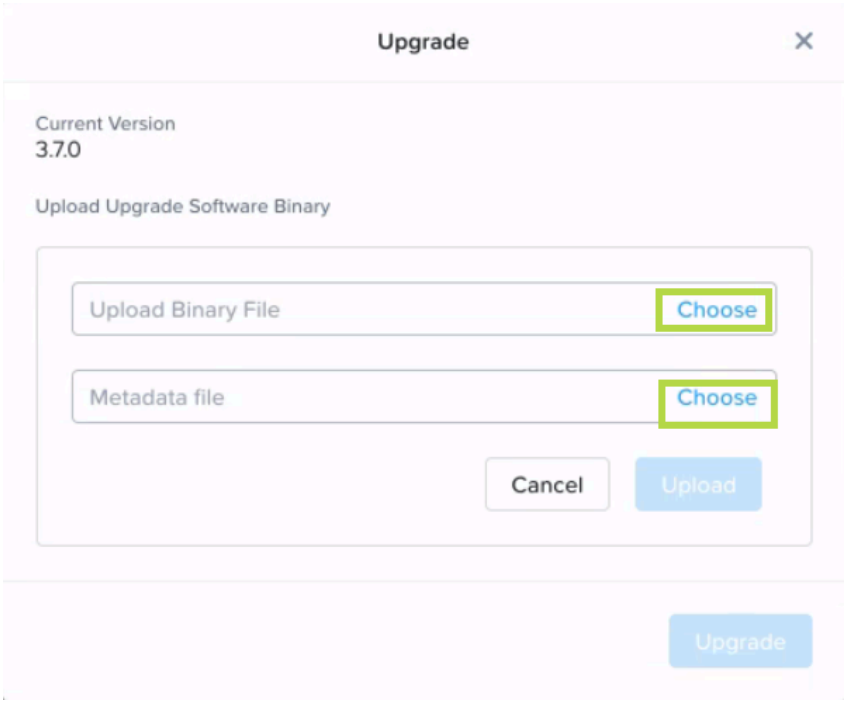

Figure 94: Upload Software Binary for Upgrade

5. Click **Upload**, and then click **Upgrade**.
   The binary and metadata files are now uploaded and upgrade is performed.
6. Log on again to access the updated version of Move.

Undeploy Move

You can undeploy Move by undeploying the VM through the CLI or deleting the VM from Prism Element UI.

> **Note:** Nutanix recommends undeploying Move with the CLI instead of deleting the Move VM. This action allows Move to automatically clean the NFS allowlist exceptions created during the Move deployment.

Undeploy Move (CLI)

You can undeploy Move by undeploying the VM through the CLI.

About this task

To undeploy the Move VM by using the CLI, do the following:

Procedure

1. Open the local CLI of your operating system.

2. Browse to the deployment utility folder location and run the deployment utility.

   ```
   $ binary_name -c cluster-virtual_ip_address
   ```

   Replace the *binary_name* with the name of the binary for your operating system.

   For more information, refer to Move Software Package on page 11.

   Replace *cluster-virtual_ip_address* with the FQDN or the IP address of the cluster.

   > **Note:** Use the `-u` parameter to log on. For more information, run the command `./binary_name --help`.

3. Log on to the Prism Element with the admin user credentials.
4. To undeploy the Move VM, run the following command and when prompted for confirmation, enter Y:

   ```
   admin@pevm$ undeploy-vm
   ```

   > **Note:** The name of the Move VM must be **Nutanix-Move**; otherwise, the following command fails.

   The Move VM undeploy can take a few minutes.

Undeploy Move (Prism Element UI)

You can undeploy Move by deleting the VM from Prism Element UI.

About this task

To delete the Move VM from the cluster, do the following:

Procedure

1. Log on to the Prism Element UI of the cluster with the admin user credentials where Move VM is deployed.
2. Go to **Entity menu** > **Virtual Infrastructure** > **VMs** .
3. Click **List** tab from the left navigation.
4. Select the VM name **Nutanix-Move**.
5. Click **Actions**, and then click **Delete**.
   The Move VM is deleted from the cluster.

Changing the Database Password

Nutanix recommends that you secure your database by changing the password of the database.

About this task

To change the password of your database, do the following:

Procedure

1. Open the Move CLI.
2. Run the following command as the root user, and then press **Enter**.

   ```
   root@move on ~ $ chdbpasswd
   ```

   The following message appears. Changing the postgresql DB password for user 'admin' and database 'datamover'.
   WARNING: This operation will restart postgresql and mgmtserver services.

3. Type the password and press **Enter**.

```
New password:
```

4. Confirm the password and press **Enter**.

```
Retype new password:
Writing updated config to file: /opt/xtract-vm/resources/config.toml
Restarting postgresql and mgmtserver to reflect the password change...
```

The password is updated successfully and the new password is stored in /resources/config.toml.

5. Verify if the password has changed.

a. Move to the Docker container of the management server.

```
root@move on ~ $ mgmtserver-shell
```

b. Move to the resources directory.

```
root@move on ~ $ cd resources
```

c. Open the config.toml file to verify if the password is changed.

Resetting Admin Password

You can reset the password for the user admin. You need the password to log into the Move CLI.

About this task

To reset the admin password for Move VM, do the following:

Procedure

1. Open the VM console from Prism Web Console
2. Restart the Move VM
3. When the VM starts up, press any one of the arrow keys to view the entry as shown in the following image:

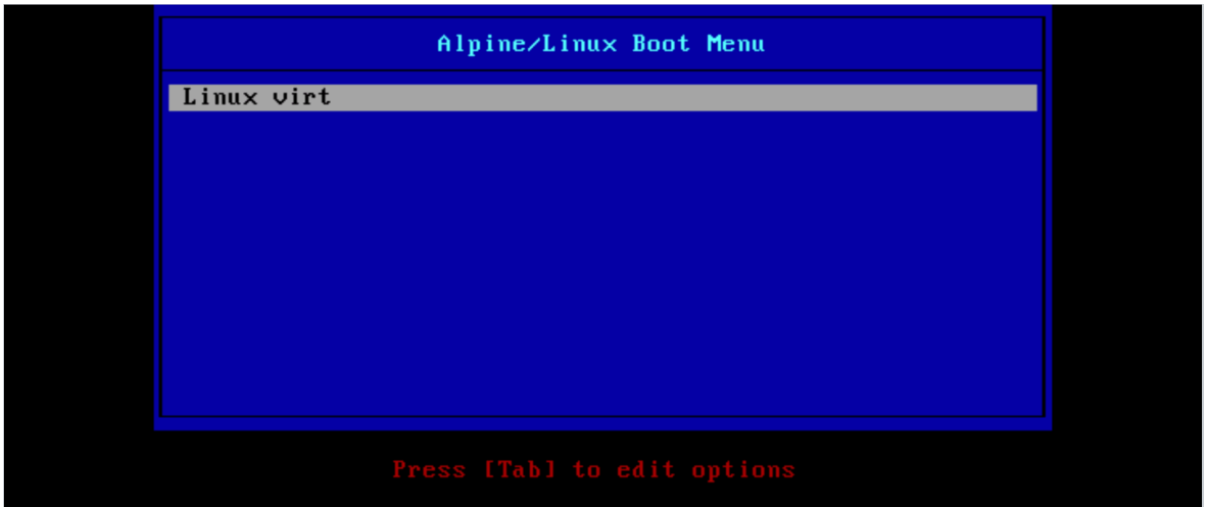

Figure 95: Boot Menu

4. Press Tab to edit the entry.

5. Append init=/bin/bash at the end of line as shown in the following image:

Figure 96: Command Append

6. Press Enter to start up the VM.
   The VM starts up at the bash prompt.
7. Remount the root file system as *read/write* using the following command as shown in the image in Step 9:

   mount -n -o remount,rw /
8. Change the password for the admin using the following command as shown in the image in Step 9:

   passwd admin
9. Enter the New password and Retype new password.

Figure 97: New Password Reset

10. Restart the VM with the following command:

    reboot -f

Resetting Web Login Password

You can reset your password if you have forgotten it.

About this task

> **Note:** The password reset command is unavailable when invoked from a remote machine.

To reset the web login password, do the following:

Procedure

1. SSH to the Move VM as an admin.
   Refer to Accessing Move VM with SSH on page 21.
2. Switch to the root user by entering the password of the Move VM.

```
admin@move on ~ $ rs
[sudo] password for admin:
```

3. Go to the Move bin.

```
root@move on ~ $ cd /opt/xtract-vm/bin
```

4. List the files within the bin and open the CLI.

```
root@move on ~ $ cli-linux-amd64
```

5. In the CLI, run the password reset command.

```
localhost (Move) » password reset
Enter new password for 'Move' user nutanix(10):
Reenter password for 'Move' user nutanix(10):
Successfully reset the password [OK]
localhost (Move) »
```

Output displays a success message, Successfully reset the password [OK].

Changing Web Login Password

If you know your password and you want to change your password, you can use the following procedure.

About this task

> **Note:** The password reset command is unavailable when invoked from a remote machine.

To change the web logon password, do the following:

Procedure

1. SSH into the Move VM as an admin.
2. Switch to the root user by entering the password of the Move VM.

```
admin@move on ~ $ rs
[sudo] password for admin:
```

3. Go to the Move bin.

```
root@move on ~ $ cd /opt/xtract-vm/bin
```

4. List the files within the bin and open the CLI.

```
root@move on ~ $ ./cli-linux-amd64
```

5. In the CLI, change the password.

```
localhost (Move) » password change
Enter password for 'Move' user nutanix(10):
Enter new password for 'Move' user nutanix(12):
Reenter password for 'Move' user nutanix(12):
Successfully changed the password [OK]
localhost (Move) »
```

Output displays a success message, Successfully changed the password [OK].

Configuring Time-Out for Source Inventory Refresh

You can change the default time-out and configure a new time-out for source inventory refresh.

About this task

To configure a time-out for source inventory, do the following:

Procedure

1. Log on to the Move CLI.

2. Switch to the root user by entering the password of the Move VM.

```
admin@move on ~ $ rs
[sudo] password for admin:
```

3. In the command line, run the following command.

```
root@move on ~ $ vi /opt/xtract-vm/bin/docker-compose.yml
```

4. In the file, go to **services** > **srcagent** > **environment**, add the following line.

```
- VC_INV_TIMEOUT_MINS=10
```

5. Save and exit the file.

6. Restart the source agent service.

```
root@move on ~ $ docker restart bin_srcagent_1
```

7. Refresh the inventory.

8. Verify the changes in the srcagent log.

Changing SSH Port

Move provides the option to access a Move VM with SSH. It also provides the option to change the SSH port number. You can do this either through an SSH client or through the VM console. The following procedure details the steps to change the SSH port number.

About this task

To change the SSH port number, perform the following:

Procedure

1. Perform one of the following depending on whether you want to use an SSH client or the VM console.
   - » (SSH client) Open an SSH client such as PuTTY (for Windows) or Terminal (for Mac).
   - » (VM console) In the Prism Element UI, select the Move VM and click **Launch Console**.
2. Log on to the Move VM as admin using SSH.
   `$ ssh admin@<move_vm_ipaddress>`
   Replace `<move_vm_ipaddress>` with the IP address of the Move VM.
   The Command Line Interface (CLI) prompts for the password of the VM.
   admin@`<move_vm_ipaddress>`'s password:
3. Type the password and press Enter.
   You will be logged on to the Move VM.
4. Switch to the root user.
   `admin@move on ~ $ rs`
   The CLI prompts for the admin password.
   [sudo] password for admin:
5. Type the admin password and press Enter.

6. Change the SSH port number of the Move appliance.

   ```
   root@move on ~ $ change-ssh-port
   ```

   The CLI asks whether you want to change the SSH port of the appliance.

7. Input y and press Enter.
   The CLI displays the current SSH port number. It also prompts you to provide the new port number for the SSH port.

8. Type the new port number for SSH and press Enter.

   > **Note:** If you provide an invalid input for the new port number, then the CLI will prompt you to enter a valid one.

   The new port number will be successfully configured as the SSH port of the Move appliance.

   - (SSH client) Since the SSH port number has changed, you will be logged out of the current session. To log on to the Move VM again using SSH, you must use the new SSH port number. Go to Step 9.

   - (VM console) You will remain logged on to the current session.

9. (Applicable to SSH client only) Open a new SSH client window and log on as admin using the new SSH port number.

