



Use the VMware planning center

Workload Factory for NetApp Console

NetApp
October 07, 2026

Table of Contents

- Use the VMware planning center 1
 - Upload VM inventory data to NetApp Workload Factory 1
- Migrate to Amazon Elastic VMware Service 3
 - Explore savings for Amazon Elastic VMware Service with NetApp Workload Factory 3
 - Manually create a deployment plan for Amazon EVS 4
 - Create a deployment plan for Amazon EVS using the migration advisor 5
 - Deploy the recommended FSx for ONTAP file system 8
- Migrate to Google Cloud VMware Engine 10
 - Create a deployment plan for Google Cloud VMware Engine using the migration advisor 10
- Migrate to Amazon EC2 15
 - Create a deployment plan for Amazon EC2 using the VMware workloads migration advisor 15
 - Deploy the recommended FSx for ONTAP file system 18
- Migrate to VMware Cloud on AWS 18
 - Create a deployment plan for VMware Cloud on AWS using the VMware workloads migration advisor. . 18
 - 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 NetApp Workload Factory for VMware 23
- Optimize VMware workloads on FSx for ONTAP 24
 - Configuration analysis for EVS configurations 24
 - Implement well-architected EVS configurations 25

Use the VMware planning center

Upload VM inventory data to NetApp Workload Factory

Before you can create a migration plan, you need to upload your VM inventory data.

Steps

1. Log in to the [NetApp Console](#).

2. Select the menu , **Workloads > VMware**.

The planning center is displayed. From the planning center, you can upload your VM inventory data, see previous inventory data uploads, and plan a migration using any of the VM inventory datasets.

3. Select **Upload VM inventory**.

4. Choose which data collection method to use:

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

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

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

6. Enter a name for the VM inventory dataset you will create.

7. Download and install RVTools 4.4.2 or greater from the [RVTools website](#)

8. Launch RVTools and log in to the target VCenter Server.

This will capture information about that VCenter Server.

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

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

10. Select  and choose the RVTools file you exported.

11. Select **Upload**.

12. Select **Finish**.

Use the Data collector (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

5. Select the **Data collector** tile, and select **Next**.

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

6. 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.

7. 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.

+ 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.

8. Select **Upload**.

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

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

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

6. Enter a name for the VM inventory dataset you will create.

7. Enter the tenant endpoint for Data Infrastructure Insights.

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

8. 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.

9. Select **Upload**.

Workload Factory gathers information from Data Infrastructure Insights.

What's next?

Manage your uploaded VM inventory data from the planning center, or create a deployment plan.

- [Explore the VMware workloads planning center](#)
- [Create an Amazon VMware Cloud deployment plan](#)
- [Create an Amazon Elastic VMware Service deployment plan](#)
- [Create an Amazon EC2 deployment plan](#)
- [Create a Google Cloud VMware Engine deployment plan](#)

Migrate to Amazon Elastic VMware Service

Explore savings for Amazon Elastic VMware Service with NetApp 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 the [NetApp Console](#).

2. Select the menu , **Workloads > VMware**.

The planning center is displayed.

3. From the VMware menu, select **Explore savings**. The savings calculator is displayed.

In the **Environment preferences** pane, you can select your EVS deployment region and adjust TCO and savings predictions accordingly.

4. From the **Region** drop-down list, select the EVS deployment region to calculate savings.

5. 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:

- **Required physical CPU (count)**
- **Required physical memory (GiB)**
- **Required VM storage (TiB)**

6. Under **EVS billing plan**, select the plan you are on.

7. Do one of the following:

- To use the migration advisor to [create a deployment plan](#) for an Amazon EVS environment, select **Plan EVS migration**.
- Select **Close** to close the TCO calculator.

Manually create a deployment plan for Amazon EVS

Log in to the NetApp Workload Factory to access the VMware planning center. From there, you can manually create a deployment plan or migration plan for Amazon EVS that is customized for your needs.

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 a user name and password to access Workload Factory. If you don't have access, create an account now. Review the instructions [here](#).
- You must have an Amazon Elastic VMware Service subscription.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > VMware**.

The planning center is displayed.
3. Select **Manually create EVS plan**.
4. 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.

