



Use the VMware migration advisor

VMware workloads

NetApp
August 04, 2025

Table of Contents

Use the VMware migration advisor	1
Migrate to Amazon Elastic VMware Service	1
Explore savings for Amazon Elastic VMware Service with BlueXP workload factory	1
Create a deployment plan for Amazon EVS using the VMware workloads migration advisor	1
Deploy the recommended FSx for ONTAP file system	2
Migrate to Amazon EC2	4
Capture your VM configurations before migrating to Amazon EC2 with VMware workloads	4
Create a deployment plan for Amazon EC2 using the VMware workloads migration advisor	8
Deploy the recommended FSx for ONTAP file system	14
Migrate to VMware Cloud on AWS	15
Capture your VM configurations before migrating to VMware Cloud with VMware workloads	15
Create a deployment plan for VMware Cloud on AWS using the VMware workloads migration advisor	17
Deploy the recommended FSx for ONTAP file system	21
Connect your FSx for ONTAP file systems to VMware Cloud on AWS	23
Migrate your data to the new infrastructure with BlueXP workload factory for VMware	23

Use the VMware migration advisor

Migrate to Amazon Elastic VMware Service

Explore savings for Amazon Elastic VMware Service with BlueXP workload factory

Explore potential savings for migrating your VMware workloads to Amazon Elastic VMware Service (EVS). The calculator enables you to compare costs of using Amazon EVS with and without Amazon FSx for NetApp ONTAP as storage.

If the savings calculator determines that the most cost-effective storage is FSx for ONTAP, you can create a detailed assessment, which provides a migration plan that you can review before using. You can then use the Codebox to generate Infrastructure-as-Code templates.

Explore savings for EVS environments

Follow these steps to explore potential savings for a planned migration to an Amazon EVS environment.



Before proceeding, review the disclaimer at the bottom of the savings calculator to understand more about how the pricing estimates are calculated.

Steps

1. Log in to workload factory using one of the [console experiences](#).
2. From the VMware workloads tile, select **Explore savings** and then **Amazon Elastic VMware Service (EVS)**. The savings calculator is displayed.

The **Environment preferences** pane provides sliders that you can adjust depending on your environment's needs.

3. Adjust the following sliders as needed to see real-time savings calculations for the values that you choose. If using a keyboard, you can make small adjustments using the arrow keys:
 - a. **Required physical CPU (count)**
 - b. **Required physical memory (GiB)**
 - c. **Required VM storage (TiB)**
 - d. **EVS billing plan**
4. Do one of the following:
 - To use the migration advisor to [create a deployment plan](#) for an Amazon EVS environment, select **Create a detailed assessment**.
 - Select **Close** to close the TCO calculator.

Create a deployment plan for Amazon EVS using the VMware workloads migration advisor

Log in to the NetApp workload factory to access the VMware migration advisor. You'll follow the steps in the wizard to create a deployment plan or migration plan that is customized for your needs.

Note that you must have a user name and password to access workload factory. If you don't have access, create an account now. See the instructions [here](#).

Create a deployment plan by manually specifying environment requirements

You can manually specify your requirements for virtual machines in Amazon Elastic VMware Service and use customized Amazon FSx for NetApp ONTAP file systems as external datastores.

Requirements

- You must have an Amazon Elastic VMware Service subscription.

Steps

1. Log in to workload factory using one of the [console experiences](#).
2. Optionally, if you accessed workload factory from the Workload Factory console, from the VMware workloads tile, select **Plan and create** and then select **Amazon Elastic VMware service**. The VMware migration advisor splash screen is displayed.
3. Select the **Amazon Elastic VMware service** tile.
4. When ready, select **Let's start**.
5. Enter the requirements for your VM environment.

Keep in mind the following VM migration performance requirements and recommendations:



- A minimum storage capacity of 10TiB is recommended for the VM inventory for performance reasons.
- A minimum throughput is required for the Amazon Elastic VMware Service datastores, depending on the required number of IOPS you specify for this deployment.
- Depending on the FSx for ONTAP filesystem configuration, a minimum number of external datastores is required for the Amazon Elastic VMware Service environment to reach optimal performance.

6. When ready, select **Next** to review the migration plan.
7. Review the plan. Expand each section to view the plan details.
8. Optionally, when you are satisfied with the migration plan, you can save it as a template or a PDF:
 - Select **Export plan** to save the migration plan as a template in a `.json` format on your computer. You can import the plan at a later time to use as a template when deploying systems with similar requirements.
 - Select **Download plan** to download the deployment plan in a `.pdf` format so you can distribute the plan for review.
9. [Deploy the VMware workloads recommended FSx for ONTAP file system](#).

Deploy the recommended FSx for ONTAP file system

After you verify that the recommended FSx for ONTAP file system (or multiple file systems in some cases) meets your exact requirements, you can use workload factory to deploy the system in your AWS environment.

Depending on the policy and permissions that you added to your workload factory account, you can deploy the FSx for ONTAP file system completely using workload factory (using read/write mode). If you have fewer

permissions (read-only mode), or no permissions (Basic mode), you'll need to use the CloudFormation template from the Codebox and deploy the FSx for ONTAP file system yourself in AWS.

Requirements for deployments to Amazon Elastic VMware Service

- You must have an Amazon Elastic VMware Service subscription.
- You must have the necessary permissions to create an FSx for ONTAP file system in your AWS account.

Steps

1. At the bottom of the **Review plan** page, select **Deploy** and the Create an FSx for ONTAP file system page is displayed.

Most of the fields that define your FSx for ONTAP file system are completed based on the information you provided, but there are a few fields that you need to complete in this page.

