



Migrate VMs to Amazon EC2 using FSxN

NetApp Solutions

NetApp
June 18, 2024

Table of Contents

Migrate VMs to Amazon EC2 using FSxN 1

Migrate VMs to Amazon EC2 using FSxN

Migrate VMs to Amazon EC2 using FSxN: Overview

Organizations are accelerating their migrations to cloud computing solutions on AWS, taking advantage of services such as Amazon Elastic Compute Cloud (Amazon EC2) instances and Amazon FSx for NetApp ONTAP (FSx for ONTAP) to modernize their IT infrastructures, achieve cost savings, and improve operational efficiency. These AWS offerings enable migrations that optimize total cost of ownership (TCO) through consumption-based pricing models, enterprise storage features, providing the flexibility and scalability to meet evolving global business demands.

Overview

For enterprises deeply invested in VMware vSphere, migrating to AWS is a cost-effective option given the current market conditions, one that presents a unique opportunity.

As these organizations transition to AWS, they seek to capitalize on the cloud's agility and cost benefits while preserving familiar feature sets, particularly when it comes to storage. Maintaining seamless operations with familiar storage protocols—especially iSCSI—processes, tools, and skillsets is crucial when migrating workloads or setting up disaster recovery solutions.

Using the AWS managed storage service FSx for ONTAP for retaining the enterprise storage capabilities, that too coming from any third-party vendor storage from on-premises, enterprises can unlock the power of AWS while minimizing disruption and maximizing their future investments.

This technical report covers how to migrate on-premises VMware vSphere VMs to an Amazon EC2 instance with data disks placed on FSx for ONTAP iSCSI LUNs using the MigrateOps “data-mobility-as-code” functionality of Cirrus Migrate Cloud (CMC).

Solution requirements

There are a number of challenges that VMware customers are currently looking to solve. These organizations want to:

1. Leverage enterprise storage capabilities, such as thin provisioning, storage efficiency technologies, zero footprint clones, integrated backups, block-level replication, and tiering. This helps optimize migration efforts and future proof deployment on AWS from Day 1.
2. Optimize storage deployments currently on AWS that use Amazon EC2 instances by incorporating FSx for ONTAP and the cost-optimizing features it provides.
3. Reduce the total cost of ownership (TCO) of using Amazon EC2 instances with block storage solutions by rightsizing Amazon EC2 instances to meet the required IOPS and throughput parameters. With block storage, Amazon EC2 disk operations have a cap on bandwidth and I/O rates. File storage with FSx for ONTAP uses network bandwidth. In other words, FSx for ONTAP has no VM-level I/O limits.

Technical components overview

FSx for ONTAP concepts

Amazon FSx for NetApp ONTAP is a fully managed AWS storage service that provides NetApp® ONTAP® file

systems with all the familiar ONTAP data management features, performance, and APIs on AWS. Its high-performance storage supports multiple protocols (NFS, SMB, iSCSI), providing a single service for workloads using Windows, Linux, and macOS EC2 instances.

Since FSx for ONTAP is an ONTAP file system, it brings a host of familiar NetApp features and services with it, including SnapMirror® data replication technology, thin clones, and NetApp Snapshot™ copies. By leveraging a low-cost capacity tier via data tiering, FSx for ONTAP is elastic and can reach a virtually unlimited scale. Plus, with signature NetApp storage efficiency technology, it reduces storage costs on AWS even further. For more, see [Getting started with Amazon FSx for ONTAP](#).

File System

The central resource of FSx for ONTAP is its file system based on solid-state drive (SSD) storage. When provisioning an FSx for ONTAP file system, the user inputs a desired throughput and storage capacity, and selects an Amazon VPC where the file system will reside.

Users also have a choice between two built-in high-availability deployment models for the file system: Multi-Availability Zone (AZ) or single-AZ deployment. Each of these options offers its own level of durability and availability, which customers can select depending on their use case's business continuity requirements. Multi-AZ deployments consist of dual nodes that replicate seamlessly across two AZs. The more cost-optimized single-AZ deployment option structures the file system in two nodes split between two separate fault domains that both reside within a single AZ.

Storage Virtual Machines

Data in the FSx for ONTAP file system is accessed through a logical storage partition which is called a storage virtual machine (SVM). An SVM is actually its own file server equipped with its own data and admin access points. When accessing iSCSI LUNs on an FSx for ONTAP file system, the Amazon EC2 instance interfaces directly with the SVM using the SVM's iSCSI endpoint IP address.