5. When ready, select **Review plan** to review the migration plan.
6. Review the plan. Expand each section to view the plan details.

By default, Workload Factory saves the migration plan to the planning center. You can deselect this option near the top of the page.

7. Optionally, you can export the migration plan as a PDF or CSV file by selecting **Manage plan** at the top right of the page and then selecting **Download a report (.pdf)** or **Download VM storage deployment (.csv)**.

The .csv file creates a map of all the VMs included in the migration plan, along with their assigned storage volumes.

8. When you are ready to provision the deployment plan, select **Provision**.

[Deploy the VMware workloads recommended FSx for ONTAP file system.](#)

Create a deployment plan for Amazon EVS using the migration advisor

From the VMware planning center, you can launch the Amazon EVS migration advisor to help create a migration plan that is customized for your needs.

You can create a deployment plan to migrate virtual machines to Amazon Elastic VMware Service and use customized Amazon FSx for NetApp ONTAP file systems as external datastores. The options in the migration

advisor might differ depending on the tool used to collect the VM inventory data.

Requirements

- You must have a user name and password to access Workload Factory. If you don't have access, create an account now. Review the instructions [here](#).
- You must have an Amazon Elastic VMware Service subscription.

Steps

1. Log in to the [NetApp Console](#).

2. Select the menu , **Workloads > VMware**.

The planning center is displayed.

3. Select a VM inventory from the list that you want to use to make a deployment plan, and select **Start planning** in that row.

4. Select **EVS** from the drop-down menu that appears.

The **Prepare for AWS Cloud onboarding** wizard appears.

5. Enter the required information.

Specify

1. In the *VM configuration upload* section, review information about the dataset you are using to create a migration plan. The **VM inventory summary** section is populated from the inventory file to reflect the number of VMs and the total storage capacity.

2. 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.

3. 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.

4. Select **Next**.

Select

1. On the **Select virtual machines** 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 *Peak read IOPS* displays a column with the peak read IOPS for each VM.

2. Optionally, you can choose to optimize the deployment for cost or recoverability.
 - **Cost:** Workload Factory chooses VMs from the list that have lower I/O density. This helps reduce resource requirements.
 - **Recoverability:** Workload Factory chooses VMs from the list that are the easiest to quickly copy locally. This offers quick recovery times in case of disruptions.

3. Select **Next**.

Design

- On the **ClassReview instance storage assignment** 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**.

Review plan

1. 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.

The migration plan is automatically saved in the list of plans in the planning center by default.

2. Optionally, you can export the migration plan as a PDF or CSV file by selecting **Manage plan** at the top right of the page and then selecting **Download a report (.pdf)** or **Download VM storage deployment (.csv)**.

The .csv file creates a map of all the VMs included in the migration plan, along with their assigned storage volumes.

3. When you are ready to proceed with the plan, select **Provision** to begin deploying the recommended Amazon FSx for NetApp ONTAP file system.

[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

- 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.

Considerations

- You can use the Quick create or Advanced create option. Advanced create offers a few additional storage parameters that you can customize. [Learn what these two options offer.](#)
- In the "Amazon Elastic VMware Service preferences" section, you can choose the EVS virtualization environment to connect to the external datastores. This auto-populates some of the fields with the best-practice options for the deployment. You can change these options as required.

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.

2. In the "File system general configuration" section, enter the following information:
 - a. **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.
 - b. **File system name:** Enter the name that you want to use for this FSx for ONTAP file system.
 - c. **Tags:** Optionally, you can add tags to categorize this FSx for ONTAP file system.
3. In the "Amazon Elastic VMware Service preferences" section, from the "Environment ID" list, select the EVS environment where the datastores will be connected.

This auto-populates the following fields:

- **VPC**
- **Availability Zone**

- **Subnet**
- In the **NFS Datastores access** section, if you have selected the **EVS host management only** option the **EVS host management (CIDR)** field is auto-populated.
- In the **Datastore mount options** section, if you have chosen to mount datastores to the EVS cluster, the **EVS Cluster VMware vCenter address** and the **vSphere administrator credentials secret ARN** fields are auto-populated.