   ```
   $ ssh -p <new_SSH_port> admin@<move_vm_ipaddress>
   ```

   Replace as follows:

   - `<new_SSH_port>` with the new SSH port number.
   - `<move_vm_ipaddress>` with the IP address of the Move VM.

   The CLI prompts you for the password of the VM.

   admin@`<move_vm_ipaddress>`'s password:

10. (Applicable to SSH client only) Type the password and press Enter.
    You will be logged on to the Move VM using the new SSH port number.

MOVE EVENTS OVERVIEW

Events dashboard displays the list of events happening in Move.

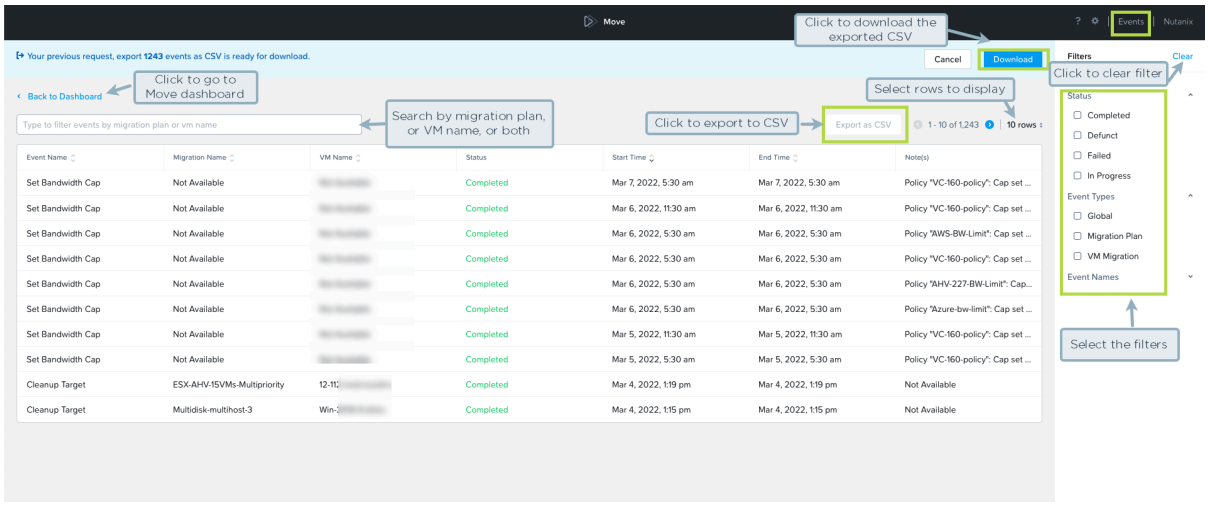

Figure 98: Events Dashboard

The following are the types of events.

- **Global**: Events happening at the appliance level that are not related to migration. For example, Service Restarted, Add Provider.
- **Migration Plan**: Events happening at the migration plan level. For example, Create Migration Plan, Start Migration Plan.
- **VM Migration**: Events related to VM migration level. Displays the tasks happening on a specific VM. For example, Start VM Migration, Create Base Snapshot.

> **Note:** Move supports 500K events as the maximum retention limit. Once this limit is reached, the oldest events get purged to store new events.

You can filter these events based on the following. Go to the **Filters** pane in the right, and select the options.

- **Status**: Choose one or more status from **Completed**, **Defunct**, **Failed** and **In Progress**.
- **Event Type**: Choose one or more event type from **Global**, **Migration Plan**, and **VM Migration**.
- **Event Name**: Choose one or more event name from the list.

  You can also search a particular **Event Name** and select it.

Additionally, you can also use the top filter to filter by either migration plan name, or VM name, or both. For example, if you want to search a VM with name VM-255 in a specific migration plan (MP1), then enter the name of the VM (VM-255) and the name of the migration plan (MP1) to filter the list of events specific to VM-255 in MP1.

Events changes the status from **In Progress** to **Complete**, when the events are successful. Events can also move to **Failed** or **Defunct** status if not successful. In case any service inside the Move appliance is restarted, Move may not be able to collect the information of the ongoing event in that service, and the status of that event will be marked as **Defunct**

For example, the **VM Migration** event of the **Migration Plan** event type will be **In Progress** state until that specific VM migration is **Completed** or **Failed**. If the **VM Migration** events fails and you retry, a new event is created.

You can also view the start and end time of a particular event.

If you want to export the events as a CSV file, click **Export to CSV**. Once the file is exported, you will see a notification to either download or cancel. You can click **Download** to download the exported file in your device. You will be notified once the download is complete. The downloaded file lists all the events (which are shown on the user interface taking into account any filters applied). For example, if on the Events page, the total number of records are shown as 1249, then the CSV file will have all 1249 events sorted by event creation times with latest events shown first. The name of the downloaded file will be move-events-*`date`*-*`time`*.csv

> **Note:** You can only request one download at a time. The **Export to CSV** button will be disabled if the previous download is already in progress.

To go back to the Move Dashboard, click **Back to Dashboard** on the top-left corner.

VIEW METRICS

View Metrics feature provides the option to view the health status of a Move appliance. Using Grafana dashboards, Move provides a set of pre-defined dashboards. Each of them conveys the health status of a Move appliance through various metrics. This topic describes how to view the health status of the Move appliance.

About this task

Perform the following steps to view the health status of a Move appliance using **View Metrics**.

Procedure

1. Log on to the Move VM.
2. Click the gear icon from the top-right corner and then select **View Metrics** from the drop-down list.
   A window opens displaying the health status of the system using Grafana dashboards. The dashboard for **Docker and system monitoring** appears by default.
3. Click the **Dashboards** icon on the left pane and select **Manage**.
   A browser window opens with the list of pre-defined dashboards.
4. Select the required dashboard from the list.
   The corresponding dashboard opens in a new browser window.

Create new Grafana Dashboards

While Move provides pre-defined Grafana dashboards, it also provides the option to create new ones. You can create new Grafana dashboards when logged on as an administrator.

About this task

Perform the following steps to modify a pre-defined Grafana dashboard.

Procedure

1. Log on to the Move VM.
2. In the address bar of the web browser, type the following URL and press **Enter**.
   https://`<nutanix-move-ip>`:3000
   Where, `<nutanix-move-ip>` is the IP address of the Move VM.
   A warning about security certificate appears in your browser.
3. Accept the security certificate warning and proceed.
   Grafana welcome page appears, prompting for the login credentials.
4. Enter admin as the username and password. Click **Log in**.
   You are logged on as administrator.