While maintaining a single SVM in a cluster is possible, the option of running multiple SVMs in a cluster has a wide range of uses and benefits. Customers can determine the optimal number of SVMs to configure by considering their business needs, including their requirements for workload isolation.

Volumes

Data within an FSx for ONTAP SVM is stored and organized in structures known as volumes, which act as virtual containers. An individual volume can be configured with a single or multiple LUNs. The data stored in each volume consumes storage capacity in the file system. However, since FSx for ONTAP thinly provisions the volume, the volume only takes up storage capacity for the amount of data being stored.

The Cirrus Migrate Cloud MigrateOps concept

CMC is a transactable software-as-a-service (SaaS) offering from Cirrus Data Solutions, Inc. which is available via the AWS Marketplace. MigrateOps is a Data-Mobility-as-Code automation feature of CMC that allows you to declaratively manage your data mobility operations at scale using simple operation configurations in YAML. A MigrateOps configuration determines how you want your data mobility tasks to be executed. To learn more about MigrateOps, see [About MigrateOps](#).

MigrateOps takes an automation-first approach, which is purpose-built to streamline the entire process, ensuring cloud-scale enterprise data mobility without operational disruptions. In addition to the already feature-rich functionalities that CMC offers for automation, MigrateOps further adds other automations that are often managed externally, such as:

- OS remediation
- Application cutover and approval scheduling

- Zero-downtime cluster migration
- Public/Private cloud platform integration
- Virtualization platform integration
- Enterprise storage management integration
- SAN (iSCSI) configuration

With the above tasks fully automated, all the tedious steps in preparing the on-prem source VM (such as adding AWS agents and tools), creation of destination FSx LUNs, setting up iSCSI and Multipath/MPIO at the AWS destination instance, and all the tasks of stopping/starting application services are eliminated by simply specifying parameters in a YAML file.

FSx for ONTAP is used to provide the data LUNs and rightsize the Amazon EC2 instance type, while providing all the features that organizations previously had in their on-premises environments. The MigrateOps feature of CMC will be used to automate all the steps involved, including provisioning mapped iSCSI LUNs, turning this into a predictable, declarative operation.

Note: CMC requires a very thin agent to be installed on the source and destination virtual machine instances to ensure secure data transfer from the storage source storage to FSx for ONTAP.

Benefits of using Amazon FSx for NetApp ONTAP with EC2 instances

FSx for ONTAP storage for Amazon EC2 instances provides several benefits:

- High throughput and low latency storage that provide consistent high performance for the most demanding workloads
- Intelligent NVMe caching improves performance
- Adjustable capacity, throughput, and IOPs can be changed on the fly and quickly adapt to changing storage demands
- Block-based data replication from on-premises ONTAP storage to AWS
- Multi-protocol accessibility, including for iSCSI, which is widely used in on-premises VMware deployments
- NetApp Snapshot™ technology and DR orchestrated by SnapMirror prevent data loss and speed up recovery
- Storage efficiency features that reduce storage footprint and costs, including thin provisioning, data deduplication, compression, and compaction
- Efficient replication reduces the time it takes to create backups from hours to just minutes, optimizing RTO
- Granular options for file back up and restores using NetApp SnapCenter®

Deploying Amazon EC2 instances with FSx ONTAP as the iSCSI-based storage layer delivers high performance, mission-critical data management features, and cost-reducing storage efficiency features that can transform your deployment on AWS.

Running a Flash Cache, multiple iSCSI sessions, and leveraging a working set size of 5%, it's possible for FSx for ONTAP to deliver IOPS of ~350K, providing performance levels to meet even the most intensive workloads.

Since only network bandwidth limits are applied against FSx for ONTAP, not block storage bandwidth limits, users can leverage small Amazon EC2 instance types while achieving the same performance rates as much larger instance types. Using such small instance types also keeps compute costs low, optimizing TCO.

The ability of FSx for ONTAP to serve multiple protocols is another advantage, one that helps standardize a