4. In the "Network & security" section, enter the following information:

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

5. 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.

6. 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.

7. 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.
8. 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 Google Cloud VMware Engine

Create a deployment plan for Google Cloud VMware Engine using the migration advisor

From the VMware planning center, you can launch the Google Cloud VMware Engine migration advisor to help create a migration plan that is customized for your needs.

You can create a deployment plan to migrate virtual machines to Google Cloud VMware Engine and use customized Google Cloud NetApp Volumes as external NFS datastores. The options in the migration advisor might differ depending on the tool used to collect the VM inventory data.

Requirements

- You must have [uploaded a VM inventory](#).
- You must have a user name and password to access Workload Factory. If you don't have access, create an account now. Review the instructions [here](#).
- You must have a Google Cloud VMware Engine subscription in a region where both Google Cloud VMware Engine and Google Cloud NetApp Volumes services are available.

Steps

1. Log in to the [NetApp Console](#).

2. Select the menu , **Workloads > VMware**.

The planning center is displayed.

3. Select **Google Cloud** from the drop-down menu.

4. From the VM inventory table, select **Start planning** next to the inventory that you want to use to make a deployment plan.

The **Prepare for Google Cloud onboarding** wizard appears.

5. Enter the required information.

Configure migration preferences

1. In the *VM configuration upload* section, review information about the dataset you are using to create a migration plan. The **VM inventory summary** section is populated from the inventory file to reflect the number of VMs and the total storage capacity.
2. In the *VM inventory considerations* section, select the options to filter the list of VMs and configure how the migration plan is designed:
 - a. **VM inventory considerations:** Select which VMs to include based on their power state:
 - **Powered on** (selected by default)
 - **Powered off**
 - **Suspended**
 - b. **VM storage to consider:** Choose how to size the Google Cloud NetApp Volumes datastores:
 - **Utilized:** Size datastores based on the actual used capacity of VMs.
 - **Provisioned** (default): Size datastores based on the total provisioned capacity of VMs.
 - c. **VM memory to consider:** Choose how to calculate the required Google Cloud VMware Engine cluster compute resources:
 - **Utilized:** Calculate based on actual memory usage by VMs.
 - **Provisioned** (default): Calculate based on total memory allocated to VMs.
 - d. **Virtual machines performance estimation:** This option appears only when using RVTools as your data source. Select the expected performance level:
 - **Low:** 50 IOPS on average per VM.
 - **Standard:** 200 IOPS on average per VM (selected by default).
 - **High:** 800 IOPS on average per VM.
3. In the *Google Cloud VMware Engine deployment region* section:
 - a. **Region:** Select the Google Cloud region where Google Cloud VMware Engine and Google Cloud NetApp Volumes will be deployed. Only regions where both services are available appear in the list.
4. In the *VM storage capacity and protection considerations* section:
 - a. **Capacity headroom percentage:** Enter the percentage of additional capacity to add for growth. The default value is 20%, which is also the recommended minimum.



If you select less than 20% headroom, snapshot protection is automatically disabled.

- b. **VM snapshot protection:** Enable this option (enabled by default) to protect VMs with snapshots using the ONTAP default snapshot policy.

5. Select **Next**.

Select a cluster configuration

1. On the **Select preferred Google Cloud VMware Engine cluster configuration** page, configure your Google Cloud VMware Engine cluster preferences and review the recommended configurations.
2. In the *Architecture preferences* section:
 - a. **Allow Google Cloud VMware Engine VE2 node types:** When enabled (default), both VE1 and VE2 node types are considered. Disable this option to limit configurations to VE1 nodes only.
 - b. **Allow Google Cloud NetApp Volumes standard volumes:** When disabled (default), only Premium and Extreme service levels are used in the configuration design, which is the recommended best practice. Premium and Extreme volumes provide regional availability, making them resilient to availability zone failures. Standard volumes provide only zonal availability and are not protected against AZ failures. Enable this option to include Standard service level volumes, which could reduce costs for lower-performance workloads where AZ-level resiliency is not required.
3. In the *Cluster configuration* section, review the available Google Cloud VMware Engine cluster configurations:

The configurations are sorted by estimated overall cost (lowest to highest) and show different combinations of:

- Google Cloud VMware Engine node types (VE1 or VE2)
 - Estimated overall cost
 - Estimated savings
 - Node count
 - VSAN fault tolerance settings: FTT (Fault Tolerance Type) and FTM (Fault Tolerance Method)
 - Excess storage capacity in TiB
 - Storage-only node type
4. Select your preferred cluster configuration by clicking the radio button next to it.