   Caution: Ensure to change the password for the Grafana dashboard after login. Anyone who knows the IP address of the Move VM can log on to the dashboard using the default login credentials.
5. Modify the dashboard as required.
   For example, you can edit or remove panels of the metrics shown in the dashboard since you have administrator rights.

6. Click **Dashboard settings** (gear icon) and then select **Save as new dashboard**.
   Save dashboard as window appears.

7. Input a name for the dashboard in the **Dashboard name** field.

8. Select a folder from the **Folder** drop-down menu.

9. Click **Save**.
   The modified dashboard is saved as a new dashboard.

MOVE BANDWIDTH THROTTLING

Bandwidth throttling feature allows you to define the maximum available bandwidth for data transfer for a specific source provider. You will also be able to schedule the application of this bandwidth threshold. Based on the specified bandwidth threshold and schedule, Move will limit the data migration throughput rate for that source provider.

> **Note:** The bandwidth threshold does not imply that the network traffic will be equal to the specified limit. It only represents the upper limit of the network traffic, that is, the network traffic from the source provider will not consume more than the specified bandwidth threshold.

You can access this feature from the Move dashboard and create bandwidth throttling policies. You will also be able to edit, deactivate, or remove these policies. Changes to the policies will be applied immediately to the source provider, including any ongoing migration plans.

Create Bandwidth Throttling Policy

Move provides the option to create a bandwidth throttling policy for a source provider. This topic details the procedure to create a policy.

About this task

To create the bandwidth throttling policy, do the following:

Procedure

1. Log on to the Move UI.
2. Click the gear icon from the top-right corner and then select **Bandwidth Throttling** from the drop-down list. The resulting page displays the existing bandwidth throttling policies, if there are any already in place. It also provides options to create new bandwidth throttling policies.
3. Click the **Create Policy** button.
   Create Bandwidth Throttling Policy window appears.
4. On the **Create Bandwidth Throttling Policy** window, do the following in the fields mentioned below:
   a. **Policy Name**: Enter a name for the policy.
   b. **Select source**: Select the source from the drop-down menu options.
      A source can be part of a single bandwidth policy only. If a source in the drop-down menu options is grayed-out, it implies that the source is already part of a policy.

      > **Note:** Hyper-V is not supported as a source for bandwidth throttling.

   c. **Frequency**: Select the frequency.
      - To activate the bandwidth policy on all the days of the week, select the **All Days** option.
      - To activate the bandwidth policy for a specific time duration (on a weekly basis), select the **Specific Duration** option. Select the time ranges in the **From** and **To** fields to specify the time frame. Select the days from the **Days of the week** drop-down options.

        > **Note:** The time zone where the Move appliance is located is considered.

   d. **Bandwidth Limit**: Enter the bandwidth limit in MBps. The highest bandwidth value that can be entered is 100000 MBps.

5. Click **Next**.
   The resulting window shows the overview of the options selected.

6. Click **Save and Activate**.
   The policy is saved and activated. The **Status** of the policy is shown as **Active**. Once activated, any migration that happens from the source will not consume more than the bandwidth limit of the policy during the scheduled time frame.

   Note: If you want to only save the bandwidth policy and not activate it, select the drop-down icon next to **Save and Activate** and click **Save Only**.

Update Bandwidth Throttling Policy

Move provides the option to update bandwidth throttling policies. This topic details the procedure to edit a policy by adding schedules.

Before you begin

Ensure that you have at least one bandwidth throttling policy created in Move.

About this task

To update the bandwidth throttling policy, do the following:

Procedure

1. Log on to the Move UI.

2. Click the gear icon from the top-right corner and then select **Bandwidth Throttling** from the drop-down list.
   The resulting page displays the existing bandwidth throttling policies. It also provides options to edit these policies and create new ones.

3. Click the **Edit** option corresponding to the bandwidth throttling policy that you want to update.
   Update Bandwidth Throttling Policy window appears for the selected policy.

4. Click **Add New Schedule**.
   Options to include an additional schedule to the bandwidth throttling policy appears.

5. Select the time ranges from the **From** and **To** fields to specify the time frame.

6. Select the days from the **Days of the week** drop-down options.

7. Enter the bandwidth limit in the **Bandwidth Limit** field in MBps. The highest bandwidth value that can be entered is 100000 MBps.

8. Click **Save**.
   The new schedule is included under **Bandwidth Throttling Schedule(s)** of the policy.

9. Click **Update and Activate**.
   The policy is updated and activated. The new schedule is included to the policy and the **Status** of the policy is shown as **Active**. Once activated, any migration that happens from the source will not consume more than the bandwidth limit during the scheduled time frame.

   Note: If you only want to update the bandwidth policy and not activate it, select the drop-down icon next to **Save and Activate** and click **Save Only**.

Monitor Bandwidth Usage

Move provides the option to monitor the bandwidth usage in the Move appliance. This topic details the procedure to navigate to the dashboard that shows this data.

About this task

To monitor the bandwidth usage, do the following:

Procedure

1. Log on to the Move UI.
2. Click the gear icon from the top-right corner and then select **Bandwidth Throttling** from the drop-down list. The resulting page displays the existing bandwidth throttling policies, if any. It also provides options to edit the existing policies and create new ones.
3. Click **View bandwidth usage**.
   A dashboard opens displaying the details of bandwidth usage and real-time throughput on the Move appliance.

Virtual Machine (VM) Priority

Nutanix Move provides the option to set the priority of virtual machines (VMs) for migration. This topic discusses the impact of VM priority on scheduling migrations and allocating bandwidth for migrations.

Move supports starting the migration of multiple migration plans simultaneously. To ensure that source providers are not overloaded and throttling errors are avoided, Move maintains rules internally to limit the number of disk migrations that can be performed simultaneously. When the number of disks being migrated from a particular source reaches this limit, any additional disks considered for migration are added to a queue. The queued disks are scheduled for migration after the ongoing migration of one or more disks is complete.

Move provides the option to set the priority of a VM while creating a migration plan. It provides three priority levels:

- **High**
- **Medium** (default priority level)
- **Low**

Once a migration plan is started, the priority of the VMs in that plan cannot be modified.