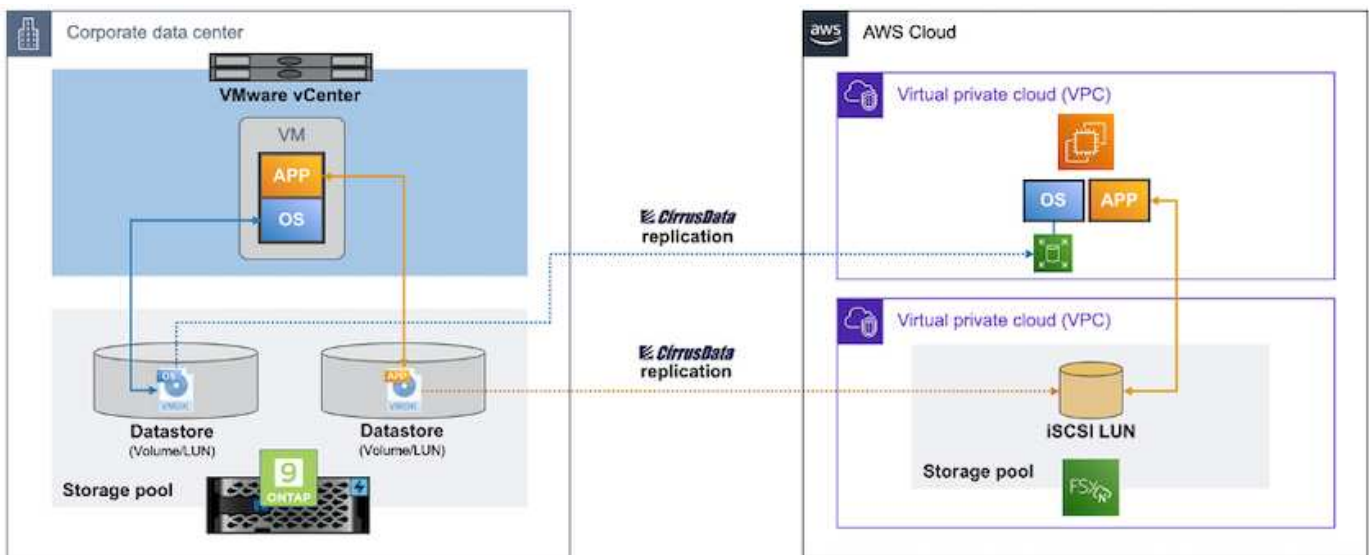
single AWS storage service for a wide range of existing data and file services requirements. For enterprises deeply invested in VMware vSphere, migrating to AWS is a cost-effective option given the current market conditions, one that presents a unique opportunity.

Migrate VMs to Amazon EC2 using FSx: Architecture and Pre-Requisites

This article shows the high-level architecture and deployment pre-requisites for completing the migration.

High level architecture

The diagram below illustrates the high-level architecture of migrating Virtual Machine Disk (VMDK) data on VMware to AWS using CMC MigrateOps:



How to migrate your VMware VMs to AWS using Amazon EC2 and FSx for ONTAP iSCSI

Prerequisites

Before starting the walkthrough steps, make sure the following prerequisites are met:

On AWS

- An AWS account. This includes permissions for subnets, VPC setup, routing tables, security rule migration, security groups, and other requirements for networking such as load balancing. As with any migration, the most effort and consideration should go into networking.
- Appropriate IAM roles that allow you to provision both FSx for ONTAP and Amazon EC2 instances.
- Route tables and security groups are allowed to communicate with FSx for ONTAP.
- Add an inbound rule to the appropriate security group (see below for more details) to allow for secure data transfer from your on-premises data center to AWS.
- A valid DNS that can resolve public internet domain names.
- Check that your DNS resolution is functional and allows you to resolve host names.
- For optimal performance and rightsizing, use performance data from your source environment to rightsize

your FSx for ONTAP storage.

- Each MigrateOps session uses one EIP, hence the quota for EIP should be increased for more parallelism. Keep in mind, the default EIP quota is 5.
- (If Active Directory-based workloads are being migrated) A Windows Active Directory domain on Amazon EC2.

For Cirrus Migrate Cloud

- A Cirrus Data Cloud account at cloud.cirrusdata.com must be created before using CMC. Outbound communication with the CDN, Cirrus Data endpoints, and software repository via HTTPS must be allowed.
- Allow communication (outbound) with Cirrus Data Cloud services via HTTPS protocol (Port 443).
- For a host to be managed by the CMC project, the deployed CMC software must initiate a one-way outbound TCP connection to Cirrus Data Cloud.
- Allow TCP protocol, Port 443 access to portal-gateway.cloud.cirrusdata.com which is currently at 208.67.222.222.
- Allow HTTP POST requests (via HTTPS connection) with binary data payload (application/octet-stream). This is similar to a file upload.
- Ensure that portal-gateway.cloud.cirrusdata.com is resolvable by your DNS (or via OS host file).
- If you have strict rules for prohibiting product instances to make outbound connections, the “Management Relay” feature of CMC can be used where the outbound 443 connection is from a single, secured non-production host.

