QWE876hjk=

Technical Reports

# Electronic health records on Amazon Web Services

## Using FSx for ONTAP

Mike D'Aiuto, Healthcare Solutions Architect, NetApp
Janruary 2023 | TR-4937

## Abstract

This document provides best practices for using Amazon FSx for NetApp® ONTAP® to run electronic health records (EHRs). FSx for ONTAP is a first-party AWS service sold and supported by Amazon Web Services (AWS). It details a complete production and disaster recovery (DR) deployment on AWS with specific performance, data protection, and migration considerations to AWS.

TABLE OF CONTENTS



LIST OF FIGURES



 

 

# Introduction

This document describes the options for managing EHRs in the AWS cloud (from backup, tiering, read-only, DR, production, or all EHR) and supporting workloads in AWS. AWS (a leader in cloud) and NetApp (a leader in data management) have partnered to deliver a unique, native, and fully managed AWS service powered by NetApp ONTAP data management software. Learn how this partnership can help accelerate your journey to the cloud (or hybrid cloud) and reduce or shut down your data center so that you can focus more on patient care.

Why should you read this technical report? Why is healthcare looking to move to the cloud?

- **Streamline software upgrades.** With FSx for ONTAP, you can rapidly deploy software releases and right-size infrastructure without lengthy planning and budgeting processes.
- **Reduce your EHR total cost of ownership (TCO).** FSx for ONTAP eliminates the 12 to 18 months of project planning needed to refresh hardware, a process that commonly occurs every 3 to 5 years.
- **Improve performance, uptime, and data protection for EHR in the cloud.** This solution is broadly accessible, highly available, and performant so that EHR customers can rest assured their most critical data is protected and always on in AWS.

FSx for ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on the popular NetApp ONTAP file system. FSx for ONTAP combines the familiar features, performance, capabilities, and API operations of NetApp file systems with the agility, scalability, and simplicity of a fully managed AWS service.

Why does FSx for ONTAP have the fastest adoption curve of any AWS?

- **ONTAP is mature.** The world's most deployed data management software.
- **FSx for ONTAP is simple and fully managed.** Launch, run, and scale storage by clicking a single button without spending time on upgrading or fixing software or hardware issues
- **Native AWS experience.** Deploy storage through a console or API, leverage IAM permissions, integrate with CloudTrail, CloudWatch, or AWS backup.
- **Operational stability.** The healthcare industry needs storage that just works with the ability to successfully scale to 5M file systems and 15M backups since the release of FSx for ONTAP.
- **Support.** One number to call for support.
- **Security.** Protect your patient records with the combined security portfolio of AWS and ONTAP. For more information, see Security in Amazon FSx for NetApp ONTAP.
- **Consolidate storage.** File and block storage, reduced cost with efficiencies and all the data management capacities available with ONTAP.
- **Disaster recovery (DR).** Cut your DR cost by up to 70%.

For cloud architects that need to migrate from Amazon Elastic Block Store (EBS), block SDS, multiprotocol file, SAN on ONTAP, or any other storage platform to the cloud, FSx for ONTAP is the answer.

## Electronic Health Records

EHR is the core of patient record storage for most healthcare organizations. An EHR environment is typically made up of multiple access protocols for NAS and SAN. Each workload has different requirements for performance, availability, DR, data protection, security, and compliance. EHR works with an integration engine to connect with other healthcare environments, as shown in Figure 1.

**Figure 1) Challenges in healthcare IT.**

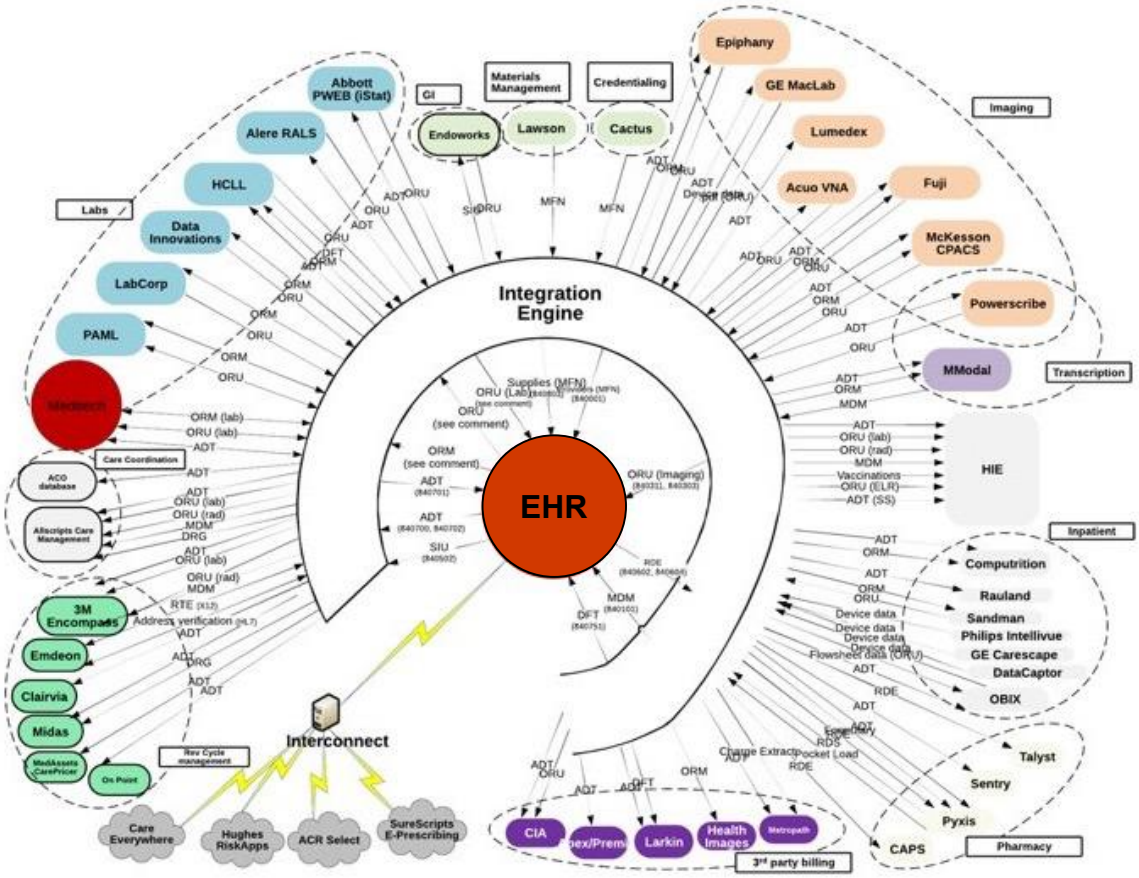


As you can see in this graphic, the challenges with healthcare IT go beyond just EHR—many other critical patient care applications also require attention.

Traditionally, healthcare has built multiple silos for applications, infrastructure, data, and even teams to manage IT. This infrastructure can lead to higher TCO, complex management, and a maintenance renewal nightmare. IT has typically functioned in a reactive firefighter mode, adding little to no strategic value to improve patient care and partnership with healthcare.

Healthcare organizations want to reduce their data center footprint and accelerate their journey to the cloud. The requirements above do not go away for EHR and supporting applications in the cloud. The best way forward is to have a single data management platform for your data, on premises or in the cloud. NetApp and FSx for ONTAP is uniquely aligned to solve this problem.