You can use the Quick create or Advanced create option. Advanced create offers a few additional storage parameters that you can customize. [See what these two options offer.](#)

2. In the "File system general configuration" section, enter the following information:
 - AWS credentials:** Select or add credentials that will give workload factory the permissions necessary to create your FSx for ONTAP file system directly. You can also select the CloudFormation code from Codebox and deploy the FSx for ONTAP file system yourself in AWS.
 - File system name:** Enter the name that you want to use for this FSx for ONTAP file system.
 - Tags:** Optionally, you can add tags to categorize this FSx for ONTAP file system.
3. In the "Network & security" section, enter the following information:
 - Region & VPC:** Select the region and the VPC where the FSx for ONTAP file system will be deployed.
 - Security group (Advanced create only):** When using the **Advanced create** option, you can select the default security group for the FSx for ONTAP VPC so that all traffic can access the FSx for ONTAP file system. You can either create a new security group or select an existing one.

If you enable the **Adjust security group configuration to EVS NFS datastores** option, workload factory configures the security group according to the settings for EVS NFS datastores.

You can add an inbound rule to the security group that restricts what other AWS services can access the FSx for ONTAP file system. This will reduce the number of services that are open. These are the minimum ports and protocols:

Protocols	Ports	Purpose
TCP, UDP	111	Portmapper (used to negotiate which ports are used in NFS requests)
TCP, UDP	635	NFS mountd (receives NFS mount requests)
TCP, UDP	2049	NFS network traffic
TCP, UDP	4045	Network Lock Manager (NLM, lockd) - Handles lock requests.
TCP, UDP	4046	Network Status Monitor (NSM, statd) - Notifies NFS clients about reboots of the server for lock management.

- c. **Availability zone:** Select the availability zone and the subnet.

You should select the same availability zone as where your VMware SDDC is deployed if you want to avoid charges for cross-AZ traffic.

- d. **Encryption** (Advanced create only): When using the **Advanced create** option, you can select the AWS encryption key name from the dropdown.
- e. **NFS Datastores access** (Advanced create only): When using the **Advanced create** option, you can select whether all hosts can access the datastores or whether only the EVS management host can access the datastores.

4. In the "File system details" section, enter the following information:

- a. **ONTAP credentials**: Enter the ONTAP user name password.
- b. **Storage VM credentials** (Advanced create only): Enter and confirm the storage VM password. The password can be specific to this file system, or you can use the same password entered for ONTAP credentials.

5. In the "EVS Cluster attachment" section, enter the following information:

- a. **Datastore mount options**: Optionally, enable the **Mount datastores to EVS cluster** option to automatically connect the datastores to the Amazon EVS cluster. This option also causes workload factory to configure VMware ESXi host settings so that they match ONTAP best practice recommendations. Before you deploy the file system, you can review the plan details in the **Summary** section to see the settings that were changed.
- b. **EVS Cluster vSphere console details**: Enter the IP address or FQDN of the VMware vCenter server that should connect to Amazon EVS.
- c. **vSphere administrator credentials secret ARN**: Choose the secret ARN for the vSphere administrator credentials. These credentials are used to mount datastores and configure recommended VMware settings.

If Amazon EVS is not available in this region, or if your account doesn't have permissions to retrieve the list of available secret ARNs, you can manually enter a secret ARN value.

6. In the **Summary** section, you can view the FSx for ONTAP file system and datastore configuration that the VMware migration advisor has designed based on your information.

7. Select **Create** to deploy the FSx for ONTAP file system. This process can take up to 2 hours.

Optionally, in the Codebox window you can select **Redirect to CloudFormation** to create the file system and recommended VM configuration using a CloudFormation stack.

In either case, you can monitor the creation progress in CloudFormation.

Result

The FSx for ONTAP file system is deployed. You can now use the AWS CloudFormation template in the Codebox to deploy the recommended VM configuration to your Amazon Elastic VMware Service environment.

Migrate to Amazon EC2

Capture your VM configurations before migrating to Amazon EC2 with VMware workloads

You can use the Migration advisor VM data collector script, RVTools, or NetApp Data Infrastructure Insights to capture your current VM configurations. We recommend that you

use the Migration advisor VM collector script or Data Infrastructure Insights because these options gather both VM configuration and performance data for the most accurate plan.

Use RVTools (quick assessment)

RVTools is a Windows application that interacts with VCenter and ESX server (5.x to 8.0) to capture information about your VMware virtual environments. It gathers information about VMs, CPU, memory, disks, clusters, ESX hosts, datastores, and more. You can export this information to an xlsx file to use with the migration advisor.

[Learn more about RVTools](#)

Steps

1. Download and install RVTools 4.4.2 or greater from [RVTools website](#)
2. Launch RVTools and log in to the target VCenter Server.

This will capture information about that VCenter Server.

3. Export your VMware information to an xlsx or xls file.

See the chapter on "commandline parameters" in the [RVTools documentation](#) for details.

Result

The xlsx or xls file that contains your results is saved to your working directory.

Use the Migration advisor VM collector script (detailed planning)

The NetApp Migration advisor VM collector script gathers VM configuration information from your vCenter similar to what RVTools collects, however, it captures current VM performance data to provide actual read and write IOPS and throughput statistics. You can use the script in one of two modes:

- Hourly data collection
- Daily data collection

You need to start the Migration advisor to copy the Migration advisor VM collector script from the Codebox window.

Before you begin

The system where you'll be running the collector script must meet the following requirements:

- Microsoft PowerShell 7.0 or greater must be installed. See the [Microsoft PowerShell documentation](#) if you need to install PowerShell.
- VMware PowerCLI must be installed. See the [VMware PowerCLI documentation](#) if you need to install PowerCLI.
- SSL certificate checking must be disabled.
- PowerShell must allow running unsigned scripts.

Note that if you plan to capture long-term statistics from your VMs (daily statistics), you must have activated statistics level 3 or above in the vSphere management console (VMware management environment).

Steps

1. Log in to workload factory using one of the [console experiences](#).
2. From the VMware workloads tile, select **Assess and plan** and then select **Migration to AWS native compute**. The VMware migration advisor splash screen is displayed.

Review the information to learn about the benefits of the migration advisor.

3. When ready, select **Let's start**.
4. Select the **Detailed planning** tile, and select **Next**.

The Codebox window appears on the right side of the page.