Note: No storage data is ever sent to the Cirrus Data Cloud endpoint. Only management metadata is sent, and this can be optionally masked so that no real host name, volume name, network IP are included.

For migrating data from on-premises storage repositories to AWS, MigrateOps automates the management of a Host-to-Host (H2H) connection. These are optimized, one-way, TCP-based network connections that CMC uses to facilitate remote migration. This process features always-on compression and encryption that can reduce the amount of traffic by up to eight times, depending on the nature of the data.

Note: CMC is designed so that no production data / I/O leaves the production network during the entire migration phase. As a result, direct connectivity between the source and destination host is required.


Migrate VMs to Amazon EC2 using FSxN: Deployment Guide

This article describes the deployment procedure for this migration solutions.

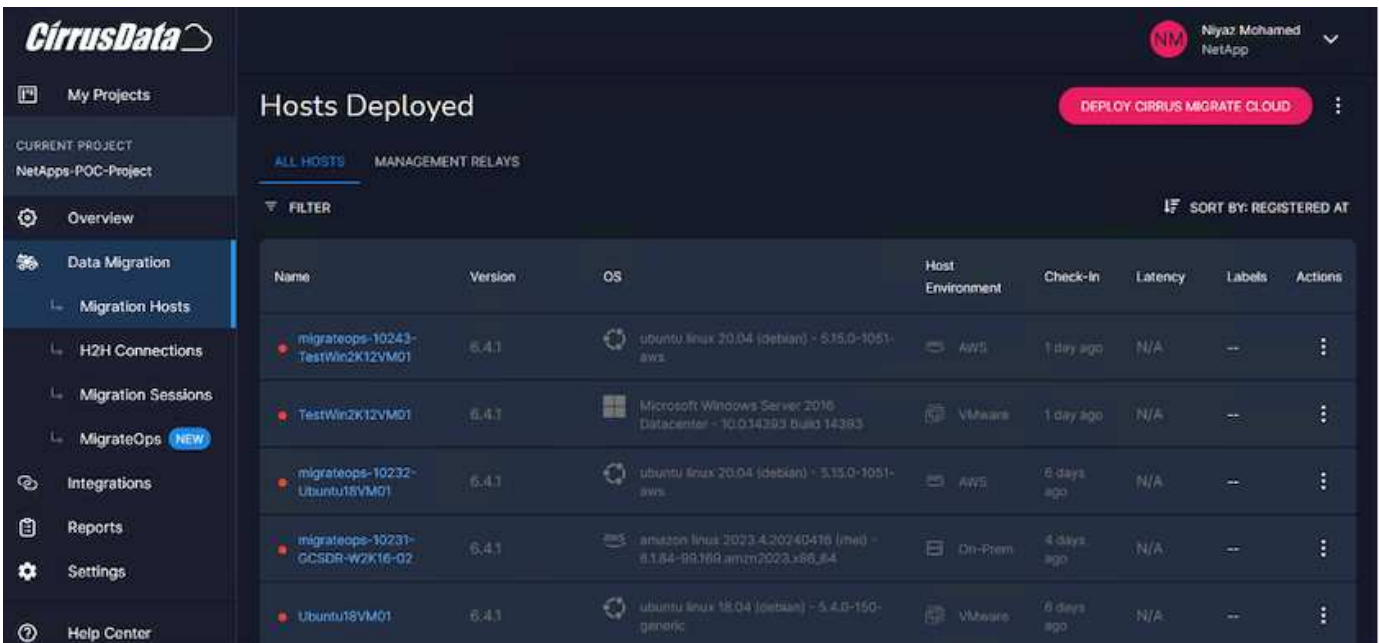
Configure FSx for ONTAP and Cirrus Data for migration operations

This [step-by-step deployment guide](#) shows how to add FSx for ONTAP volume to a VPC. Since these steps are sequential in nature, make sure they are covered in order.

For the purposes of this demonstration, “DRaaS Demo” is the name of the file system created.



Once your AWS VPC is configured and FSx for ONTAP is provisioned based on your performance requirements, log in to cloud.cirrusdata.com and [create a new project](#) or access an existing project.



Before creating the recipe for MigrationOps, AWS Cloud should be added as an integration. CMC provides built-in integration with FSx for ONTAP and AWS. The integration for FSx for ONTAP provides the following automated functionalities:

Prepare your FSx for ONTAP file system:

- Create new volumes and LUNs that match the source volumes

Note: A destination disk in the FSx for ONTAP FS model is a “LUN” that is created on a “Volume” that has enough capacity to contain the LUN plus a reasonable amount of overhead for facilitating snapshots and metadata. The CMC automation takes care of all these details to create the appropriate Volume and the LUN with optional user-defined parameters.