Scheduling

Move uses VM priority for scheduling. It schedules the disks of higher-priority VMs first for migration.

If additional disks are considered for migration when the number of disks being migrated has reached the threshold limit, Move performs the following steps:

1. Preempt the migration of the lowest-priority disk.
2. Schedule the migration of the highest-priority disk.

> **Note:** When the cutover of a VM is initiated, the disks of that VM are set to the highest-priority level. This ensures that the cutover is not blocked due to other ongoing migrations.

Example: If three VMs are scheduled for migration, with their priority set to **High**, **Medium**, and **Low** individually, Move first schedules the high-priority VM for migration, which is followed by the medium-priority VM and the low-priority VM in sequence.

Allocating Bandwidth

Move uses VM priority for allocating bandwidth. It allocates more of the available bandwidth to the disks of higher-priority VMs. If the throughput of higher-priority VMs consumes all the available bandwidth, the lower-priority VMs are queued. The queued VMs wait for bandwidth until the migration of higher-priority VMs is complete.

If the source for a migration plan is ESXi, Move allocates more data streams to the higher-priority VMs than to the lower-priority VMs. The data streams are allocated in the ratio of 4:3:2:1 for the highest-, high-, medium-, and low-priority VMs, respectively. This enables the higher-priority VMs to transfer data faster than the lower-priority VMs.

> **Note:** The highest-priority level is for VMs in the cutover phase or the test migration phase.

Example: Assume that three VMs are available for which the priority is set to **High**, **Medium**, and **Low** individually, and the data seeding for these VMs is under way with ESXi as the source. If 16 data streams are available, Move allocates them as follows:

- High-priority VM: 8 data streams
- Medium-priority VM: 5 data streams
- Low-priority VM: 3 data streams

This prioritization results in the higher-priority VMs getting a larger share of the available bandwidth than the lower-priority VMs. Once the data transfer for the high-priority VM is complete, Move reallocates the data streams as follows:

- Medium-priority VM: 11 data streams
- Low-priority VM: 5 data streams

MOVE APPLIANCE SETTINGS

The **Move Appliance Settings** page allows you to view and configure the settings of the Move appliance. It provides the following options:

- **Snapshot Configuration** provides the option to configure the interval at which snapshots are taken for any migration environment.
- **Docker Bridge IP** displays the Docker Bridge IP address.
- **NTP Servers** provides the option to configure NTP servers.

Snapshot Configuration

Move provides the option to configure the interval at which snapshots are taken for any migration environment. This topic details the procedure to configure the snapshot interval.

About this task

To configure the snapshot interval, do the following.

Procedure

1. Log on to Move.
2. Click the gear icon on the top-right and select **Appliance Settings** from the drop-down menu.
   Move Appliance Settings page appears.
3. From the **Settings** pane on the left, select **Snapshot Configuration**.
   Snapshot Configuration page opens. The table displays the duration of the snapshot interval for each of the source providers.
4. Edit the interval for the source providers as necessary.
   Under **Actions**, click the edit icon, edit the interval as necessary, then click the tick icon.

   Note: The maximum allowed snapshot interval for the source providers is as follows:

   - **VMware ESXi**: 72 hours (4320 minutes)
   - **Others**: 24 hours (1440 minutes)

   The table displays the updated values.
5. Ensure that the updates made to the **Duration** column are saved.
   a. Click **Back to Dashboard** on the top left.
   b. Open **Snapshot Configuration** settings page again by following Step 2 and Step 3.
   The table reflects the updates made to the **Duration** column.

Docker Bridge IP

This page displays the Docker Bridge IP address in Move.

About this task

The following procedure details the steps to view the Docker Bridge IP address in Move.

Procedure

1. Log on to Move.
2. Click the gear icon on the top-right and select **Appliance Settings** from the drop-down menu.
   Move Appliance Settings page appears.
3. From the **Settings** pane on the left, select **Docker Bridge IP**.
   Docker Bridge IP page opens and displays the Docker Bridge IP address.

Note:

- The Docker Bridge IP address must not be the same as any other IP address in the cluster.
- To edit the Docker Bridge IP address, see KB-7135.

NTP Servers

Move provides the option to configure NTP servers.

About this task

The following procedure details the steps to configure NTP servers for Move.

Procedure

1. Log on to Move.
2. Click the gear icon on the top-right and select **Appliance Settings** from the drop-down menu.
   Move Appliance Settings page appears.
3. From the **Settings** pane on the left, select **NTP Servers**.
   NTP Servers page opens. Move has pool.ntp.org configured as the default NTP server.
4. (Optional) To add another NTP server, enter the hostname or the IP address of that server in the **NTP Server** field and click **Add**.
   The new NTP server appears in the **Hostname or IP Address** table.

Note:

- Move does not allow adding an NTP server if the hostname or the IP address entered in the **NTP Server** field is invalid. Be sure to enter a valid hostname or IP address while adding an NTP server.
- To delete an NTP server, click the **x** icon of that server.
- Move requires at least one active NTP server. If only one NTP server is configured, it cannot be deleted.

MOVE TROUBLESHOOTING

This section guides you through the troubleshooting steps for some of the issues you might face while using Move.

Move Support Bundle Collection

You can generate and download a support bundle that you can send to Nutanix Support for assistance. Support bundle contains a collection of logs from the following sources:

- Nutanix-Move
- Nutanix-Move Hyper-V agent logs

> **Note:** From Move 1.1.3 release, the support bundle includes the ping statistics of the source ESXi hosts and target AHV hosts.

Downloading Support Bundle (UI)

You can generate and download a support bundle from the Move dashboard that you can send to Nutanix Support for assistance.