5. In the Codebox window, select either  to save the Migration advisor data collector script (named "list-vms.ps1") to the target system, or select  to copy the script so you can paste it on the appropriate system.
6. Follow these steps to capture the configurations of your current VMs:

- a. Log in to the system where you downloaded the data collector, on which PowerShell and PowerCLI are installed.
- b. Connect to your VMware vCenter server by running the following command:

```
Connect-VIServer -server <server_IP>
```

Replace <server_IP> with the IP address or hostname of your VMware server.

- c. Run the data collector script you downloaded and specify the data collection option for "hourly" or "daily" VM statistics.

```
./list-vms.ps1 -PerformanceStatisticsTimePeriod <LastHour | LastDay>
```

where:

- **LastHour** collects IOPS and throughput data for the past 1 hour with 20-second intervals.
- **LastDay** collects IOPS and throughput data for the past 24 hours with 5-minute intervals.

Optionally, you can run the data collector script without specifying any options; this enables you to select the data collection period from the interactive menu.

Result

The script outputs a compressed file with a timestamp in the name in the current working directory. The .zip file contains a list of all VMs and their properties, along with their IOPS and throughput data.

Use NetApp Data Infrastructure Insights (detailed planning)

NetApp Data Infrastructure Insights is a cloud infrastructure monitoring tool that gives you visibility into your complete infrastructure. You need to enable temporary API communication between the migration advisor and Data Infrastructure Insights so that the migration advisor can gather information about your VMware infrastructure.

[Learn more about Data Infrastructure Insights](#)

Steps

1. Log in to your Data Infrastructure Insights account.

2. In the side navigation menu, select **Admin**.
3. In the resulting menu, select **API Access**.
4. In the **API Access Tokens** tab, select **+ API Access Token**.

The Create an API Access Token dialog appears.

5. Provide a name and description for the API token.
6. Under **What type of APIs will this token be used to call** select one of the following:
 - Acquisition Unit
 - Assets
 - Data Collection
7. Under **Permissions** select **Read Only**.
8. Under **Token expires in**, choose the length of time that you need the API token to be valid.
9. Uncheck **Automatically rotate tokens for Kubernetes**.
10. Select **Save**.
11. Select **Copy API Access Token**.
12. Save this token in preparation to use with the workload factory migration advisor.

What's next?

[Create an Amazon EC2 deployment plan using the migration advisor.](#)

Create a deployment plan for Amazon EC2 using the VMware workloads migration advisor

Log in to the NetApp workload factory to access the VMware migration advisor. You'll follow the steps in the wizard to create a deployment plan or migration plan that is customized for your needs.

Note that you must have a user name and password to access workload factory. If you don't have access, create an account now. See the instructions [here](#).

Create a deployment plan based on an on-premises vSphere environment

You can migrate your current virtual machine configurations in on-premises vSphere environments to virtual machines in Amazon EC2 and use customized Amazon FSx for NetApp ONTAP file systems as external datastores.

Requirements

- You must have created the inventory file from your existing systems by using the migration advisor VM collector (.zip file) or by using RVTools (.xlsx file).
- You must have access to the inventory file from the system on which you are logging in to workload factory.

Steps

1. Log in to workload factory using one of the [console experiences](#).
2. From the VMware workloads tile, select **Assess and plan** and then select **Migration to AWS native compute**. The VMware migration advisor splash screen is displayed.

Review the information to learn about the benefits of the migration advisor.

3. When ready, select **Let's start**.
4. Choose the level of detail and method you prefer during the process of creating a deployment plan:
 - **Quick assessment:** This option uses an RVtools inventory file to retrieve information about your VMs during an accelerated process, resulting in a design blueprint for your cloud migration.
 - **Detailed planning:** These options use either the migration advisor data collector inventory file or direct API access to NetApp Data Infrastructure Insights to collect detailed information, resulting in a deployment-ready design.

RVTools (quick assessment)

Steps

5. Select the **RVTools** tile and select **Next**.

The **Prepare for AWS Cloud onboarding** page is displayed.

6. In the *VM configuration upload* section, select  and select the .xlsx file created by RVTools.

The **Data collection details** section provides the date range and number of hours of data that was collected using the Onboarding Advisor data collector.

The **VM inventory summary** section is populated from the inventory file to reflect the number of VMs and the total storage capacity.

7. In the *VM inventory considerations* section, select the options to filter the list of VMs that you want to migrate.

- a. **Region:** Select the region where Amazon FSx for NetApp ONTAP file systems will be deployed.

For optimal performance and cost efficiency, this is typically the same region as where your existing Amazon EC2 SDDC is deployed.

- b. Choose a predicted performance level for the VMs in this region. We recommend that you begin with a smaller IOPS setting. You can increase your provisioned SSD IOPS after the file system is created as workloads are migrated or deployed:

- **Standard to high performance:** For VMs with average IO rates between 20 and 5000 IOPS.
- **Very high performance:** For VMs with average IO rates of greater than 5000 IOPS.
- **Very low performance:** For VMs with average IO rates lower than 20 IOPS.

8. In the *Target capacity and protection considerations* section, select from a few storage options.

- a. **VM Storage to consider:** Select whether the datastores created for each onboarded VM are sized based on their currently utilized size (recommended) or their provisioned size.

The external datastores will be implemented using Amazon FSx for NetApp ONTAP file system volumes.

- b. **Average data reduction ratio:** Choose from among the three common data reduction ratios. Select "1:1 - No reduction", "1:1.25 - 20% reduction", or "1:1.5 - 33% reduction".

Select **Help me decide** if you're unsure which ratio to choose. The *Data reduction ratio assistant* dialog appears. Select any statements that apply to your VM inventory and storage estate. The assistant will recommend an appropriate data reduction rate. Select **Apply** to use the recommended ratio.

- c. **Headroom percentage:** Enter the percentage of capacity growth that is added to the capacity for your FSx for ONTAP file systems.

Note that if you select an amount less than 20% you won't be able to create volume snapshots for protection and long-term backups.

- d. **VM snapshot protection:** Enable this option to protect the VMs with snapshots.