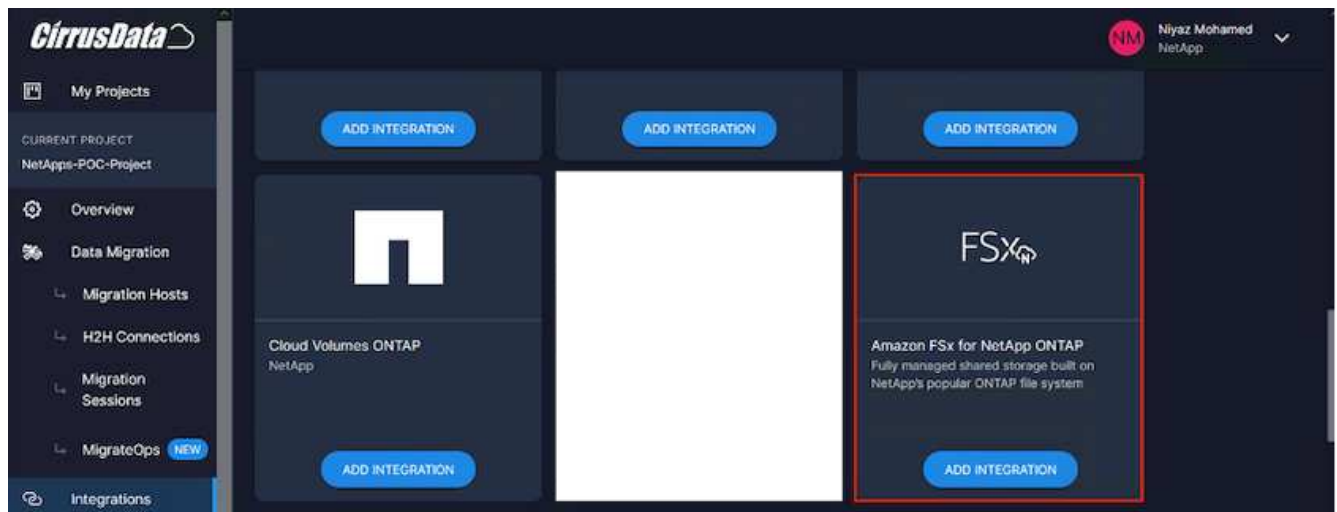
- Create Host entity (called iGroups in FSx) with the Host Initiator IQN
- Map newly created volumes to appropriate host entities using mappings
- Create all other necessary configurations

Prepare Production Host for iSCSI connection:

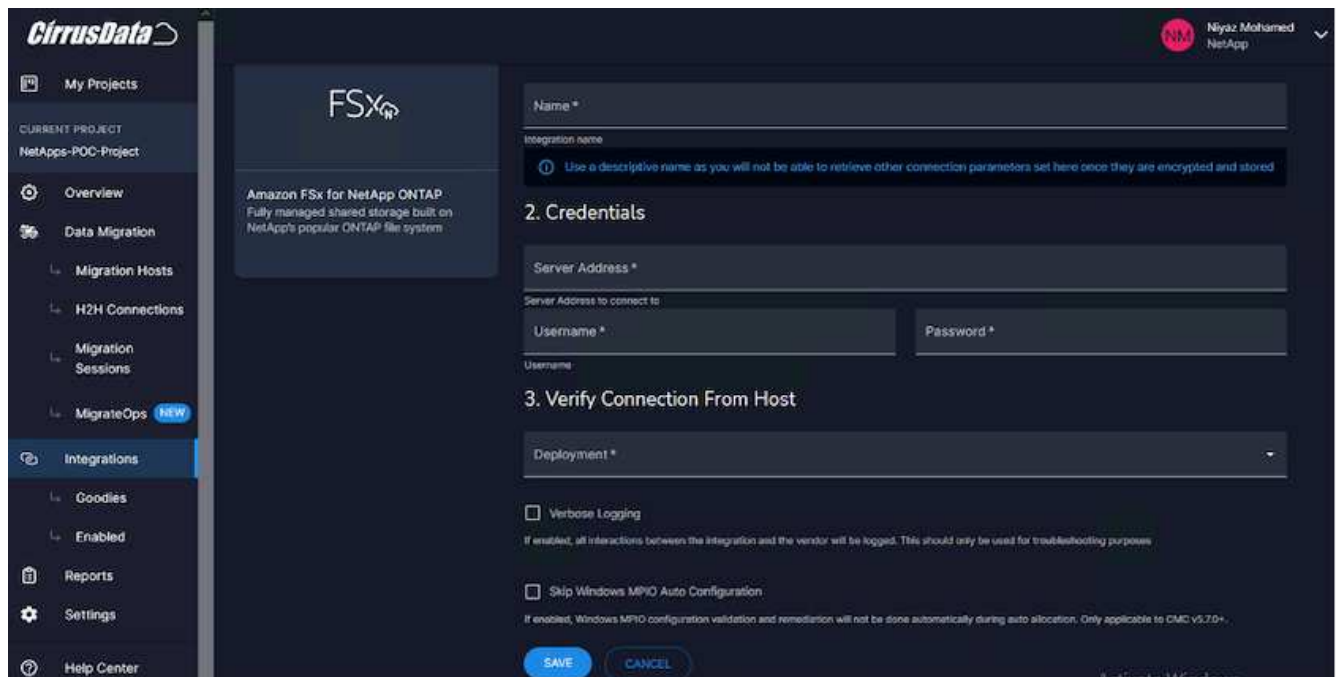
- If necessary, install and configure iSCSI feature and set up Initiator.
- If necessary, install and configure multipath (MPIO for Windows) with proper vendor identifiers.
- Adjust system settings, if necessary, according to vendor best practices, e.g. with udev settings on Linux.
- Create and manage iSCSI connections such as persistent/favorite iSCSI targets on Windows.