About this task

To download the support bundle from Move UI, do the following:

Procedure

1. Log on to the Move UI.
2. Click the **gear** icon from the Move dashboard, and then click **Download Support Bundle.**

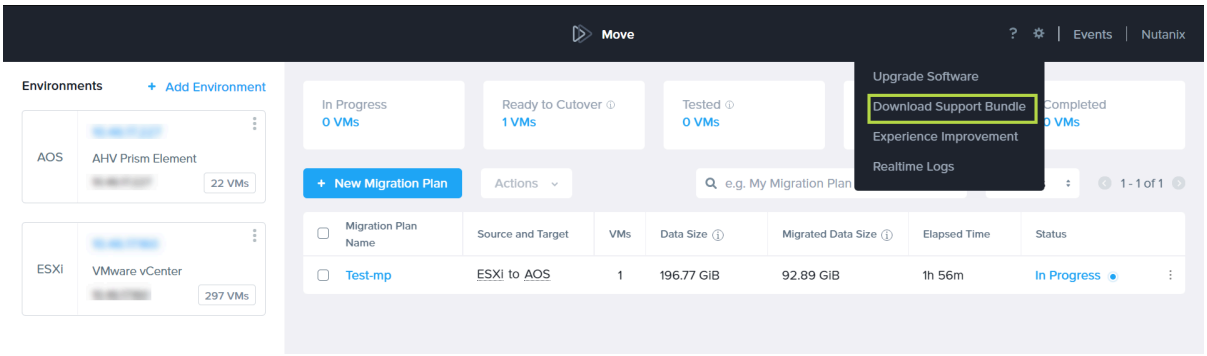

Figure 99: Support Bundle Download

> **Note:** Download might take a few moments to begin.

A diagnostic bundle is generated and downloaded as move-support-bundle.`<date-of-download>`.tar.gz. For example, move-support-bundle.Sep-27-2021-01_39_17.tar.gz

Downloading Support Bundle (CLI)

You can generate and download a support bundle from the Move CLI that you can send to Nutanix Support for assistance.

About this task

To download the support bundle from Move CLI, do the following:

Procedure

1. Log on as root user into the Move VM.
2. To download the support bundle, run the following command:

```
root@move on ~ $ support-bundle
```

> **Note:** Download might take a few moments to begin.

A diagnostic bundle is generated under the default directory or the dump path you specified in the support-bundle --dump-path `target-directory` script.

Checking Real Time Logs

Move provides an option to check the real time logs for all the Move components. If you face any issues with a specific VM or migration, you can track a log from this window.

About this task

To check the real time logs, do the following:

Procedure

1. Log on to the Move UI.

2. Click the **gear** icon from the Move dashboard, and then click **Realtime Logs**

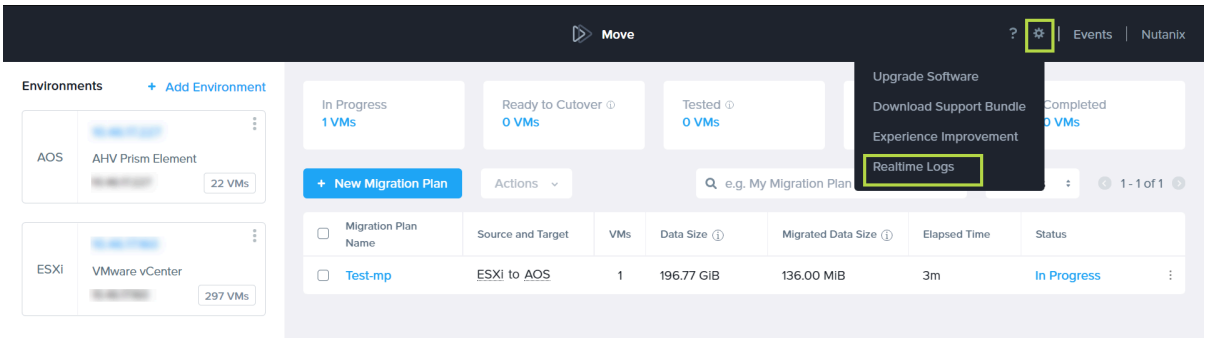

Figure 100: Real Time Logs Sub-menu

You can do the following:

- Filter the logs based on the component from the **Filter Pane**.
- Scroll up or select a log to pause the logs for better readability.
- Use the browser-level search to search specific logs.

A list of real time logs are displayed in a new window.

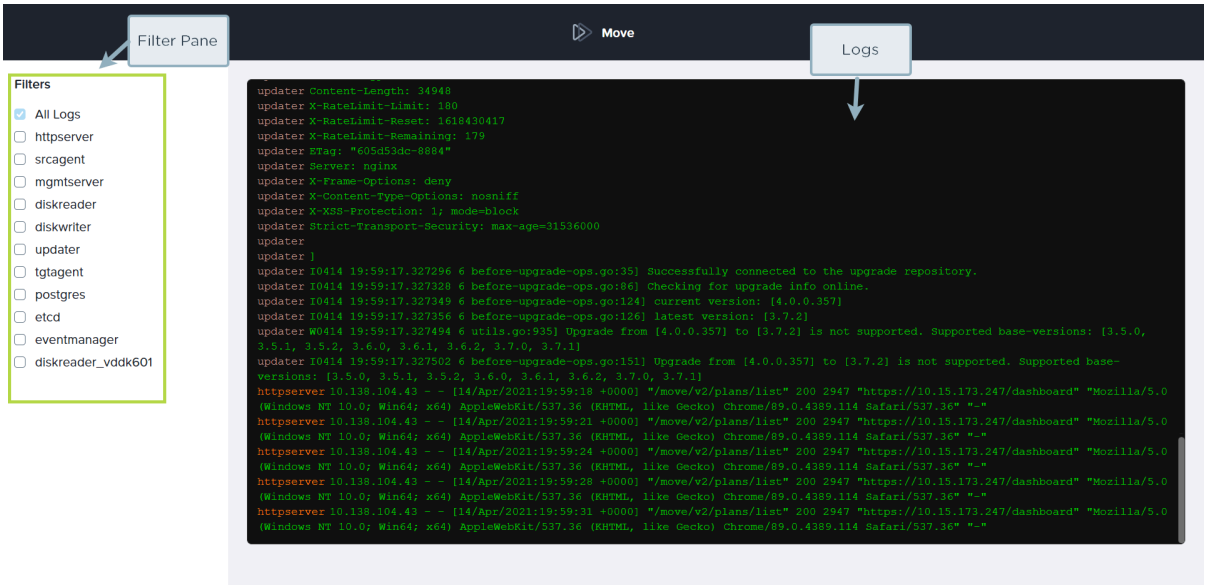

Figure 101: Real Time Logs

Error Adding Provider

This section guides you through the troubleshooting steps for some of the issues you might face while adding the provider in the Move UI.