9. Select **Next**.

10. On the **Scope** page, select the VMs from the list that you want to include in the AWS migration. You can filter the list by the power state of each VM, and which data center and cluster the VM resides in.

In the VM list, you can select which types of VM information to display as columns. For example, selecting *Estimated instance type* displays a column with the estimated Amazon EC2 instance type for each VM, based on the VM resource requirements.

11. Select **Next**.
12. On the **Classify** page, review the VM information, volume classification rules, volume assignments, and list of volumes that will be migrated as part of deployment, and then select **Next**.



If a volume has capacity or performance requirements that exceed the capabilities of an Amazon FSx for NetApp ONTAP filesystem in a specific region, migration advisor recommends that the volume be deployed on an Amazon EBS filesystem.

13. On the **Package** page, review the EC2 instances and the volumes assigned to different FSx for ONTAP clusters, and then select **Next**.
14. On the **Review plan** page, review the estimated monthly savings and cost estimates for all the VMs that you plan to migrate.

The top of the page estimates the monthly savings for FSx for ONTAP file systems and EBS volumes. You can expand each section to view details for the suggested filesystem configuration, estimated savings breakdown, assumptions, and technical disclaimers.

Migration advisor data collector (detailed planning)

Steps

5. Select the **Migration advisor data collector** tile and select **Next**.

The **Prepare for AWS Cloud onboarding** page is displayed.

6. In the *VM configuration upload* section, select and select the .zip file created by the migration advisor data collector.

The **Data collection details** section provides the date range and number of hours of data that was collected using the Onboarding Advisor data collector.

The **VM inventory summary** section is populated from the inventory file to reflect the number of VMs and the total storage capacity.

7. In the *VM inventory considerations* section, select the region where Amazon FSx for NetApp ONTAP file systems will be deployed. For optimal performance and cost efficiency, this is typically the same region as where your existing Amazon EC2 SDDC is deployed.
8. In the *Target capacity and protection considerations* section, select from a few storage options.

- a. **VM Storage to consider:** Select whether the datastores created for each onboarded VM are sized based on their currently utilized size (recommended) or their provisioned size.

The external datastores will be implemented using Amazon FSx for NetApp ONTAP file system volumes.

- b. **Average data reduction ratio:** Choose from among the three common data reduction ratios. Select "1:1 - No reduction", "1:1.25 - 20% reduction", or "1:1.5 - 33% reduction".

Select **Help me decide** if you're unsure which ratio to choose. The *Data reduction ratio assistant* dialog appears. Select any statements that apply to your VM inventory and storage estate. The assistant will recommend an appropriate data reduction rate. Select **Apply** to use the

recommended ratio.

- c. **Headroom percentage:** Enter the percentage of capacity growth that is added to the capacity for your FSx for ONTAP file systems.

Note that if you select an amount less than 20% you won't be able to create volume snapshots for protection and long-term backups.

- d. **VM snapshot protection:** Enable this option to protect the VMs with snapshots.

9. Select **Next**.

10. On the **Scope** page, select the VMs from the list that you want to include in the AWS migration. You can filter the list by the power state of each VM, and which data center and cluster the VM resides in.

In the VM list, you can select which types of VM information to display as columns. For example, selecting *Estimated instance type* displays a column with the estimated Amazon EC2 instance type for each VM, based on the VM resource requirements.

11. Select **Next**.

12. On the **Classify** page, review the VM information, volume classification rules, volume assignments, and list of volumes that will be migrated as part of deployment, and then select **Next**.



If a volume has capacity or performance requirements that exceed the capabilities of an Amazon FSx for NetApp ONTAP filesystem in a specific region, migration advisor recommends that the volume be deployed on an Amazon EBS filesystem.

13. On the **Package** page, review the EC2 instances and the volumes assigned to different FSx for ONTAP clusters, and then select **Next**.

14. On the **Review plan** page, review the estimated monthly savings and cost estimates for all the VMs that you plan to migrate.

The top of the page estimates the monthly savings for FSx for ONTAP file systems and EBS volumes. You can expand each section to view details for the suggested filesystem configuration, estimated savings breakdown, assumptions, and technical disclaimers.

NetApp Data Infrastructure Insights (detailed planning)

Steps

5. Select the **NetApp Data Infrastructure Insights** tile and select **Next**.

The **Prepare for AWS Cloud onboarding** page is displayed.

6. In the *NetApp Data Infrastructure Insights connection configuration* section, enter the tenant endpoint for Data Infrastructure Insights.

This should be the URL you use to access Data Infrastructure Insights.

7. If you have not yet generated an API access token for Data Infrastructure Insights, follow the instructions on the page to create one. Otherwise, paste your API access token in the **Enter API Access Token** text box.

8. Select **Connect**.

Workload factory gathers information from Data Infrastructure Insights.

The **VM inventory summary** section is populated from the gathered information to reflect the number of VMs and the total storage capacity.

9. In the *VM inventory considerations* section, select the region where Amazon FSx for NetApp ONTAP file systems will be deployed. For optimal performance and cost efficiency, this is typically the same region as where your existing Amazon EC2 SDDC is deployed.

10. In the *Target capacity and protection considerations* section, select from a few storage options.

- VM Storage to consider:** Select whether the datastores created for each onboarded VM are sized based on their currently utilized size (recommended) or their provisioned size.

The external datastores will be implemented using Amazon FSx for NetApp ONTAP file system volumes.

- Average data reduction ratio:** Choose from among the three common data reduction ratios. Select "1:1 - No reduction", "1:1.25 - 20% reduction", or "1:1.5 - 33% reduction".

Select **Help me decide** if you're unsure which ratio to choose. The *Data reduction ratio assistant* dialog appears. Select any statements that apply to your VM inventory and storage estate. The assistant will recommend an appropriate data reduction rate. Select **Apply** to use the recommended ratio.

- Headroom percentage:** Enter the percentage of capacity growth that is added to the capacity for your FSx for ONTAP file systems.

Note that if you select an amount less than 20% you won't be able to create volume snapshots for protection and long-term backups.