To configure CMC Integration for FSx for ONTAP and AWS, perform the following steps:

1. Log in to the Cirrus Data Cloud portal.
2. Go to the Project for which you want to enable the integration.
3. Navigate to Integrations → Goodies.
4. Scroll to find FSx for NetApp ONTAP and click ADD INTEGRATION.



5. Provide a descriptive name (strictly for display purposes) and add the appropriate credentials.




- Once the integration is created, during the creation of a new migration session, select Auto Allocate Destination Volumes to automatically allocate new volumes on FSx for ONTAP.

Note: New LUNs will be created with the same size as the source volume's size, unless "Migrate to Smaller Volumes" is enabled for the migration.

Note: If a host entity (iGroup) doesn't already exist, a new one will be created. All host iSCSI Initiator IQNs will be added to that new host entity.

Note: If an existing host entity with any of the iSCSI initiators already exists, it will be reused.

- Once done, add the integration for AWS, following the steps on the screen.



Note: This integration is used while migrating virtual machines from on-premises storage to AWS along with FSx for ONTAP integration.

Note: Use management relays to communicate with Cirrus Data Cloud if there is no direct outbound connection for production instances to be migrated.

With Integrations added, it's time to register hosts with the Project. Let's cover this with an example scenario.

Host registration scenario

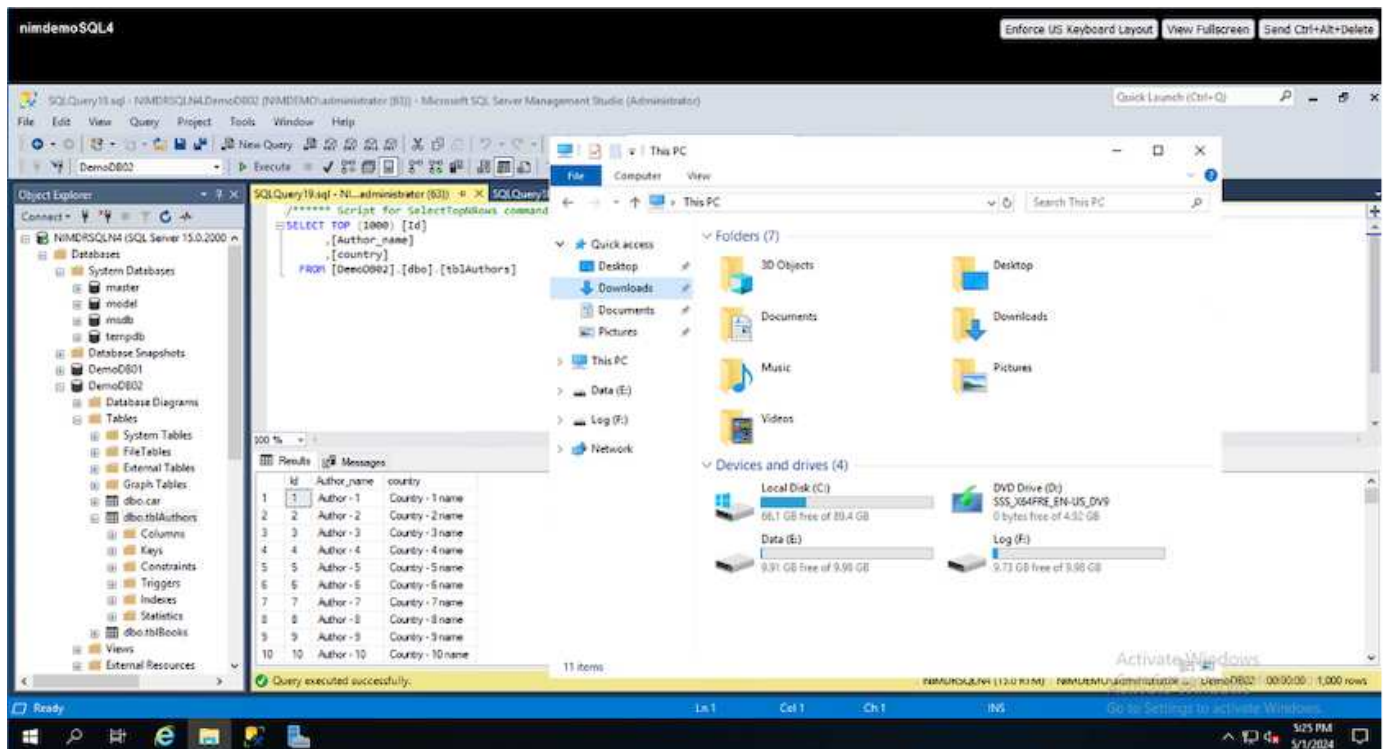
Guest VMware VMs residing on vCenter in on-premises data center:

- Windows 2016 running with SQL Server with three VMDKs including OS and data disks. It is running an active database. The database is located on a data volume backed by two VMDKs.

Note: Since the source is a VMware environment and VMDKs are used, the Windows iSCSI Initiator software is not currently configured on this guest VM. To connect to our destination storage via iSCSI, both iSCSI and MPIO will have to be installed and configured. Cirrus Data Cloud integration will perform this installation automatically during the process.

Note: The Integration configured in the previous section automates the configuration of the new destination storage in creating the new disks, setting up the host entities and their IQNs, and even remediation of the

application VM (host) for iSCSI and multipath configurations.



This demonstration will migrate the application VMDKs from each VM to an automatically provisioned and mapped iSCSI volume from FSx for ONTAP. The OS VMDK in this case will be migrated to an Amazon EBS volume as Amazon EC2 instances support this Amazon EBS only as the boot disk.