The first (lowest cost) configuration is selected by default.

5. Select **Next**.

Select a VM layout

1. On the **Automatic VM selection** page, review the optimal virtual machine layout automatically selected based on your required external storage capacity and preferences.
2. Review the *VM inventory* section, which displays:
 - Overall VM count
 - Overall capacity (provisioned or utilized based on your earlier selection)

You can expand the section to view the following details:

- Aggregated average IOPS
- Aggregated throughput (MB/s)
- Aggregated peak IOPS
- Aggregated peak throughput (MB/s)

3. In the *VM selection preferences* section, choose the optimization method:

- **Cost** (default): Assigns VMs with the highest I/O density to vSAN storage, which helps minimize Google Cloud NetApp Volumes costs while meeting performance requirements.
- **Protection**: Assigns VMs with the highest storage capacity to Google Cloud NetApp Volumes external datastores, which provides better recoverability through Google Cloud NetApp Volumes snapshot and replication capabilities.

4. Review the *External datastore configuration* section, which shows how VMs and capacity are distributed between Google Cloud NetApp Volumes and vSAN storage.

You can expand the section to view the following details: * Aggregated average IOPS * Aggregated throughput (MB/s) * Aggregated peak IOPS * Aggregated peak throughput (MB/s)

5. In the *VMs selected for Google Cloud NetApp Volumes NFS datastores* table, you can:

- View the automatically selected VMs based on your chosen optimization method.
- Manually select or deselect individual VMs by using the checkboxes.
- Filter VMs by cluster, operating system, or datacenter.
- Sort VMs by any visible column (such as VM name, provisioned capacity, average read IOPS).
- Show or hide additional columns by clicking .



If you manually modify the VM selection, you can revert to the automatic selection based on your optimization preference by selecting **Reset**.

6. Select **Next**.

Review NFS datastore assignments

1. On the **Review instance store assignment** page, review the detailed datastore layout and VM assignments.
2. Use the **Total** and **Per datastore** tabs to switch between an aggregate view and per-datastore details.
3. In the **Total** tab:
 - a. Review the *Storage Capacity analysis* chart that displays:
 - Google Cloud NetApp Volumes logical capacity
 - VM provisioned capacity
 - VM level free space (if provisioned capacity was selected)
 - Overall capacity headroom
 - Snapshots buffer capacity (if snapshot protection is enabled)
 - b. Review the *Detailed information* section showing:
 - Overall VMs provisioned capacity

- Overall capacity headroom
- Free space (if provisioned capacity was selected)
- Overall logical capacity for all Google Cloud NetApp Volumes storage pools
- Google Cloud NetApp Volumes storage pool count and types
- Overall VMs count and provisioned capacity
- Aggregated VM average IOPS (when available)
- Average I/O density (with tooltip explaining calculation)
- Aggregated VM average throughput (when available)

4. In the **Per datastore** tabs:

- a. Select a datastore tab to view its details.
- b. Review the details for that datastore, which includes:
 - Capacity breakdown chart
 - Datastore properties:
 - Overall VMs utilized capacity
 - Overall capacity headroom
 - Free space (if provisioned capacity was selected)
 - Datastore logical capacity
 - Datastore service level (Standard, Premium, Extreme)
 - Overall VMs count and provisioned capacity
 - Aggregated VM average IOPS (when available)
 - Average datastore I/O density (when available)
 - Aggregated VM average throughput (when available)
- c. Review the *Virtual machines* table showing all VMs assigned to this datastore with columns for VM name, capacity, cluster, operating system, IOPS, throughput, and other properties.

5. Select **Next**.

Review and save your deployment plan

1. On the **Review your migration plan** page, review and customize your deployment plan before saving it.
2. From the drop-down, choose your pricing plan:
 - **On-demand** (default): Pay-as-you-go pricing



The pricing displayed throughout the review page adjusts based on your selected plan.