- VM snapshot protection:** Enable this option to protect the VMs with snapshots.

11. Select **Next**.

12. On the **Scope** page, select the VMs from the list that you want to include in the AWS migration. You can filter the list by the power state of each VM, and which data center and cluster the VM resides in.

In the VM list, you can select which types of VM information to display as columns. For example, selecting *Estimated instance type* displays a column with the estimated Amazon EC2 instance type for each VM, based on the VM resource requirements.

13. Select **Next**.

14. On the **Classify** page, review the VM information, volume classification rules, volume assignments, and list of volumes that will be migrated as part of deployment, and then select **Next**.



If a volume has capacity or performance requirements that exceed the capabilities of an Amazon FSx for NetApp ONTAP filesystem in a specific region, migration advisor recommends that the volume be deployed on an Amazon EBS filesystem.

15. On the **Package** page, review the EC2 instances and the volumes assigned to different FSx for ONTAP clusters, and then select **Next**.

16. On the **Review plan** page, review the estimated monthly savings and cost estimates for all the VMs that you plan to migrate.

The top of the page estimates the monthly savings for FSx for ONTAP file systems and EBS volumes. You can expand each section to view details for the suggested filesystem configuration, estimated savings breakdown, assumptions, and technical disclaimers.

When you are satisfied with the migration plan, you have a few options:

- Select **Manage plan > Save a plan** to save the deployment plan data to your account, enabling you to import the plan at a later time to use as a template when deploying systems with similar requirements. You can name the plan before you save it (the username and timestamp are added to the name you provide).
- Select **Manage plan > Export a plan** to save the migration plan as a template in a .json format on your computer. You can import the plan at a later time to use as a template when deploying systems with similar requirements.
- Select **Manage plan > Download a report** to download the deployment plan in a .pdf format so you can distribute the plan for review.
- Select **Manage plan > Download instance storage deployment** to download the external datastore deployment plan in a .csv format so you can use it to create your new cloud-based intelligent data infrastructure.

You can select **Done** to return to the VMware migration advisors page.

Create a deployment plan based on an existing plan

If you are planning a new deployment that is similar to an existing deployment plan that you've used in the past, you can import that plan, make changes, and then save it as a new deployment plan.

Requirements

You must have access to the .json file for the existing deployment plan from the system on which you are logging in to workload factory.

Steps

1. Log in to workload factory using one of the [console experiences](#).
2. From the VMware workloads tile, select **Assess and plan** and then select **Migration to AWS native compute**.
3. Select **Import plan**.
4. Do one of the following:
 - Select **Load saved plan**.
 - a. From the list, select the plan you want to import.
 - b. Select **Load**.
 - Select **From my computer**.
 - a. Select the existing .json plan file that you want to import in the migration advisor, and then select **Open**.

The **Review plan** page is displayed.

5. You can select **Previous** to access previous pages and modify the settings for the plan as described in the preceding section.
6. After you have customized the plan to your requirements, you can save the plan or download the plan report as a PDF file.

Deploy the recommended FSx for ONTAP file system

After you verify that the recommended FSx for ONTAP file system (or multiple file

systems in some cases) meets your exact requirements, you can use workload factory to deploy the system in your AWS environment.

Follow the instructions in [Create an FSx for ONTAP file system in BlueXP workload factory](#) to deploy the FSx for ONTAP file system according to the configuration recommended by workload factory. You can use the **Advanced create** steps in the instructions to fully customize the file system.

Migrate to VMware Cloud on AWS

Capture your VM configurations before migrating to VMware Cloud with VMware workloads

You can use the Migration advisor VM collector script or RVTools to capture your current VM configurations. We recommend that you use the Migration advisor VM collector script because it gathers both VM configuration and performance data for the most accurate plan.

If you plan to deploy a new system based on an existing deployment plan that you previously saved locally, you can skip this step and select the existing deployment plan. [Learn how to use an existing plan to build a new plan.](#)

Use the Migration advisor VM collector script

The NetApp Migration advisor VM collector script gathers VM configuration information from your vCenter similar to what RVTools collects, however, it captures current VM performance data to provide actual read and write IOPS and throughput statistics.

You need to start the Migration advisor to copy the Migration advisor VM collector script from the Codebox window.

Before you begin

The system where you'll be running the collector script must meet the following requirements:

- Microsoft PowerShell 7.0 or greater must be installed. See the [Microsoft PowerShell documentation](#) if you need to install PowerShell.
- VMware PowerCLI must be installed. See the [VMware PowerCLI documentation](#) if you need to install PowerCLI.
- SSL certificate checking must be disabled.
- PowerShell must allow running unsigned scripts.

Note that if you plan to capture long-term statistics from your VMs (daily statistics), you must have activated statistics level 3 or above in the vSphere management console (VMware management environment).

Steps

1. Log in to workload factory using one of the [console experiences](#).
2. From the VMware workloads tile, select **Assess and plan** and then select **Migration to VMware Cloud on AWS**.

The VMware migration advisor main page is displayed.

3. Select **Create a new deployment plan**, select the option for **Use the migration advisor VMware data collector**, and select **Next**.

The **Prepare for VMware Cloud on AWS migration** page is displayed.

4. In the Codebox window, select either  to save the Migration advisor VM collector script (named "list-vms.ps1") to the target system, or select  to copy the script so you can paste it on the appropriate system.
5. Follow these steps to capture the configurations of your current VMs:
 - a. Log in to the system where you downloaded the data collector, on which PowerShell and PowerCLI are installed.
 - b. Connect to your VMware vCenter server by running the following command:

```
Connect-VIServer -server <server_IP>
```

Replace <server_IP> with the IP address or hostname of your VMware server.

- c. Run the data collector script you downloaded and specify the data collection option for "daily" or "hourly" VM statistics.

```
./list-vms.ps1 -isLongTermDataCollectionEnabled <true | false>
```

where:

- **true** collects IOPS and throughput data for the past 24 hours with 5-minute intervals
- **false** collects IOPS and throughput data for the past 1 hour with 20-second intervals.