## NetApp

NetApp ONTAP offers a single data management platform to consolidate all healthcare workloads and a long history of transforming healthcare organizations. ONTAP is the world's #1 most widely deployed data management platform with 30 years of development and trusted by over 4,000 healthcare organizations. The same ONTAP software that transformed EHR on premises for years is also available in FSx for ONTAP, allowing you to unlock your data and start your journey to the cloud.

**Figure 2) Scale and simplicity for healthcare.**

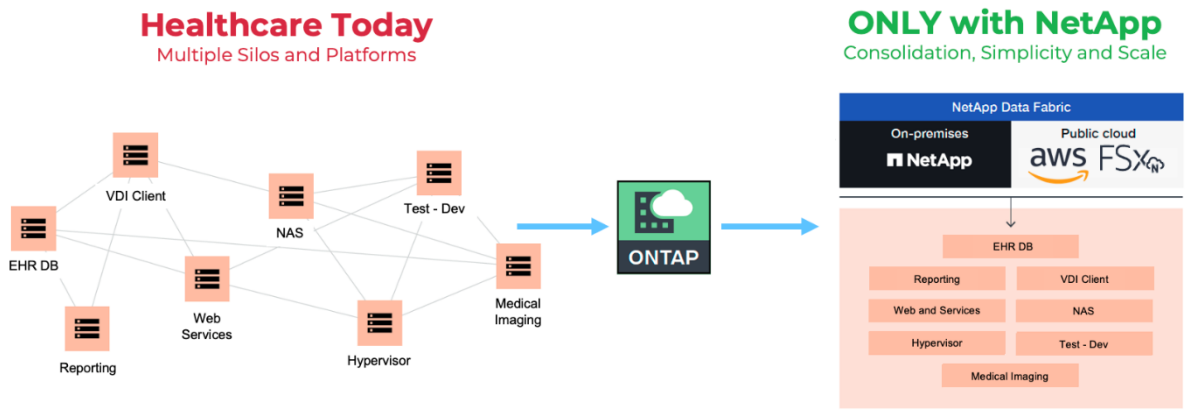


## AWS

NetApp is a cloud-first data management software company that has enjoyed a complementary 10-year partnership with AWS. AWS and NetApp have aligned to help healthcare organizations reduce their data center footprint and accelerate their journey to the cloud. FSx for ONTAP enables the lift and shift of critical enterprise applications such as EHR to the cloud with no refactoring of data. FSx for ONTAP is the catalyst to move applications long considered boat anchors to the cloud.

After your EHRs have been moved to the cloud, hundreds of other applications can be lifted and shifted to the cloud.

**Figure 3) AWS and NetApp partnership.**

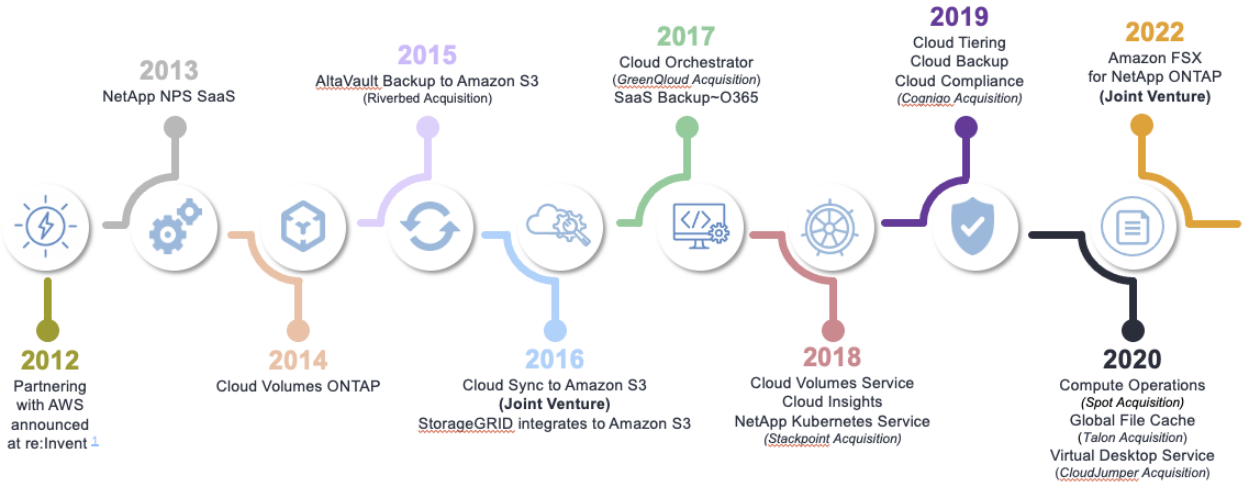


# Cloud challenges

Migrating to the cloud presents the following challenges for critical enterprise applications:

- Limited scale on native AWS and Azure
- Refactoring data and complex migrations
- High TCO

- Low 99.9% SLA
- Lack of data management required to run EHR at scale
- Security and compliance
- Multiple silo protocols and tiers
- Low comfort of independent software vendors (ISVs)

# FSx for ONTAP considerations

When setting up the environment for your data, depending on your preference, you can manage FSx for ONTAP by using the following tools from either NetApp or AWS:

- NetApp management tools:
  - NetApp BlueXP™ hybrid multicloud management
  - ONTAP CLI and API
- AWS tools:
  - AWS Management Console
  - AWS CLI
  - Amazon FSx API and SDK

**Figure 4) Key NetApp portfolio integration.**

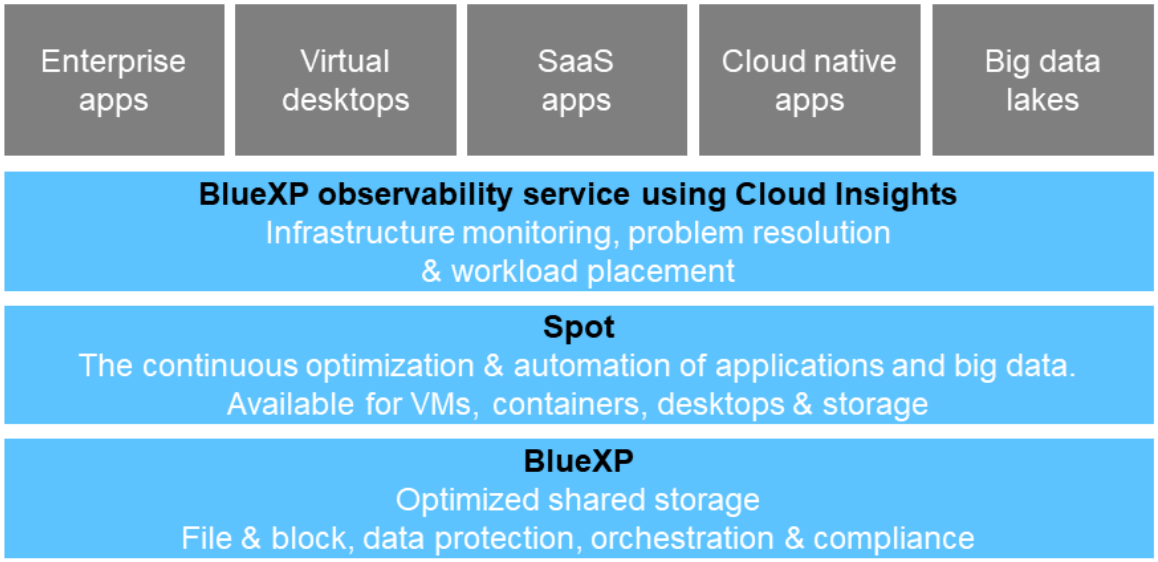


NetApp offers a complete portfolio of solutions to enhance EHR management in AWS. This portfolio includes tools for optimization, cost reduction, monitoring, security, and data protection for the environment. Figure 5 shows how these options integrate with FSx for ONTAP.