- Verify that you are using correct IP address or FQDN.
- Verify that the provider credentials are correct.
- Verify that the provider is DNS resolvable to IP address or IP address can be pinged from the Move VM.

- Verify service account used has admin privileges.

Unable to Uninstall VMware Tools

This section lists the KB articles if you face any issues while uninstalling the VMware Tools from your VMs after migration from ESXi to AHV.

> **Note:** This section is only applicable if you cannot uninstall VMware Tools from the Windows VMs. This section is not applicable for the Linux VMs.

Refer to the following KB articles if you face any issues while uninstalling the VMware Tools.

- Nutanix KB article 10294
- VMware KB article 1001354

Manual Cleanup for VM Migrations

This section describes the steps one needs to perform for manual cleanup after VM migration actions such as Cancel or Discard fail and the UI displays the migration status as Cleanup Failed.

A VM migration can have different types of providers as Source and/or Target. The message in the info icon besides the Cleanup Failed error clearly states whether only the Source or Target cleanup has failed or if both have failed.

Move might have performed some of the cleanup steps during its failed attempt as part of the Discard or Cancel action. Proceed with the next steps in case a step has already been performed.

> **Note:** *Guest VM* refers to the VM being migrated.

The following sections list down the cleanup steps for different Source types and Target types:

AWS as a source

1. Remove the Guest VMs public IP address from the security group of the AWS agent of Move.
2. Delete all the snapshots taken by Move for the volumes of the Guest VMs. Snapshots are named as *Move-Snapxxxx*.

   Also search and delete the EBS volumes created from those snapshots by using snapshot-id.
3. Login to the Guest VMs and run the following cleanup scripts for Linux and Windows respectively:

   For Linux: /opt/Nutanix/Uninstall/scripts/cleanup_installation.sh

   For Windows: C:\Nutanix\Uninstall\scripts\cleanup_installation.ps1

AHV as a source

Delete all the snapshots taken by Move for the Guest VMs. Snapshots are named as *Move-Snapxxxx*.

> **Note:** This can be done using the V3 APIs of Prism.

Hyper-V as a source

Delete all the checkpoints taken by Move for the Guest VMs using the Hyper-V manager. The checkpoints are named as *Move-Snapxxxx*.

1. Login to the source VM.
2. Run the following cleanup scripts for Linux and Windows respectively:

   For Linux: /opt/Nutanix/Uninstall/scripts/cleanup_installation.sh

   For Windows: C:\Nutanix\Uninstall\scripts\cleanup_installation.ps1

ESXi as a source

1. Go to vCenter and delete all the snapshots taken by Move for the UVMs. Snapshots are named as *Move-Snapxxxx*.
2. Go to vCenter and disable CBT for the Guest VMs if it was originally disabled.
3. For ESXi Guest VM cleanup, download the following files from the HTTP server of Move and run it in the Guest VM.

   - For Linux Guest VM, download *cleanup_installation.sh*
   - For Windows Guest VM, download *cleanup_installation.ps1*

   For example, if `move_ipaddress` is the IP address of the Nutanix-Move VM, then download the following:

   - For Linux: https://`move_ipaddress`/resources/scripts/linux/cleanup_installation.sh
   - For Windows: http://`move_ipaddress`/resources/scripts/win/cleanup_installation.ps1

AOS as a target

1. On the source Guest VMs, uninstall the VirtIO driver if the target is AHV.
2. Delete the partially copied vDisks from the AOS containers which are mounted inside the Nutanix-Move VM to free up unwanted container space as shown in the following steps:

   > **Note:** If VM is successfully created on target (that is, VM migration is successful), Move automatically deletes these files.

   1. Perform a SSH login to the Move VM and enter the root shell with the rs command.
   2. If the target is AHV, then on the source Guest VM uninstall the VirtIO driver.
   3. Find the vDisks for the VM(s) for which the cleanup has failed.
      1. Change directory using cd to /opt/xtract-vm/datamover
      2. Run [find . -name "*vdiskname*" |xargs ls -ltr] for each vDisk and copy the full path and then run the rm command on that path to delete the file.

         Example: In the following example, a find with the VM name as the keyword is performed and both the vDisk for that VM came in the output since the vDisk name had the VM name in it.

         The UUID in bold is the migration plan UUID. Before deleting this, you can match it in the Move UI with the migration plan under which the VMs are present. On the Move UI, navigate inside the migration plan by clicking the plan name and in the address bar. The last token is the migration plan UUID (https://MOVE_IP/plan-details/d43008ed-72d8-43c5-a404-4f93916ac538).

         > **Note:** find . -name "*RHEL-63-64bit*" |xargs ls -ltr
         >
         > -rwxr-xr-x 1 diskwriter vmxtract 2147483648 May 15 2020 ./CVM_10.136.74.14/1dd6f8b3-c4c6-42e7-928e-b5d2a6aa27f9/**d43008ed-72d8-43c5-a404-4f93916ac538**/5554e915-16ca-5bcb-9a70-b0cfc693d0b9/2001_RHEL-63-64bit_1.vmdk.raw
         >
         > -rwxr-xr-x 1 diskwriter vmxtract 2147483648 May 15 2020 ./CVM_10.136.74.14/1dd6f8b3-c4c6-42e7-928e-b5d2a6aa27f9/**d43008ed-72d8-43c5-a404-4f93916ac538**/0a6d5b57-074d-5ba1-8c25-9a749a969c6c/2000_RHEL-63-32bit_1.vmdk.raw

      3. Once the files are deleted, the Curator full scan is run automatically to reclaim the free space.

         > **Note:** The space will be reclaimed after 8 hours. If you want to reclaim the space immediately, contact Nutanix support.

Testing Network Performance of Move

If you are facing any performance issues such as slow migration, you can run iPerf to troubleshoot network related issues.

Pre-requisites: Exposing the iPerf3 port for the Throughput Test

To expose the iPerf3 port, do the following:

1. Expose the iPerf port by adding the rule in /opt/xtract-vm/bin/docker-compose.yml.

   In the docker-compose.yml file, in the specific container section, under ports, add the rule to allow the port.

   For example, to run the iPerf3 from the diskwriter container, add the port mapping in the diskwriter - ports section of the docker-compose.yml file. Exposing port 9100 in the following example.