Result

The script outputs a CSV file named `list-vms-yyyy-MM-dd-HH-mm-ss.csv` in the current working directory. The CSV file contains a list of all VMs and their properties, along with their IOPS and throughput data.

Use RVTools software

RVTools is a Windows application that interacts with VCenter and ESX server (5.x to 8.0) to capture information about your VMware virtual environments. It gathers information about VMs, CPU, memory, disks, clusters, ESX hosts, datastores, and more. You can export this information to an `xlsx` file to use with the migration advisor.

[Learn more about RVTools](#)

Steps

1. Download and install RVTools 4.4.2 or greater from [RVTools website](#)
2. Launch RVTools and log in to the target VCenter Server.

This will capture information about that VCenter Server.

3. Export your VMware information to an `xlsx` or `xls` file.

See the chapter on "commandline parameters" in the [RVTools documentation](#) for details.

Result

The `xlsx` or `xls` file that contains your results is saved to your working directory.

What's next?

[Create a VMware Cloud on AWS deployment plan using the migration advisor.](#)

Create a deployment plan for VMware Cloud on AWS using the VMware workloads migration advisor

Log in to the NetApp workload factory to access the VMware migration advisor. You'll follow the steps in the wizard to create a deployment plan or migration plan that is customized for your needs.

When migrating to VMware Cloud, you can use the migration advisor to create a deployment plan for the following scenarios:

- [To migrate your current on-premises vSphere environment to VMWare cloud.](#)
- [To deploy a system in the cloud based on an existing deployment plan that has similar requirements.](#)

Note that you must have a user name and password to access workload factory. If you don't have access, create an account now. See the instructions [here](#).

Create a deployment plan based on an on-premises vSphere environment

You can migrate your current virtual machine configurations in on-premises vSphere environments to virtual machines in VMware Cloud on AWS and use customized Amazon FSx for NetApp ONTAP file systems as external datastores.

Requirements

- You must have created the inventory file from your existing systems by using the migration advisor VM collector (.csv file) or by using RVTools (.xlsx file).
- You must have access to the inventory file from the system on which you are logging in to workload factory.

Steps

1. Log in to workload factory using one of the [console experiences](#).
2. From the VMware workloads tile, select **Assess and plan** and then select **Migration to VMware Cloud on AWS**.

The VMware migration advisor main page is displayed.

3. Select **Create a new deployment plan**.
4. Select the type of inventory file you'll be using to populate workload factory with your current VM configuration and select **Next**.

- Select **Use the migration advisor VMware data collector** to use the .csv file that you created using the VMware data collector.
- Select **Use RVTools** to use the .xlsx file that you created using RVTools.

The "Prepare for VMware Cloud onboarding" page is displayed.

5. In the *Upload VM configuration* section, select  and select the file that you want to use.
 - Select the .csv file when using the Migration advisor VM collector.
 - Select the .xlsx file when using RVTools.

The VMs summary section is populated from the inventory file to reflect the number of VMs and the total storage capacity.

6. In the *VM inventory considerations* section, select the options to filter the list of VMs that you want to migrate.
 - a. **Consider VMs:** Indicate which VMs will be extracted from the .csv file based on their operating power state. You can bring in all VMs, or only those that are On, Off, or Suspended.
 - b. **VM Storage to consider:** Select whether the datastores created for each onboarded VM are sized based on their currently utilized size (recommended) or their provisioned size.

The external datastores will be implemented using Amazon FSx for NetApp ONTAP file system volumes.

7. In the *VMware Cloud on AWS deployment configuration* section, enter details about the required VMware

Cloud on AWS configuration.

- a. **Region:** Select the region where the VMs and Amazon FSx for NetApp ONTAP file systems will be deployed.

For optimal performance and cost efficiency, this is typically the same region as where your existing VMware Cloud on AWS SDDC is deployed.

- b. **VM estimated performance requirements:** This option is available only when using RVTools. The Migration advisor VM collector captures this information from your environment. Provide the following per virtual machine average performance parameters that you want to be applied to your new VMs that will be deployed:

- **Average IOPS per VM:** Enter the number of IOPS required for your file systems. If you are unsure, you can use the default of 3 IOPS per GiB of SSD storage for Amazon FSx for ONTAP file systems. For example, if you deploy 2,000 GiB of capacity, this will be translated to 6,000 IOPS. We recommend that you begin with a smaller IOPS setting. You can increase your provisioned SSD IOPS after the file system is created as workloads are migrated or deployed.
- **Average I/O block size:** The size of each block containing read or write operations. The default size is 4 KB. A larger block size may be better for large sequential read and write workloads. A smaller block size may offer better performance for workloads that perform small random writes to sparse files or to large files.
- **Average write ratio:** The percentage of operations that are write operations for your workloads. The default ratio is 30% writes and 70% reads.

8. In the *VM storage capacity considerations* section, select from a few storage options.

- a. **Average data reduction ratio:** Choose from among the three common data reduction selection values. Select "1:1 - No reduction", "1:1.25 - 20% reduction", or "1:1.5 - 33% reduction".
- b. **Headroom percentage:** Enter the percentage of capacity growth that is added to the capacity for your FSx for ONTAP file systems.

Note that if you select an amount less than 20% that you won't be able to create volume snapshots for protection and long-term backups.

9. Select **Next** and the "VMware cloud on AWS node configuration" page is displayed.

This page enables you to define the VMware cloud on AWS cluster configuration using an estimated savings analysis and the recommended node type. You can configure the following:

- a. **vSAN architecture:** Select whether you want to use vSAN Express Storage Architecture (ESA) or vSAN Original Storage Architecture (OSA) architecture.
- b. **vSAN Fault Tolerance:** Select the level of fault tolerance that is required for the VMs. You can choose "Auto", which is recommended, or from among a variety of RAID levels.
 - RAID-1 (FTT 1): consists of an exact copy (or mirror) of a set of data on 2 or more disks.
 - RAID-5 (FTT 1): consists of block-level striping with distributed parity - parity information is distributed among 3 or more drives, and it can survive a single disk failures.
 - RAID-5 (FTT 2): consists of block-level striping with distributed parity - parity information is distributed among 4 or more drives, and it can survive any two concurrent disk failures.
 - RAID-6 (FTT 2): extends RAID 5 by adding another parity block; thus, it uses block-level striping with two parity blocks distributed across all member disks. It requires 4 or more drives, and it can survive any two concurrent disk failures.
- c. **Nodes configuration selection list:** Select an EC2 instance type for the nodes.