**Figure 5) Key NetApp portfolio integration.**

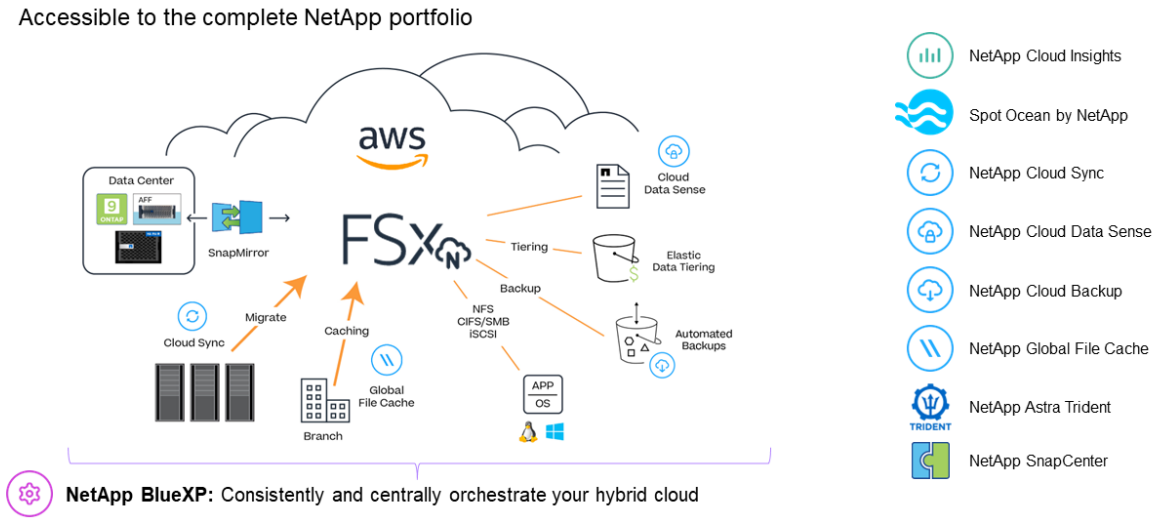


## Availability zones

Each region in AWS is made up of three to six availability zones (AZs). An AZ is made up of one or more data-center buildings with fully redundant power and connectivity. FSx for ONTAP is widely available in most regions and AZs worldwide.

For individuals familiar with NetApp, an FSx for ONTAP file system is equivalent to a fully redundant ONTAP HA pair. FSx for ONTAP can be deployed in a single AZ or across multi AZs. See the following considerations when selecting AZ deployments:

- Multi AZs provide higher availability and can withstand losing an entire AZ.
- Multi AZs require more storage and add to the overall cost.
- Multi AZs have the same read but lower write latency.
- A single AZ is less expensive with lower write latency.

For the majority of EHR workloads, a single AZ is optimal.

Figure 6 illustrates NetApp portfolio integration.

**Figure 6) Key NetApp portfolio integrations.**

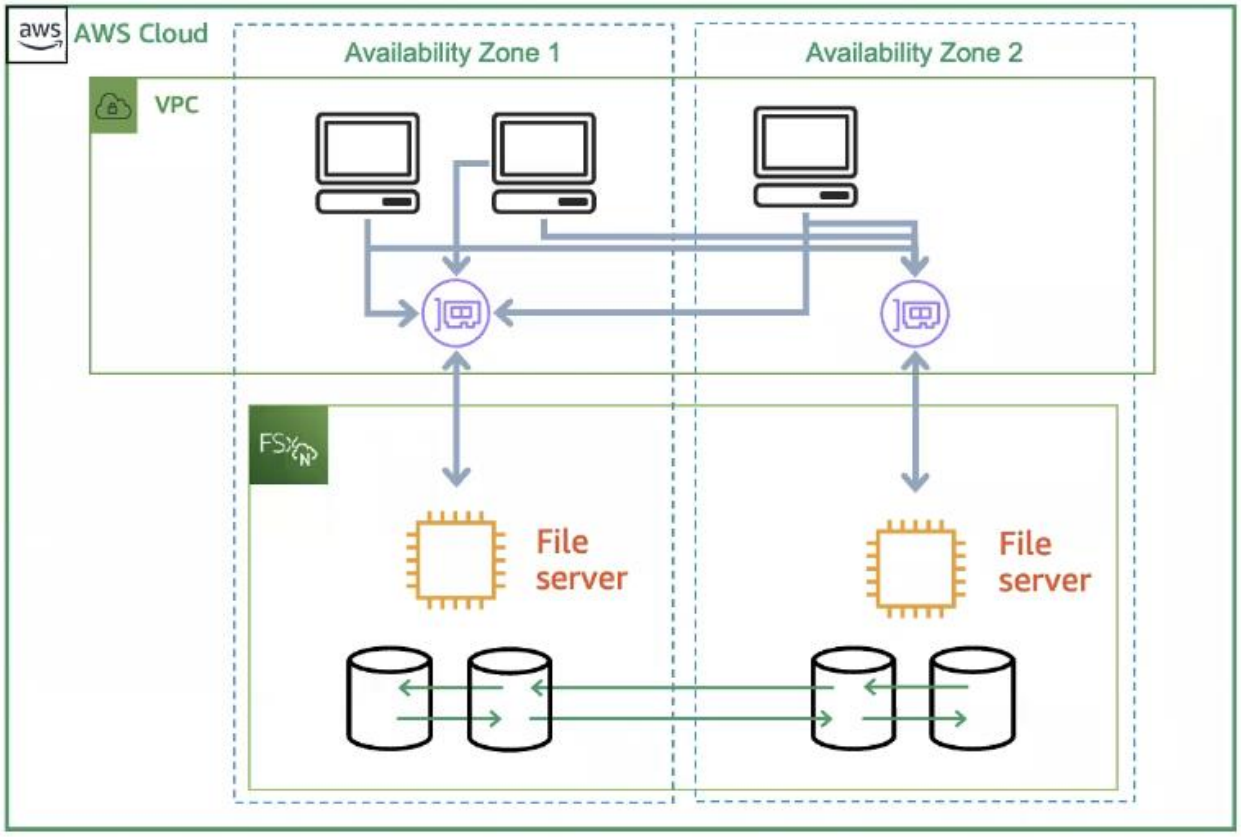


Figure 7 through Figure 9 illustrate EHR architectures.