3. Review the *Estimated savings* section:

- a. The *Google Cloud VMware Engine configuration alternatives comparison* charts compare:
 - Google Cloud VMware Engine vSAN-only configuration estimated monthly cost (compute nodes + storage-only nodes)
 - Google Cloud VMware Engine vSAN + external datastore estimated monthly cost (compute nodes + Google Cloud NetApp Volumes external datastores)

- Savings percentage and dollar amount
 - b. The *Estimated savings breakdown* table compares the two configurations with details on node counts, storage capacity, Google Cloud NetApp Volumes pool information, and monthly and yearly cost savings.
4. Review the *Google Cloud NetApp Volumes suggested configuration* section. It highlights:
- Storage pool count and types
 - Overall storage pool capacity
 - Datastores (volumes) count
 - Datastores aggregated throughput
 - Estimated monthly cost
 - a. The *Suggested configuration details* section provides a breakdown by service level (Standard, Premium, Extreme) with storage pool counts, capacities, datastore counts, throughput details, capacity headroom, snapshot buffer, and total estimated monthly cost.
 - b. The *Estimated cost analysis* section breaks down costs by service level (Standard, Premium, Extreme) including capacity costs, cold capacity, replicated data, backup usage and capacity, with a total estimated monthly cost.
 - c. The *Volume configuration* table lists each planned Google Cloud NetApp Volumes volume with details such as name, pool type, protocol (NFSv3), capacity, headroom, snapshot policy, and tiering policy.
5. Review the *Solution design assumptions* and *Disclaimer* sections.
6. When you are satisfied with the migration plan, you have a few options:
- Select **Manage plan** at the top of the page.
 - Choose one of the following downloads:
 - **Download a report:** A comprehensive PDF document with all plan details
 - **Download instance storage deployment:** A CSV file mapping VMs to their assigned datastores
7. Select **Provision** to provision the Google Cloud NetApp Volumes NFS datastores using Terraform Infrastructure-as-Code.

A codebox displays the generated Terraform code. You can copy the code to your clipboard or download it as a file to execute in your own Terraform environment.

8. Select **Create** to create the Google Cloud NetApp Volumes NFS datastores.

Migrate to Amazon EC2

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. Review 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 [uploaded a VM inventory](#).

Steps

1. Log in to the [NetApp Console](#).

2. Select the menu , **Workloads > VMware**.

The planning center is displayed.

3. If you have not yet uploaded a VM inventory, select **Upload VM inventory** and follow the instructions in [Upload VM inventory data](#).

4. Choose a VM inventory dataset from the list, and open the **Start planning** menu for the dataset.

5. Select **AWS native compute**.

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.

- 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.
- 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.

- 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.

- 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.

The migration plan is automatically saved in the list of plans in the planning center by default.

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

- Select **Download plan > Download a report** to download the deployment plan in a .pdf format so you can distribute the plan for review.
- Select **Download 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.
- Select **Provision** to begin deploying the recommended Amazon FSx for NetApp 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.

Follow the instructions in [Create an FSx for ONTAP file system in NetApp 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

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 migrate your current on-premises vSphere environment to VMWare cloud.

Note that you must have a user name and password to access Workload Factory. If you don't have access, create an account now. Review 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 the [NetApp Console](#).

2. Select the menu  **Workloads > VMware**.

The planning center is displayed.

3. Select **Plan and create VMC**.

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.

- 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.
- 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.

- VM Memory to consider:** Select whether the memory allocated for each onboarded VM is sized based on their currently utilized size (recommended) or their provisioned size.

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

- 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.

- 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.

- 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".
- 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.

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. [Learn 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:

- a. **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.

- b. **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.

- 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:** When using the **Advanced create** option, you can select the AWS encryption key name from the dropdown.
- e. **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.

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

- a. **ONTAP credentials:** Enter and confirm the ONTAP 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.

7. 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.

8. 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 NetApp 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 backup and protection capabilities you can use.