   ```
   ports:
         - "127.0.0.1:3020:3020"
         - "9100:9100"
   ```

2. Restart the Move service.

   ```
   > svcrestart
   ```

   Caution: Active migrations will fail on service restart.

Running Iperf in Server Mode from Nutanix-Move VM

To run iPerf in the server mode, do the following:

1. Connect to the container shell.

   For example, to enter diskwriter container, run the `diskwriter-shell` command from the root shell.
2. Run the iPerf3 command.

   ```
   > iperf3 -s -p port
   ```

   For example,

   ```
   > iperf3 -s -p 9100
          -----------------------------------------------------------
          Server listening on 9100
          -----------------------------------------------------------
   ```

3. Run the iPerf3 command in the client mode with the same port from the VM where we try to check the network throughput.

Running Iperf in Client Mode from Nutanix-Move VM

To run iPerf3 in the client mode, do the following:

1. Make sure the iPerf3 is running in server mode in the VM to which the network throughput will be measured.
2. Enter the Move container from which the network performance will be measured.

   For example, to test the performance from the diskwriter, enter diskwriter shell using command `diskwriter-shell`.
3. Run the iperf3 command in the client mode.

   ```
   > iperf3 -c iperf3-server-ip -p -port
   ```

   For example,

   ```
   > iperf3 -c 10.1.1.100 -p 9100
   ```

Debugging Stats

This section guides you through the steps for debugging the stats.

Move is composed of multiple internal services for reading, writing data, and coordination. Each service exposes stats through REST/HTTP to localhost. Stats are also available in the support bundle for offline analysis.

```
[admin@Nutanix-Move ~]$ service iptables stop
#Query stats from disk reader,[admin@Nutanix-Move ~]$ curl localhost:8099/debug/vars

# disk writer, [admin@Nutanix-Move ~]$ curl localhost:2000/debug/vars

# source agent [admin@Nutanix-Move ~]$ curl localhost:8085/debug/vars

#and target agent [admin@Nutanix-Move ~]$ curl localhost:8086/debug/vars
[admin@Nutanix-Move ~]$ service iptables start
```

Troubleshooting UI Issues

This section guides you through the troubleshooting steps for some of the issues you might encounter in the Move UI.

- Source vCenter VM list is empty when selecting VMs for migration.

  Troubleshooting step: Ensure to wait for a few minutes for Move VM to build inventory database and click the **Refresh** link.

- If you refresh (F5) is the browser, and if the GUI loads with no data.

  Troubleshooting step: In the web browser, type https://`ip-address` and logon again.

- When UI shows an error message without any details.

  Troubleshooting step: Download the Support Bundle logs containing extra error information.

Licensing Window Pops-Up

After the relocation of a VM from one physical host to another, the Microsoft Licensing mechanism checks for the server hardware or Hypervisor change and requires the Microsoft Windows license to be reactivated. This is a general issue and not specifically limited to Move because the Microsoft Windows operating system is not hardware-agnostic by design. So, there are no ways to preserve license or reactivate automatically using APIs or any other mechanisms. Workaround: Reactivate Microsoft Windows license.

Workaround to fix this issue is to reactivate Microsoft Windows license.

Missing Static IP Address Post Migration

If the static IP address is not assigned or the static IP address is lost after migration, you can check the logs.

After the VM migration, previous assigned static IP addresses are lost. In Windows, you can collect the log files at C:\Nutanix\Temp. In Linux, you can collect the log files at /opt/Nutanix/logs.

> **Note:** The directory which stores the log files is available only if:
>
> - you selected **Automatic Preparation**, or
> - you run the generated *VM preparation script* after selecting **Manual Preparation**.

Contact Nutanix Support at https://portal.nutanix.com/ and send the log files.

Setting Up Multihomed Environment

Move does not support multihomed environment setup, however, you can perform the following workaround.

About this task

To setup a multihomed environment, do the following workaround:

Procedure

1. Create a new migration plan with the target container, and save it.
2. Find the UUID of the chosen target storage container.
3. Allowlist the Move IP address in the target container from Prism Element.

   ```
   <move-ip/255.255.255.255>
   ```

4. SSH to the Move VM, and type `tgtagent-shell`.
5. Create a directory.

   ```
   mkdir /opt/xtract-vm/datamover/CVM_<cvm_ip>/<tgt_storage_container_uuid>
   ```

6. Mount the target container.

   ```
   sudo mount -t nfs -o soft <CVM_IP>:/<tgt_contr_name> /opt/xtract-vm/datamover/
   CVM_<cvm_ip>/<tgt_storage_container_uuid>
   ```

7. Exit from **tgtagent-shell** and start the migration.

VMs Reaching Maintenance Mode (AWS to AHV and AHV to AWS)

For AWS to AHV and AHV to AWS migrations, you might encounter incorrect login loop issue on the migrated VM for Linux instances.

About this task

If the EC2 instance reaches the maintenance mode (incorrect login loop) in the source (AWS to AHV), restart the EC2 instance once or twice till the EC2 instance starts in the default mode.

> **Note:** Sometimes, you might require to restart the EC2 instance 4-5 times as well.

If the VM reaches the maintenance mode (incorrect login loop) in the target (AHV or AWS), do the following:

Procedure

1. To launch the VM Console in default mode, press the **Ctrl+D** key.
2. Press the **Enter** key.
3. To login to the VM, enter the user name and password.
4. Type dracut -f.
5. Restart the VM

FreeBSD VMs are Not Starting on Target

FreeBSD (10.x versions) VMs are not starting on target after migration.

About this task

To start the VM, do the following workaround:

Procedure

1. Shut down the VM.
2. Set the clock to False to disable it.

   ```
   acli vm.update <vm_uuid> extra_flags=enable_hyperv_clock=0
   ```

3. Start the VM.

Bringing Disks Online Post Migration (Windows)

You can bring the Windows disks online after the VM migration is complete. After the migration of VM disks, they appear as new disks. Move brings the disks online automatically. However, if you chose to bypass the guest operations, then perform the following tasks to bring the disks online after migration.

About this task

To bring the Windows disks online after migration, do the following:

Procedure

1. Log on to the Windows guest operating system in the VM.
2. Open a command window or PowerShell console, and then log on as an admin.
3. Change the SAN policy to bring all the disks online.

   ```
   >> echo san policy=OnlineAll | diskpart
   ```

   All the Windows disks become online.

COPYRIGHT