**Figure 7) EHR hybrid cloud architecture.**

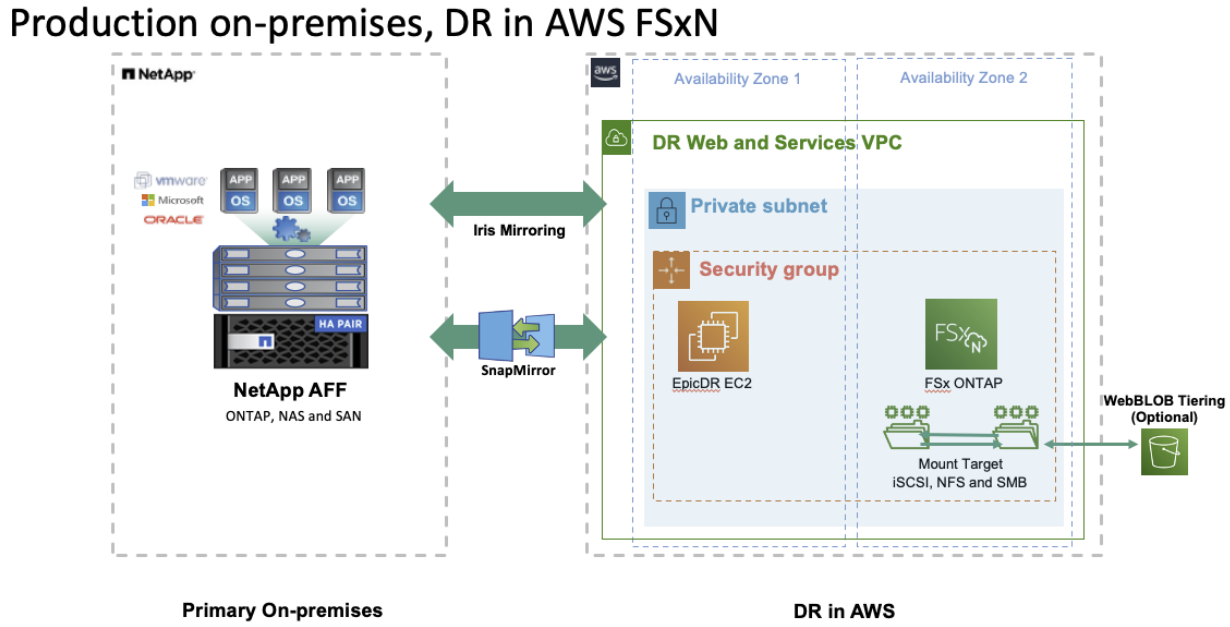

**Figure 8) EHR cloud-first architecture.**

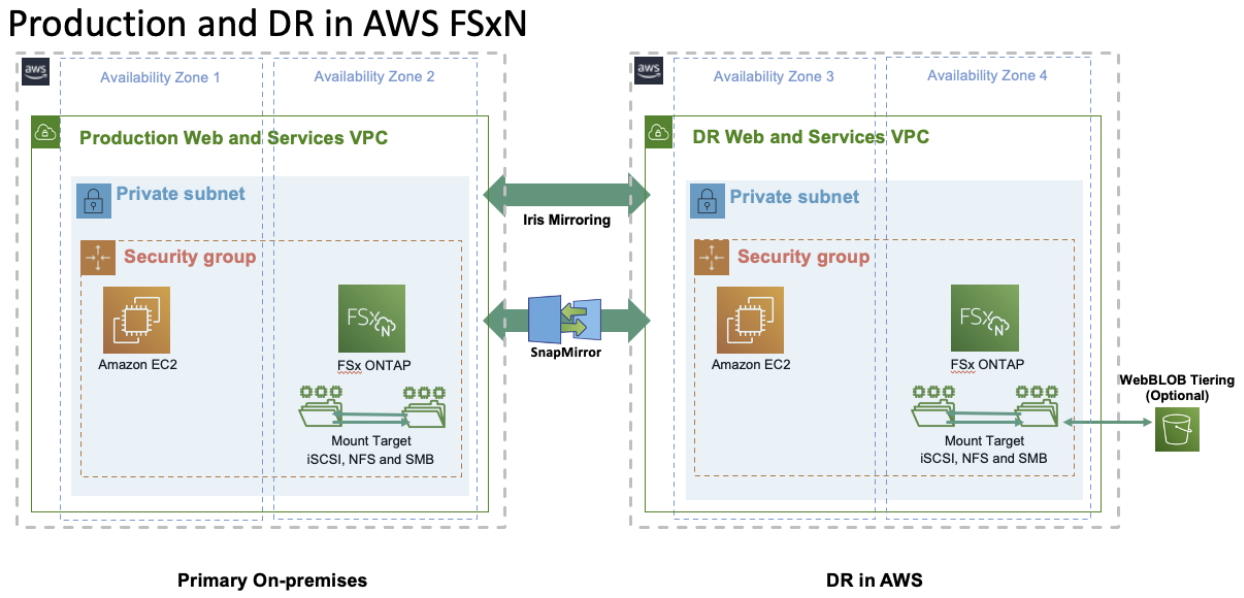


**Figure 9) EHR backup and ransomware architecture.**

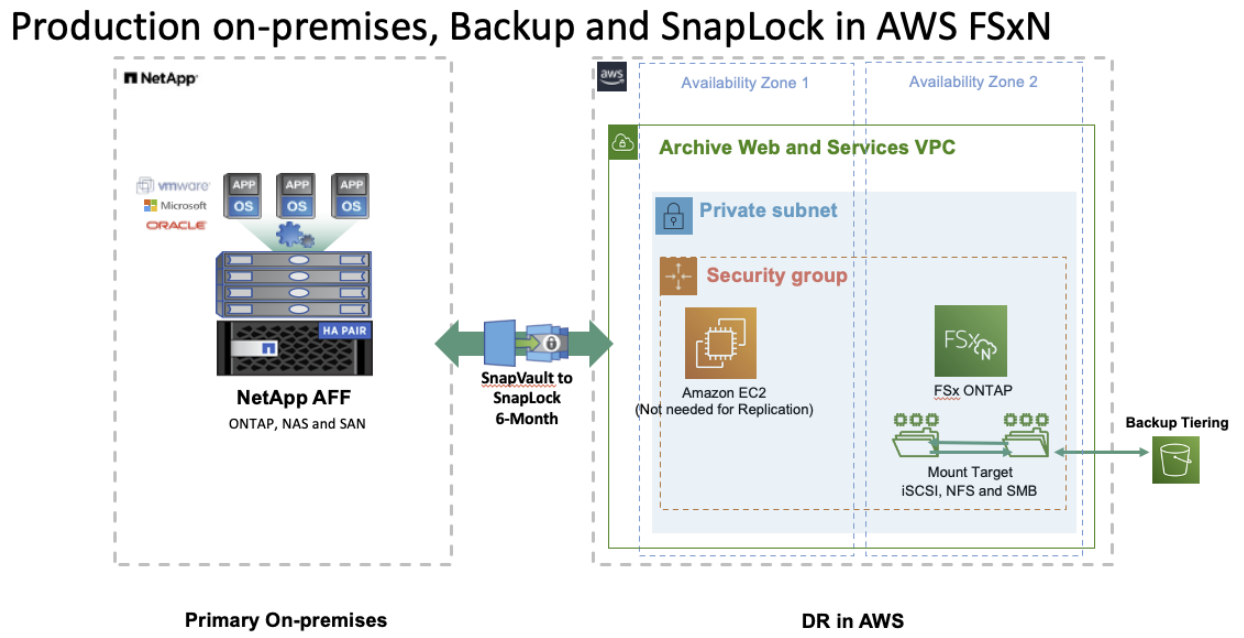


## Performance and tiering

FSx for ONTAP comes with two tiers of storage; you can leverage automatic tiering depending on each of the workloads with file and block. Block data is automatically moved between tiers based on access patterns (hot or cold).