Optimize VMware workloads on FSx for ONTAP

Configuration analysis for EVS configurations

Workload Factory analyzes Amazon Elastic VMware Service (EVS) configurations regularly to identify misalignments with best practices. Use the results to improve performance, cost efficiency, and compliance.

Key capabilities include:

- Daily configuration analysis
- Automatic best practice validations
- Proactive observability
- Insights to action

Unresolved directive in fsx-vmware/configuration-analysis.adoc - include::_include/well-architected-workloads.adoc[]

Analysis requirements

For a complete configuration analysis, you must do the following:

- Associate a link. Link connectivity lets Workload Factory analyze all file system configurations like data protection and performance.

[Learn how to associate an existing link or to create and associate a new link.](#)

- Grant *view, planning, and analysis* permissions for VMware workloads.

[Learn how to grant permissions to an AWS account](#)

- You must have at least one discovered Amazon Elastic VMware Service environment in your AWS account.

Best practices and recommendations for EVS workloads

Workload Factory provides best practices and recommendations for operating well-architected Amazon Elastic VMware Service (EVS) workloads. The well-architected analysis assesses EVS configurations to help ensure your VMware environments are optimized for reliability, security, operational excellence, cost optimization, and performance efficiency. From the well-architected status tab in VMware, you'll find insights and recommendations to help implement well-architected best practices for your EVS environments.

The well-architected analysis categorizes configurations in the following pillars of the framework: *reliability* and *security*.

Reliability

Reliability ensures that workloads perform their intended functions correctly and consistently, even when there are disruptions.

- **EVS environment resiliency**

Ensure that your EVS cluster nodes are properly distributed across partition placement groups. All nodes should be members of a single partition placement group configured with four or more partitions. Proper partition placement ensures that your EVS cluster nodes are distributed across multiple fault-isolated hardware partitions within an AWS availability zone. Misalignment can result in significant loss of processing power or downtime if a partition fails.

Security

Security emphasizes protecting data, systems, and assets through risk assessments and mitigation strategies.

- **Cluster node management**

Ensure that your EVS cluster nodes have appropriate EC2 stop and termination protection configured. EVS ESXi nodes should be managed exclusively using vCenter or other VMware-level management tools. Without proper EC2-level protections, nodes could be accidentally stopped or terminated from the EC2 console, which can lead to virtual machine data unavailability or data loss.

What's next

[Implement well-architected configurations](#)

Implement well-architected EVS configurations

Use Workload Factory configuration analysis to review well-architected status for your Amazon Elastic VMware Service (EVS) configurations and remediate issues that affect reliability, security, and cost.

About this task

Automatic daily scans of all discovered EVS environments using AWS APIs analyze your EVS configuration and identify potential issues that could impact availability, resiliency, security, or cost optimization. Findings are organized by configuration area, with each finding including status, severity levels, impacted resource details, and step-by-step remediation procedures.

Key features include:

- **Automatic daily scans:** All discovered EVS environments are automatically scanned once a day to ensure insights remain current.
- **AWS API-based scanning:** Scans use AWS APIs and do not require vSphere credentials or connectivity to your vCenter.
- **Detailed guidance on issue resolution:** Each identified issue includes a clear explanation, severity level, and step-by-step resolution procedures.
- **View-only insights:** Provides detailed findings and recommendations without automated issue resolution options.

Understanding well-architected insights

The Well-architected tab displays the following:

- **Configuration name:** The configuration area being assessed.
- **Tags:** Labels indicating the areas of impact (such as Availability, Resiliency, Security).
- **Status:** Either "Optimized" (no issues found) or "Not optimized" (issues found).
- **Severity:** The importance level of the finding (for example, Warning).
- **Resource type:** The type of AWS resource being assessed.
- **Impacted resources count:** The number of resources affected by the issue.

Scan frequency

Well-architected scans are performed automatically for all discovered EVS configurations. Key details about scan scheduling:

- Scans occur once per day for each EVS configuration.
- Scans for different configurations can occur at different times.
- If a scan fails for one configuration, scans for other configurations in the same account will still be attempted.
- The timestamp card on the Well-architected status tab shows when the last scan was completed for the current configuration.



On-demand execution of well-architected scans is not currently supported. All scans are performed automatically on the daily schedule.

