



Migrate to VMware Cloud on AWS

Workload Factory for NetApp Console

NetApp
October 07, 2026

This PDF was generated from <https://docs.netapp.com/us-en/data-services-workloads/fsx-vmware/launch-onboarding-advisor.html> on October 07, 2026. Always check docs.netapp.com for the latest.

Table of Contents

- Migrate to VMware Cloud on AWS 1
 - Create a deployment plan for VMware Cloud on AWS using the VMware workloads migration advisor. 1
 - Create a deployment plan based on an on-premises vSphere environment 1
 - Deploy the recommended FSx for ONTAP file system 4
 - Connect your FSx for ONTAP file systems to VMware Cloud on AWS 5
 - Migrate your data to the new infrastructure with NetApp Workload Factory for VMware 6

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.
 - 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.

- c. **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.

- 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.

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.

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.