- **Primary tier.** Comes with SSD-based storage (up to 196TB).
- **Capacity pool tier.** Scales elastically to petabytes; cost-effective for less accessed files.

You should match the tiering policy to the workload. Tiering policies include Auto, Snapshot-only, All, and None:

- Use Snapshot-only for high-performance workloads so that the EHR database stays on hot SSD storage and Snapshot copy backups are automatically moved to the capacity pool.
- For workloads such as BLOB that mostly contain cold data, you can configure the system to keep 20% of the data hot on the SSD disk and automatically tier cold data to the capacity pool.
- For FSx for ONTAP as a backup target, use the All policy.

  **Note:** The flexibility to manage workloads and reduce cost is built into FSx for ONTAP.

## Reduce cost

FSx for ONTAP can help reduce storage costs with the following features while adding reliability and data management capabilities:

- Inline efficiency (deduplication, zero block deduplication, compression, and compaction)
- Thin provisioning (EBS disk is thin.)
- Instant Snapshot copies and clones
- Automatic tiering

Although storage costs vary for each workload, if you have a thin-provisioned EBS disk with 100TB of capacity that is only 50% used, you could see the following savings:

- Thin provisioning: 50%
- Inline efficiency: 50%
- Archive data: 50%
- Backup data: 25%

You could end up with only 25TB of data on disk, 6.25TB of data on the SSD, and 18.75TB of archive and backup data on the capacity pool. Use the AWS pricing calculator to review pricing for your workload.

## Migrations

Figure 10 depicts the migration of data from a NetApp ONTAP cluster with NetApp SnapMirror in the corporate data center to the destination AWS FSx instance.

**Figure 10) Migrations to the cloud.**

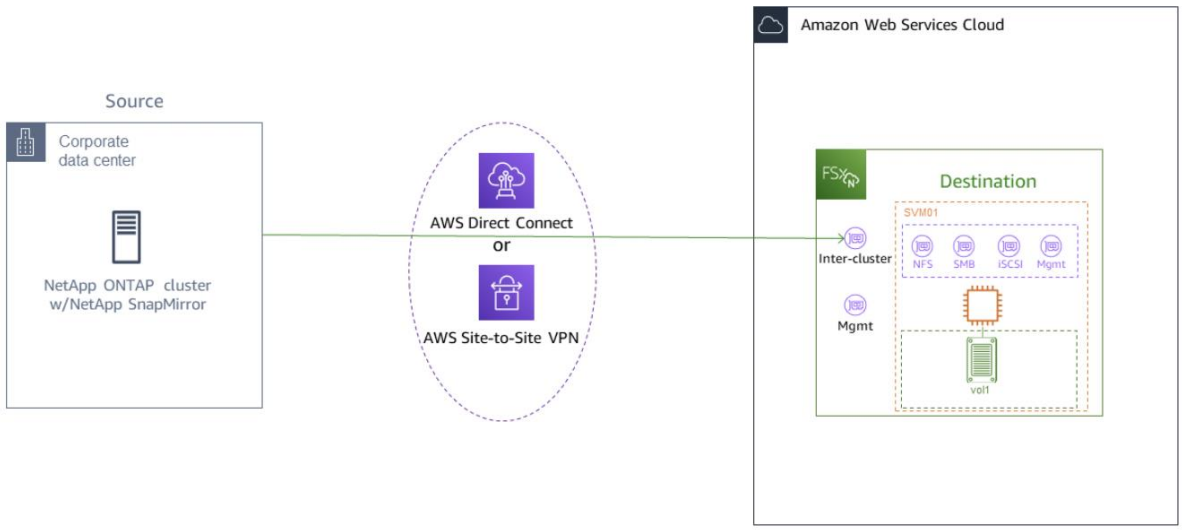

# Architecture

The journey to the cloud doesn't usually include cutting everything over at once. You can start with backups and tiering to the cloud or DR in the cloud before running production workloads in the cloud. This section reviews some deployment options for running EHR in AWS.

## Production on NetApp, DR on FSx for ONTAP

Some of the key advantages of running EHR workloads with ONTAP on premises are the same reasons ONTAP is available in the cloud with FSx for ONTAP:

- Lift and shift NAS and SAN workloads from ONTAP on premises to FSx for ONTAP without refactoring your data.
- Use SnapMirror for migration of data and DR.
- Use the same data management tools and strategy.
- Reduce EHR DR costs by up to 70%:
  - Pay only for what is used—scale up or down.
  - Turn off HER DR when not active to save on compute costs.
- Perform simplified EHR DR testing.

**Figure 11) EHR production on ONTAP.**

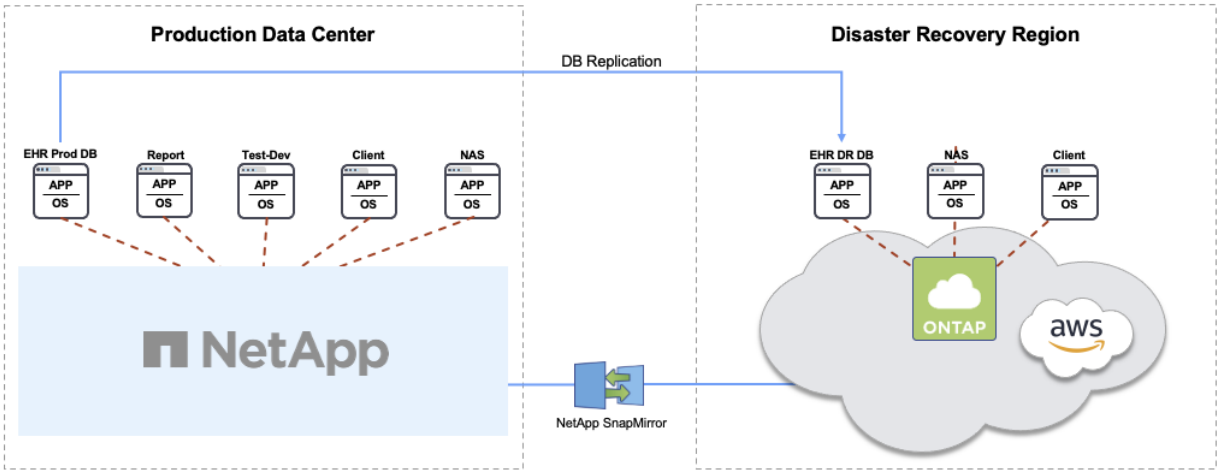


## Production on non-NetApp, DR on FSxN

If you are running other storage on premises, you can use either database replication or NetApp Cloud Sync to migrate and replicate data to the cloud. Regardless of what you run on premises, FSx for ONTAP can be used to run DR in the cloud.

Figure 12 illustrates EHR production on other storage.