Note: The scale factor with this migration approach is the network bandwidth and the pipe connecting on-premises to AWS VPC. Since each VM has 1:1 host session configured, the overall migration performance depends on two factors:



- Network bandwidth
- Target instance type and ENI bandwidth

The migration steps are as follows:

1. Install CMC agent on each host (Windows and Linux) designated for the migration wave. This can be performed by executing a one-line installation command.

To do this, access Data Migration > Migration Hosts > Click on “Deploy Cirrus Migrate Cloud” and click to select “Windows”.

Then, copy the `ixex` command to the host and run it using PowerShell. Once the deployment of the agent is successful, the host will be added to the Project under “Migration hosts”.



2. Prepare the YAML for each virtual machine.

Note: It is a vital step to have a YAML for each VM that specifies the necessary recipe or blueprint for the migration task.

The YAML provides the operation name, notes (description) along with the recipe name as MIGRATEOPS_AWS_COMPUTE, the host name (system_name) and integration name (integration_name) and the source and destination configuration. Custom scripts can be specified as a before and after cutover action.

```
operations:
  - name: Win2016 SQL server to AWS
```

```

notes: Migrate OS to AWS with EBS and Data to FSx for ONTAP
recipe: MIGRATEOPS_AWS_COMPUTE
config:
  system_name: Win2016-123
  integration_name: NimAWSHybrid
  migrateops_aws_compute:
    region: us-west-2
    compute:
      instance_type: t3.medium
      availability_zone: us-west-2b
    network:
      vpc_id: vpc-05596abe79cb653b7
      subnet_id: subnet-070aeb9d6b1b804dd
      security_group_names:
        - default
    destination:
      default_volume_params:
        volume_type: GP2
      iscsi_data_storage:
        integration_name: DemoDRaaS
      default_volume_params:
        netapp:
          qos_policy_name: ""
    migration:
      session_description: Migrate OS to AWS with EBS and
Data to FSx for ONTAP
      qos_level: MODERATE
    cutover:
      stop_applications:
        - os_shell:
            script:
              - stop-service -name 'MSSQLSERVER'
-Force
              - Start-Sleep -Seconds 5
              - Set-Service -Name 'MSSQLSERVER'
-StartupType Disabled
              - write-output "SQL service stopped
and disabled"
        - storage_unmount:
            mountpoint: e
        - storage_unmount:
            mountpoint: f
      after_cutover:
        - os_shell:
            script:

```

```

-Force
mount disks..." > log.txt

E and F for SQL..." >>log.txt
  - storage_unmount:
    mountpoint: e
  - storage_unmount:
    mountpoint: f
  - storage_mount_all: {}
  - os_shell:
    script:
      - write-output "Waiting 60 seconds to
restart SQL Services..." >>log.txt
      - Start-Sleep -Seconds 60
      - stop-service -name 'MSSQLSERVER'

-Force
      - Start-Sleep -Seconds 3
      - write-output "Start SQL Services..."

>>log.txt
      - Set-Service -Name 'MSSQLSERVER'

-StartupType Automatic
      - start-service -name 'MSSQLSERVER'
      - write-output "SQL started" >>log.txt

```

3. Once the YAMLS are in place, create MigrateOps configuration. To do this, go to Data Migration > MigrateOps, click on "Start New Operation" and enter the configuration in valid YAML format.

4. Click "Create operation".


Note: To achieve parallelism, each host needs to have a YAML file specified and configured.

5. Unless the `scheduled_start_time` field is specified in the configuration, the operation will start immediately.


6. The operation will now execute and proceed. From the Cirrus Data Cloud UI, you can monitor the progress with detailed messages. These steps automatically include tasks that are normally done manually, such as performing auto allocation and creating migration sessions.



Note: During the host-to-host migration, an additional security group with a rule allowing Inbound 4996 port will be created, which will allow the required port for communication and it will be automatically deleted once the synchronization is complete.



7. While this migration session is synchronizing, there is a future step in phase 3 (cutover) with the label “Approval Required.” In a MigrateOps recipe, critical tasks (such as migration cutovers) require user approval before they can be executed. Project Operators or Administrators can approve these tasks from the UI. A future approval window can also be created.



8. Once approved, the MigrateOps operation continues with the cutover.
9. After a brief moment, the operation will be completed.



Note: With the help of Cirrus Data cMotion™ technology, the destination storage has been kept up-to-date with all the latest changes. Therefore, after approval is given, this entire final cutover process will take a very short time—less than a minute—to complete.

Post-migration verification

Let's look at the migrated Amazon EC2 instance running the Windows Server OS and the following steps that have completed:

1. Windows SQL Services are now started.

2. The database is back online and is using storage from the iSCSI Multipath device.
3. All new database records added during migration can be found in the newly migrated database.
4. The old storage is now offline.

Note: With just one click to submit the data mobility operation as code, and a click to approve the cutover, the VM has successfully migrated from on-premises VMware to an Amazon EC2 instance using FSx for ONTAP and its iSCSI capabilities.

Note: Due to AWS API limitation, the converted VMs would be shown as “Ubuntu.” This is strictly a display issue and does not affect functionality of the migrated instance. An upcoming release will address this issue.

Note: The migrated Amazon EC2 instances can be accessed using the credentials that were used on the on-premises side.

Migrate VMs to Amazon EC2 using FSxN: Other Possibilities and Conclusion

This article highlights other possibilities for this migration solution as well as concluding the topic.

Other possibilities

The same approach can be extended to migrate VMs using in-guest storage on on-premises VMs. The OS VMDK can be migrated using CMC and the in-guest iSCSI LUNs can be replicated using SnapMirror. The process requires breaking the mirror and attaching the LUN to the newly migrated Amazon EC2 instance, as depicted in the diagram below.



Conclusion

This document has provided a complete walkthrough of using the MigrateOps feature of CMC to migrate data stored in on-premises VMware repositories to AWS using Amazon EC2 instances and FSx for ONTAP.

The following video demonstrates the migration process from start to finish:

[Migrate VMware VMs to Amazon EC2](#)

To check out the GUI and basic Amazon EBS to FSx for ONTAP local migration, please watch this five-minute demo video:



Migrating to any storage in scale with Cirrus Migrate Cloud

Copyright information

Copyright © 2024 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.