10. Select **Next** and the "Select virtual machines" page displays the VMs that match the criteria you provided in the previous page.
 - a. In the *Selection criteria* section, select the criteria for the VMs that you plan to deploy:
 - Based on cost and performance optimization
 - Based on the ability to easily restore your data with local snapshots for recovery scenarios
 - Based on both sets of criteria: the lowest cost while still providing good recovery options
 - b. In the *Virtual machines* section, the VMs that matched the criteria you provided in the previous page are selected (checked). Select or deselect VMs if you want to onboard/migrate fewer or more VMs on this page.

The **Recommended deployment** section will be updated if you make any changes. Note that by selecting the checkbox in the heading row you can select all VMs on this page.

 - c. Select **Next**.
11. On the **Datastore deployment plan** page, review the total number of VMs and datastores that have been recommended for the migration.
 - a. Select each Datastore listed across the top of the page to see how the datastores and VMs will be provisioned.

The bottom of the page shows the source VM (or multiple VMs) for which this new VM and datastore will be provisioned.

 - b. Once you understand how your datastores will be deployed, select **Next**.
12. On the **Review deployment plan** page, review the estimated monthly cost for all the VMs that you plan to migrate.

The top of the page describes the monthly cost for all deployed VMs and FSx for ONTAP file systems. You can expand each section to view details for "Recommended Amazon FSx for ONTAP file system configuration", "Estimated cost breakdown", "Volume configuration", "Sizing assumptions", and technical "Disclaimers".

13. When you are satisfied with the migration plan, you have a few options:
 - Select **Deploy** to deploy the FSx for ONTAP file systems to support your VMs. [Learn how to deploy an FSx for ONTAP file system](#).
 - Select **Download plan > VM deployment** to download the migration plan in a .csv format so you can use it to create your new cloud-based intelligent data infrastructure.
 - Select **Download plan > Plan report** to download the migration plan in a .pdf format so you can distribute the plan for review.
 - Select **Export plan** to save the migration plan as a template in a .json format. You can import the plan at a later time to use as a template when deploying systems with similar requirements.

Create a deployment plan based on an existing plan

If you are planning a new deployment that is similar to an existing deployment plan that you've used in the past, you can import that plan, make edits, and then save it as a new deployment plan.

Requirements

You must have access to the .json file for the existing deployment plan from the system on which you are logging in to workload factory.

Steps

1. Log in to workload factory using one of the [console experiences](#).
2. From the VMware workloads tile, select **Assess and plan** and then select **Migration to VMware Cloud on AWS**. The VMware migration advisor main page is displayed.
3. Select **Import an existing deployment plan**.
4. Select  and select the existing plan file that you want to import in the migration advisor.
5. Select **Next** and the Review plan page is displayed.
6. You can select **Previous** to access the *Prepare for VMware Cloud onboarding* page and the *Select VMs* page to modify the settings for the plan as described in the previous section.
7. After you have customized the plan to your requirements, you can save the plan or begin the deployment process for your datastores on FSx for ONTAP file systems.

Deploy the recommended FSx for ONTAP file system

After you verify that the recommended FSx for ONTAP file system (or multiple file systems in some cases) meets your exact requirements, you can use workload factory to deploy the system in your AWS environment.

Depending on the policy and permissions that you added to your workload factory account, you can deploy the FSx for ONTAP file system completely using workload factory (using read/write mode). If you have fewer permissions (read-only mode), or no permissions (Basic mode), you'll need to use the CloudFormation information from the Codebox and deploy the FSx for ONTAP file system yourself in AWS.

Requirements for deployments to VMware Cloud on AWS

- You must be using VMware Cloud on AWS Software-Defined Data Center (SDDC) version 1.20 or greater to deploy FSx for ONTAP file systems.
- You must not deploy the FSx for ONTAP file system in the same VPC used during the SDDC deployment. Instead, you must deploy it in a new Amazon VPC that you own to enable VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP.
- You must deploy the FSx for ONTAP file system within the same AWS region as your SDDC.

Steps

1. At the bottom of the **Review plan** page, select **Deploy** and the Create an FSx for ONTAP file system page is displayed.

Most of the fields that define your FSx for ONTAP file system are completed based on the information you provided, but there are a few fields that you need to complete in this page.

You can use the Quick create or Advanced create option. Advanced create offers a few additional storage parameters that you can customize. [See what these two options offer](#).

2. **AWS credentials:** Select or add credentials that will give workload factory the permissions necessary to create your FSx for ONTAP file system directly. You can also select the CloudFormation code from Codebox and deploy the FSx for ONTAP file system yourself in AWS.
3. **File system name:** Enter the name that you want to use for this FSx for ONTAP file system.
4. **Tags:** Optionally you can add tags to categorize this FSx for ONTAP file system.
5. In the "Network & security" section, enter the following information:

- Region & VPC:** Select the Region and the VPC where the FSx for ONTAP file system will be deployed.

If you are deploying to VMware Cloud on AWS, ensure you deploy it in a VPC that is different than the VPC where the VMware Cloud on AWS is deployed.

- Security group:** When using the **Advanced create** option, you can select the default security group for the FSx for ONTAP VPC so that all traffic can access the FSx for ONTAP file system.

You can add an inbound rule that restricts what other AWS services can access the FSx for ONTAP file system. This will block the amount of services that are open. These are the minimum ports and protocols:

Protocols	Ports	Purpose
TCP, UDP	111	Portmapper (used to negotiate which ports are used in NFS requests)
TCP, UDP	635	NFS mountd (receives NFS mount requests)
TCP, UDP	2049	NFS network traffic
TCP, UDP	4045	Network Lock Manager (NLM, lockd) - Handles lock requests.
TCP, UDP	4046	Network Status Monitor (NSM, statd) - Notifies NFS clients about reboots of the server for lock management.