**Figure 12) EHR production on other storage.**

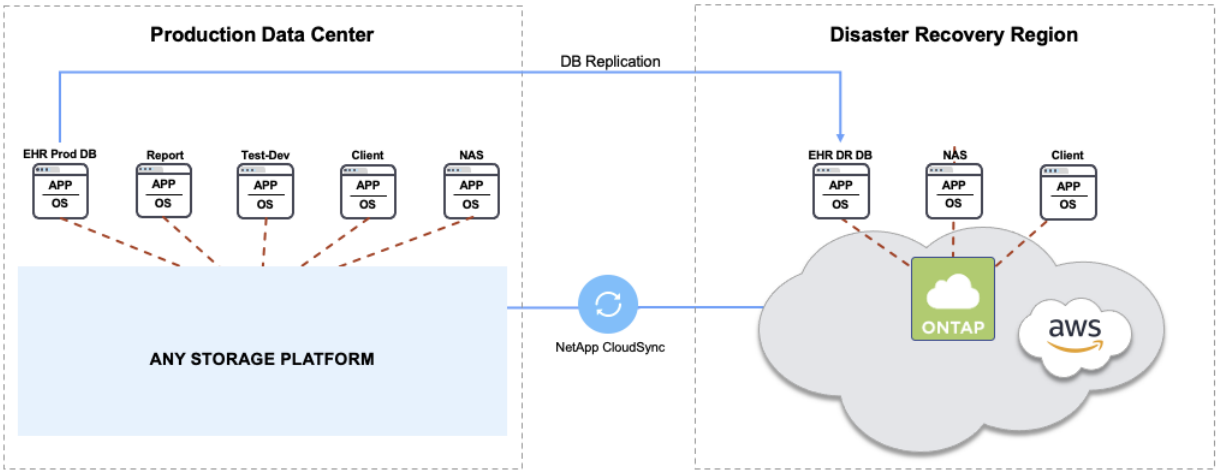


## Production and DR on FSx for ONTAP

**Figure 13) EHR production and DR in AWS.**

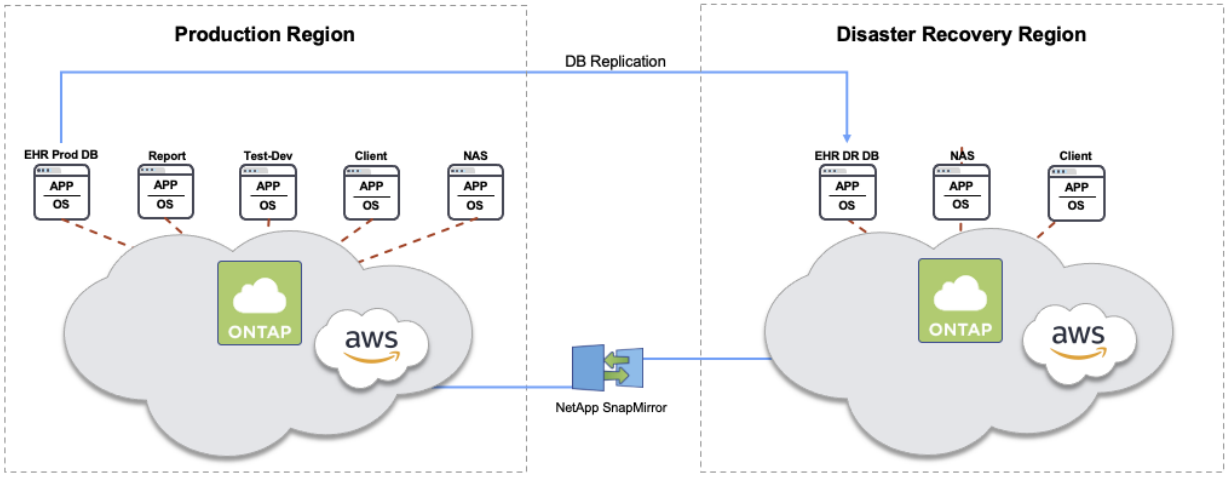


## Consistency groups (One or more FSxN systems)

Beginning in ONTAP 9.11.1, consistency groups support two-phase commits for consistency group Snapshot creation. This feature is only available with the ONTAP REST API. A consistency group should contain all the of volumes for the EHR database within the cluster.

When creating a Snapshot copy of the consistency group, the first Snapshot creation phase (action=start) should be executed on each cluster prior to starting the Snapshot commit phase (action=commit). This allows a database to be split across multiple clusters while maintaining write-order consistency.

For more information, see "Protect a consistency group."

## Storage configuration and layout

### Configuration summary

The storage configuration includes the following components:

- FSx for NetApp ONTAP file system count: 1
  - Volume count: 8
  - Size per volume: 3.75TiB
  - LUN count: 8
  - LUN size: 3.75TiB
- Compute VM count: 1
  - Compute VM SKU: m6in.32xlarge
  - Compute OS: RHEL8
  - Compute EBS disks: 2
    - •OS - 1 gp2 volume (20 GB, default IOPS)
    - •Journal - 1 io2 volume (100GB, 100,000 IOPS)