Before you begin

- You must have [added AWS credentials](#) with *View, planning, and analysis* permissions for VMware workloads.
- You must have at least one discovered Amazon Elastic VMware Service environment in your AWS account.

Access the well-architected status tab

Steps

1. Log in to the [NetApp Console](#).

2. Select the menu , **Workloads > VMware**.

The planning center is displayed.

3. From the VMware menu, select **Inventory**.

4. From the **Virtualization environments** list, select the discovered EVS environment you want to view well-architected insights for.

5. Select the **Well-architected status** tab.

The following elements are displayed:

- **Automatic daily analysis timestamp:** Shows when the last scan was performed for this environment.
- **Configurations:** Organizes findings by configuration area and displays their status and details.

View well-architected assessments

Cluster node management

This assesses whether your EVS cluster nodes have appropriate EC2 stop and termination protection configured.

Status:

- **Optimized:** All EVS nodes have both EC2 stop protection and termination protection configured.
- **Not optimized:** At least one EVS node does not have EC2 stop protection or termination protection configured.

Why this matters:

EVS ESXi nodes should be managed exclusively using vCenter or other VMware-level management tools. Without proper EC2-level protections, nodes could be accidentally stopped or terminated from the EC2 console, which can lead to virtual machine data unavailability or data loss.

To view detailed findings:

1. In the Well-architected status tab, locate **Cluster node management**.
2. Select **View** to open the findings dialog.

The dialog displays:

- **Findings summary:** A detailed explanation of the issue discovered in your environment.
- **Resource grid:** A table showing all EVS nodes and their protection status, including:
 - Node identifier
 - EC2 stop protection status
 - EC2 termination protection status
- **Action required:** Step-by-step issue resolution procedures.
- **Recommendation:** Best practice guidance.

Remediation:

To remediate this issue, enable stop and termination protection for your EVS nodes:

- Follow the procedure specified in [AWS documentation for enabling stop protection](#).
- Follow the procedure specified in [AWS documentation for enabling termination protection](#).

EVS environment resiliency

This assesses whether your EVS cluster nodes are properly distributed across partition placement groups.

Status:

- **Optimized:** All nodes are members of a single partition placement group configured with four or more partitions.
- **Not optimized** if any of the following is true:
 - Nodes are members of more than one placement group.
 - At least one node is a member of a non-partitioned placement group.
 - All nodes are members of a partitioned placement group with fewer than four partitions.

Why this matters:

Proper partition placement ensures that your EVS cluster nodes are distributed across multiple fault-isolated hardware partitions within an AWS availability zone. Misalignment can result in significant loss of processing power or downtime if a partition fails.

To view detailed findings:

1. In the Well-architected status tab, locate **EVS environment resiliency**.
2. Select **View** to open the findings dialog.

The dialog displays:

- **Findings summary:** A detailed explanation of the partitioning misalignment.
- **Resource grid:** A table showing EVS environment nodes with:
 - Node identifier
 - Placement group name
 - Placement group type
 - Placement group partitions count
- **Action required:** Step-by-step remediation procedures
- **Recommendation:** Best practice guidance

Remediation:

To remediate partition placement issues:

- When adding new nodes to the EVS environment, provision the new nodes using a partitioned placement group with at least four partitions.
- If cluster nodes are being replaced, ensure that the replacement nodes are provisioned using a partitioned placement group with at least four partitions.
- Try to consolidate all EVS nodes to a single placement group aligned with the above recommendations.

Best practice recommendation:

When creating or expanding an EVS environment, provision all cluster nodes using a single partitioned placement group configured with four partitions or higher.

What's next

After reviewing your well-architected insights and implementing recommended changes:

- Monitor the well-architected status tab daily to stay informed about your environment's status.
- Follow the remediation procedures for any "Not optimized" findings.
- Review AWS and NetApp documentation for additional best practices.
- Consider implementing the recommendations before expanding your EVS environment.

Related links

- [Create a deployment plan for Amazon EVS using the migration advisor](#)
- [Deploy the recommended FSx for ONTAP file system](#)
- [AWS placement groups documentation](#)

Copyright information

Copyright © 2026 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.