- Availability zone:** Select the Availability Zone and the Subnet.

You should select the same availability zone as where your VMware SDDC is deployed if you want to avoid charges for cross-AZ traffic.

- Encryption:** When using the **Advanced create** option, you can select the AWS encryption key name from the dropdown.
- Datastore access control:** When using the **Advanced create** option, you can select whether all hosts can access the datastores or whether only certain vSphere cluster nodes on a specific subnet can access the datastores.

- In the "File system details" section, enter the following information:
 - ONTAP credentials:** Enter and confirm the ONTAP password.
 - Storage VM credentials (Advanced create only):** Enter and confirm the storage VM password. The password can be specific to this file system, or you can use the same password entered for ONTAP credentials.
- In the **Summary** section, you can view the FSx for ONTAP file system and datastore configuration that the VMware migration advisor has designed based on your information.
- Select **Create** to deploy the FSx for ONTAP file system. This process can take up to 2 hours.

Optionally, in the Codebox window you can select **Redirect to CloudFormation** to create the file system using a CloudFormation stack.

In either case, you can monitor the creation progress in CloudFormation.

Result

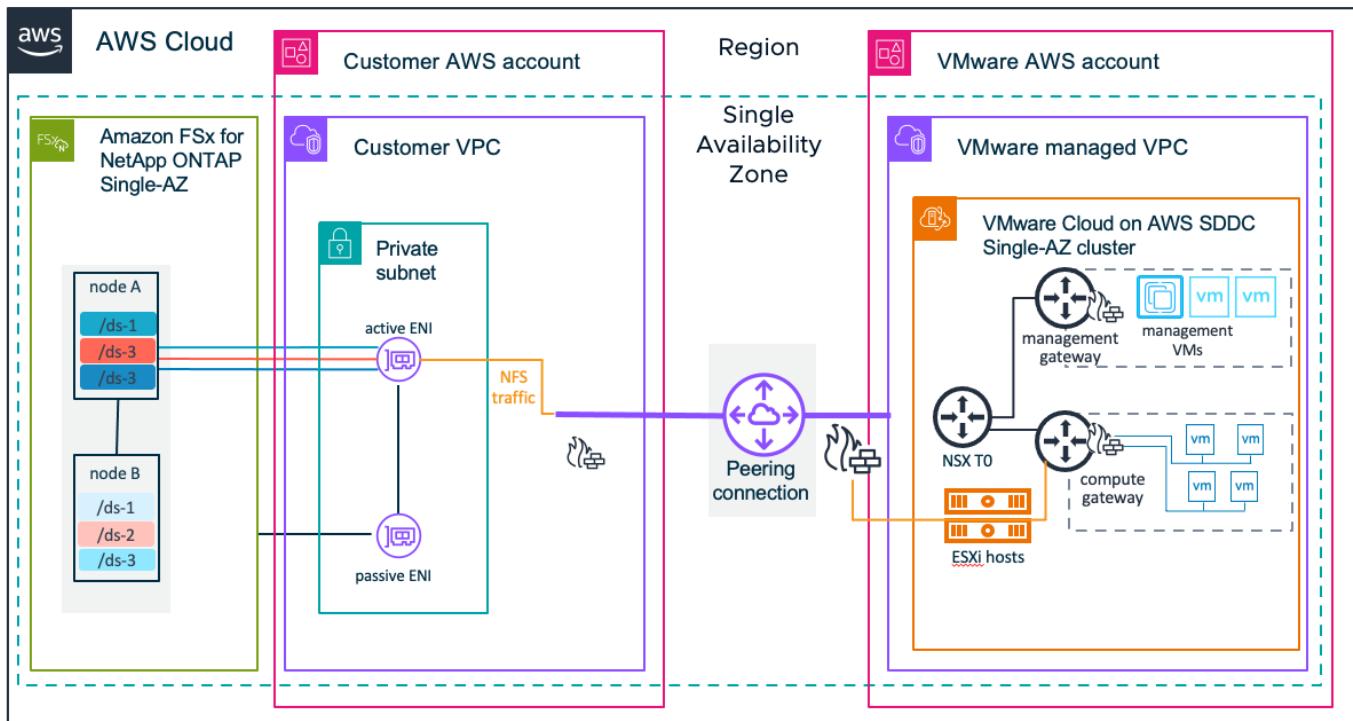
The FSx for ONTAP file system is deployed.

Connect your FSx for ONTAP file systems to VMware Cloud on AWS

After you deploy FSx for ONTAP file systems, you need to connect that system to your VMware Cloud on AWS infrastructure. Your Software-Defined Data Center (SDDC) offers network options for connecting to the FSx for ONTAP file system by using the VPC peering capability to extend network connectivity to the external NFS storage volumes.

[Review the Amazon VPC peering documentation for details](#)

The following diagram shows how to use VPC peering to connect single Availability Zone (AZ) SDDC clusters to FSx for ONTAP single-AZ deployments.



Migrate your data to the new infrastructure with BlueXP workload factory for VMware

Use a tool such as VMware Hybrid Cloud Extension (HCX) to move your data from your old virtual machine storage to the FSx for NetApp ONTAP volumes connected to your new virtual machines. VMware HCX core functionality enables you to migrate workloads from your on-premises data center to your Software-Defined Data Center (SDDC) transparently.

[Review the VMware HCX documentation for details.](#)

What's next?

Now that you've migrated your data to VMware Cloud on AWS and Amazon FSx for NetApp ONTAP external datastores, you can back up and protect the important data on your FSx for ONTAP file system to make sure your data is always available.

For information about managing your FSx for ONTAP file system, go to the [Amazon FSx for NetApp ONTAP documentation](#) to view the the backup and protection capabilities you can use.

Copyright information

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

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

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

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

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

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

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

Trademark information

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