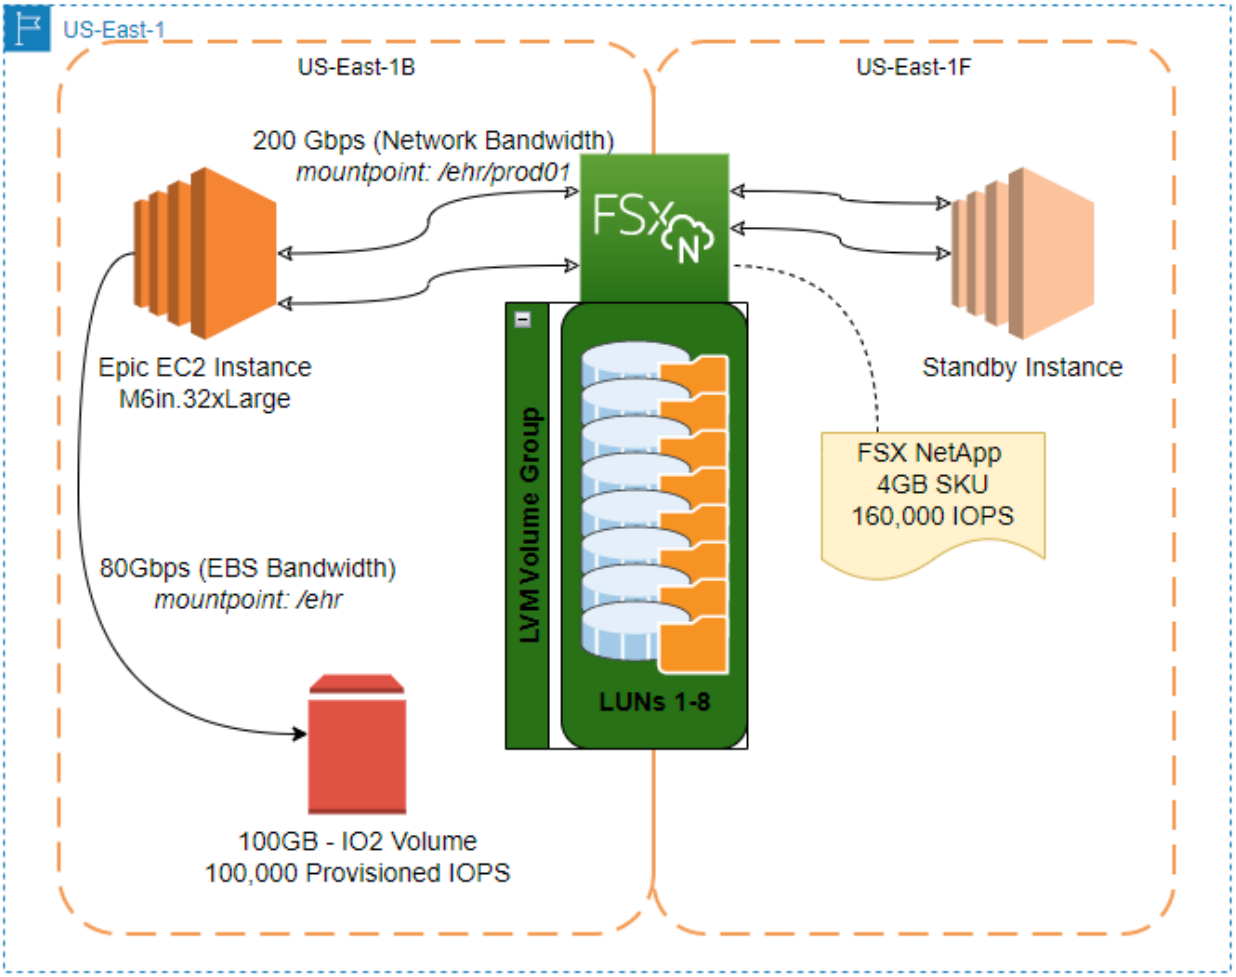


## ONTAP commands

Create eight volumes, eight LUNs (one per volume), configure the igroup, and map the LUNs.

1. Create eight volumes.

```
vol create -volume vol1 -aggregate aggr1 -size 3750g -state online  -snapshot-policy none
vol create -volume vol2 -aggregate aggr1 -size 3750g -state online  -snapshot-policy none
vol create -volume vol3 -aggregate aggr1 -size 3750g -state online  -snapshot-policy none
vol create -volume vol4 -aggregate aggr1 -size 3750g -state online  -snapshot-policy none
```

```
vol create -volume vol5 -aggregate aggr1 -size 3750g -state online  -snapshot-policy none
vol create -volume vol6 -aggregate aggr1 -size 3750g -state online  -snapshot-policy none
vol create -volume vol7 -aggregate aggr1 -size 3750g -state online  -snapshot-policy none
vol create -volume vol8 -aggregate aggr1 -size 3750g -state online  -snapshot-policy none
```

2. Create eight LUNS.

```
lun create -path /vol/vol1/lun1 -size 3750g -ostype linux -space-reserve disabled -space-
allocation disabled -class regular -caching-policy all
lun create -path /vol/vol2/lun2 -size 3750g -ostype linux -space-reserve disabled -space-
allocation disabled -class regular -caching-policy all
lun create -path /vol/vol3/lun3 -size 3750g -ostype linux -space-reserve disabled -space-
allocation disabled -class regular -caching-policy all
lun create -path /vol/vol4/lun4 -size 3750g -ostype linux -space-reserve disabled -space-
allocation disabled -class regular -caching-policy all
lun create -path /vol/vol5/lun5 -size 3750g -ostype linux -space-reserve disabled -space-
allocation disabled -class regular -caching-policy all
lun create -path /vol/vol6/lun6 -size 3750g -ostype linux -space-reserve disabled -space-
allocation disabled -class regular -caching-policy all
lun create -path /vol/vol7/lun7 -size 3750g -ostype linux -space-reserve disabled -space-
allocation disabled -class regular -caching-policy all
lun create -path /vol/vol8/lun8 -size 3750g -ostype linux -space-reserve disabled -space-
allocation disabled -class regular -caching-policy all
```

3. Create an igroup.

```
igroup create -igroup ehr  -protocol iscsi -ostype linux -initiator iqn.1994-
05.com.redhat:f68ae8197b2c
```

4. Map the LUN.

```
lun map -path /vol/vol1/lun1 -igroup ehr -lun-id 1
lun map -path /vol/vol2/lun2 -igroup ehr -lun-id 2
lun map -path /vol/vol3/lun3 -igroup ehr -lun-id 3
lun map -path /vol/vol4/lun4 -igroup ehr -lun-id 4
lun map -path /vol/vol5/lun5 -igroup ehr -lun-id 5
lun map -path /vol/vol6/lun6 -igroup ehr -lun-id 6
lun map -path /vol/vol7/lun7 -igroup ehr -lun-id 7
lun map -path /vol/vol8/lun8 -igroup ehr -lun-id 8
```

## Linux commands

1. Install the required packages.

```
yum install -y device-mapper-multipath lvm2 sysstat iscsi-initiator-utils git
```

2. Update the `sysctl.conf` tunables.

```
sunrpc.tcp_max_slot_table_entries = 128
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.ipv4.tcp_rmem = 4096 131072 16777216
net.ipv4.tcp_wmem = 4096 16384 16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_no_metrics_save=1
net.ipv4.tcp_slow_start_after_idle=0
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_sack = 1
```

3. Configure the io elevator.

```
for i in `ls /sys/block/ | egrep 'nvme|sd'`; do echo "mq-deadline" >
/sys/block/${i}/queue/scheduler; done
```

4. Configuring the io simulator depends on your EHR perform validation with approved I/O tools.

### Preparing iSCSI for multiple connections

To prepare iSCSI for multiple connections, complete the following steps:

1. Edit `/etc/iscsi/iscsid.conf`.
2. Update `node.session.nr_sessions` to equal `8`. This can be tuned depending on the number of LUNs.

```
node.session.nr_sessions=8
```

3. Start iscsid and enable and start multipathd.

```
sudo service iscsid restart
sudo mpathconf --enable
sudo service multipathd reload
```

### iSCSI configuration

To configure iSCSI, complete the following commands:

```
sudo iscsiadm -m discovery -t sendtargets -p 172.31.43.132:3260
sudo iscsiadm -m discovery -t sendtargets -p 172.31.0.222:3260
sudo iscsiadm -m node -T iqn.1992-08.com.netapp:sn.02c0fae707ef11ecbeec1feb5be3f554:vs.3 -p
172.31.43.132:3260 -l
sudo iscsiadm -m node -T iqn.1992-08.com.netapp:sn.02c0fae707ef11ecbeec1feb5be3f554:vs.3 -p
172.31.0.222:3260 -l
```

### Logical Volume Manager (LVM) configuration

1. Identify the disks on the Linux host using lsblk.

```
sudo lsblk -ll
# Following the reference configuration there will be one 100GB nvme device and eight 3.7TB
multipathed devices.
```

2. Configure the journal volume.

```
sudo pvcreate /dev/nvme1p1
sudo vgcreate -s 256K ehrjournalvg /dev/nvme1p1
sudo lvcreate -n ehrjournallv -l 100%FREE -i 1 -I 256K ehrjournalvg
sudo mkfs.xfs -K /dev/mapper/ehrjournalvg-ehrjournallv
sudo mkdir /ehr
sudo echo "/dev/mapper/ehrjournalvg-ehrjournallv /ehr xfs defaults,noatime 0 2" >> /ets/fstab
sudo mount /ehr
```

3. Configure the database volume.

```
sudo pvcreate /dev/mapper/3600a09806c111111111111111a32425[0-7]
sudo vgcreate -s 2M ehrdbvg /dev/mapper/3600a09806c111111111111111a32425[0-7]
sudo lvcreate -n ehrdblv -l 100%FREE -i 1 -I 2M ehrdbvg
sudo mkfs.xfs -K /dev/mapper/ehrdbvg-ehrdblv
sudo mkdir /ehr/prod01
sudo echo "/dev/mapper/ehrdbvg-ehrdblv /ehr/prod01 xfs _netdev,noatime 0 2" >> /ets/fstab
sudo mount /ehr/prod01
```

# Performance results

### Single FSx for ONTAP system performance

A single FSx for ONTAP system was able to produce ~116K IOPS with the total DB time staying under 45 seconds and the average read latency staying under 2ms.

**Figure 14: Single FSx for ONTAP system performance.**

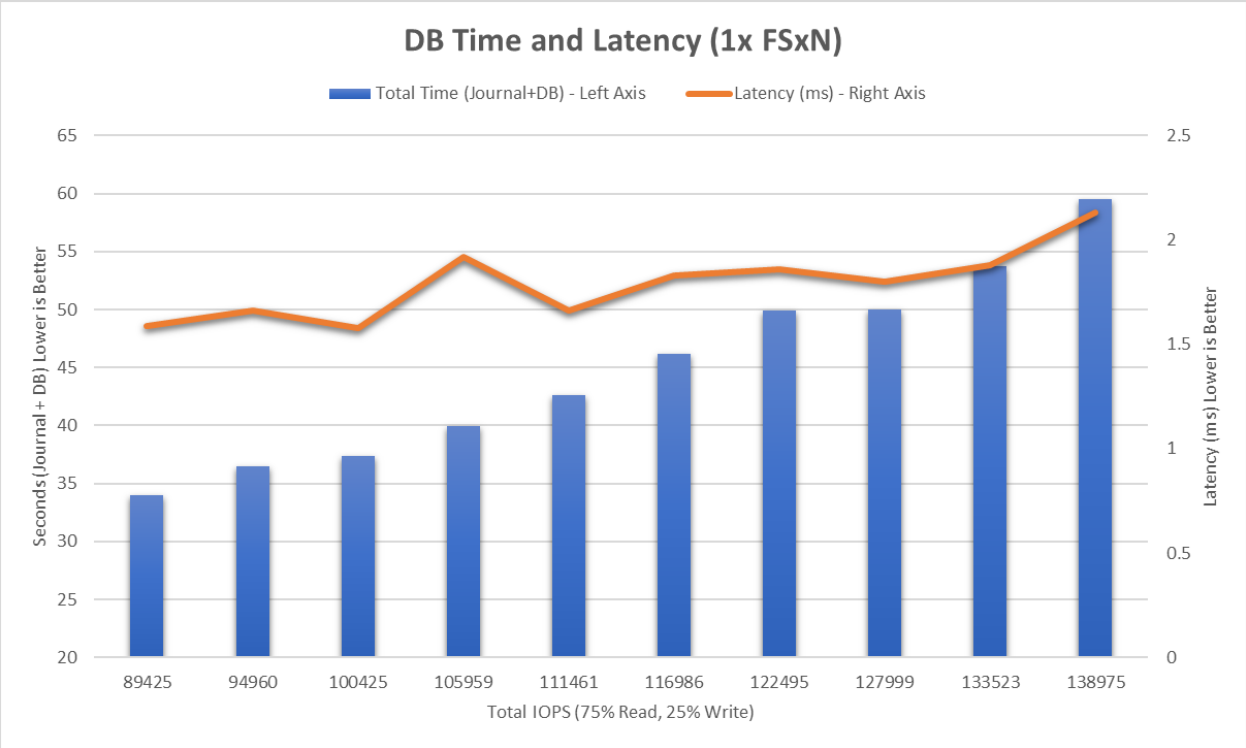


## Two FSx for ONTAP system performance

Two FSx for ONTAP systems with four LUNs each (equal to the capacity of one FSx for ONTAP instance with eight LUNs) were able to produce ~146K IOPS. The total DB time stayed under 45 seconds, and the average read latency stayed under 2ms.

**Figure 15: Two FSx for ONTAP system performance.**

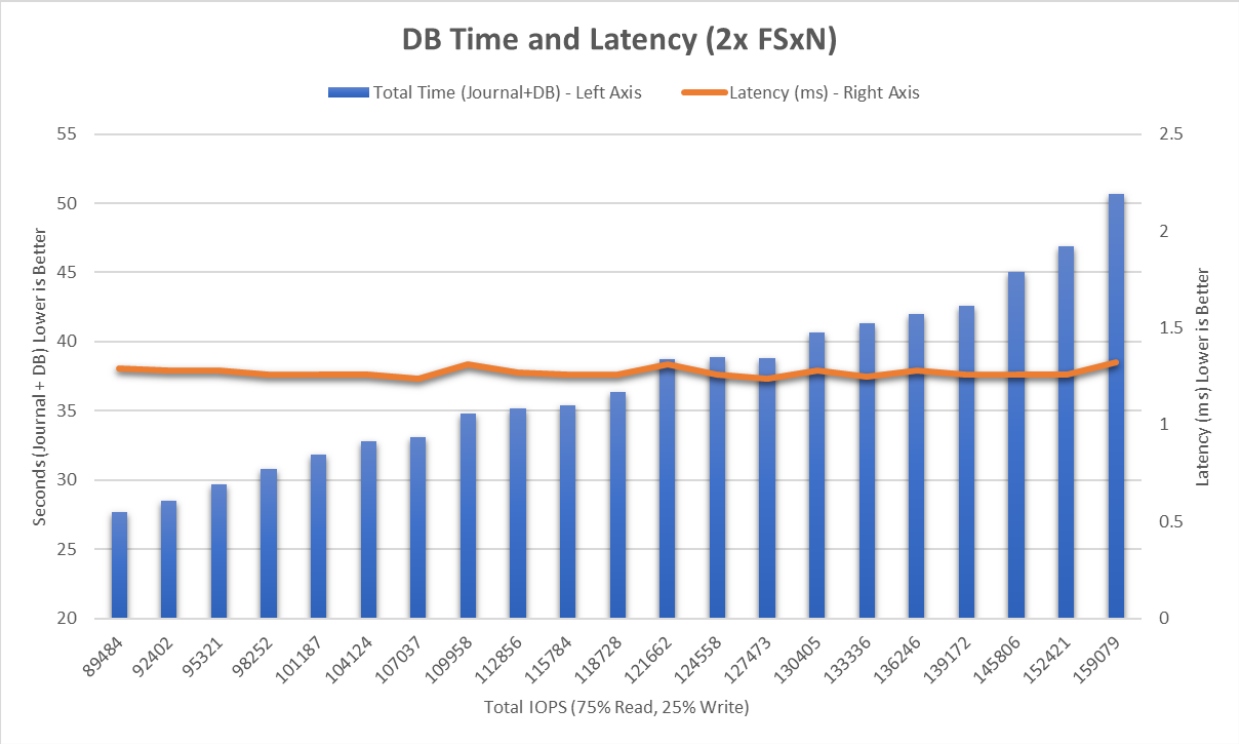


Figure 16 provides a comparison between one and two FSx for ONTAP systems.

**Figure 16: One FSx for ONTAP filesystem versus two.**

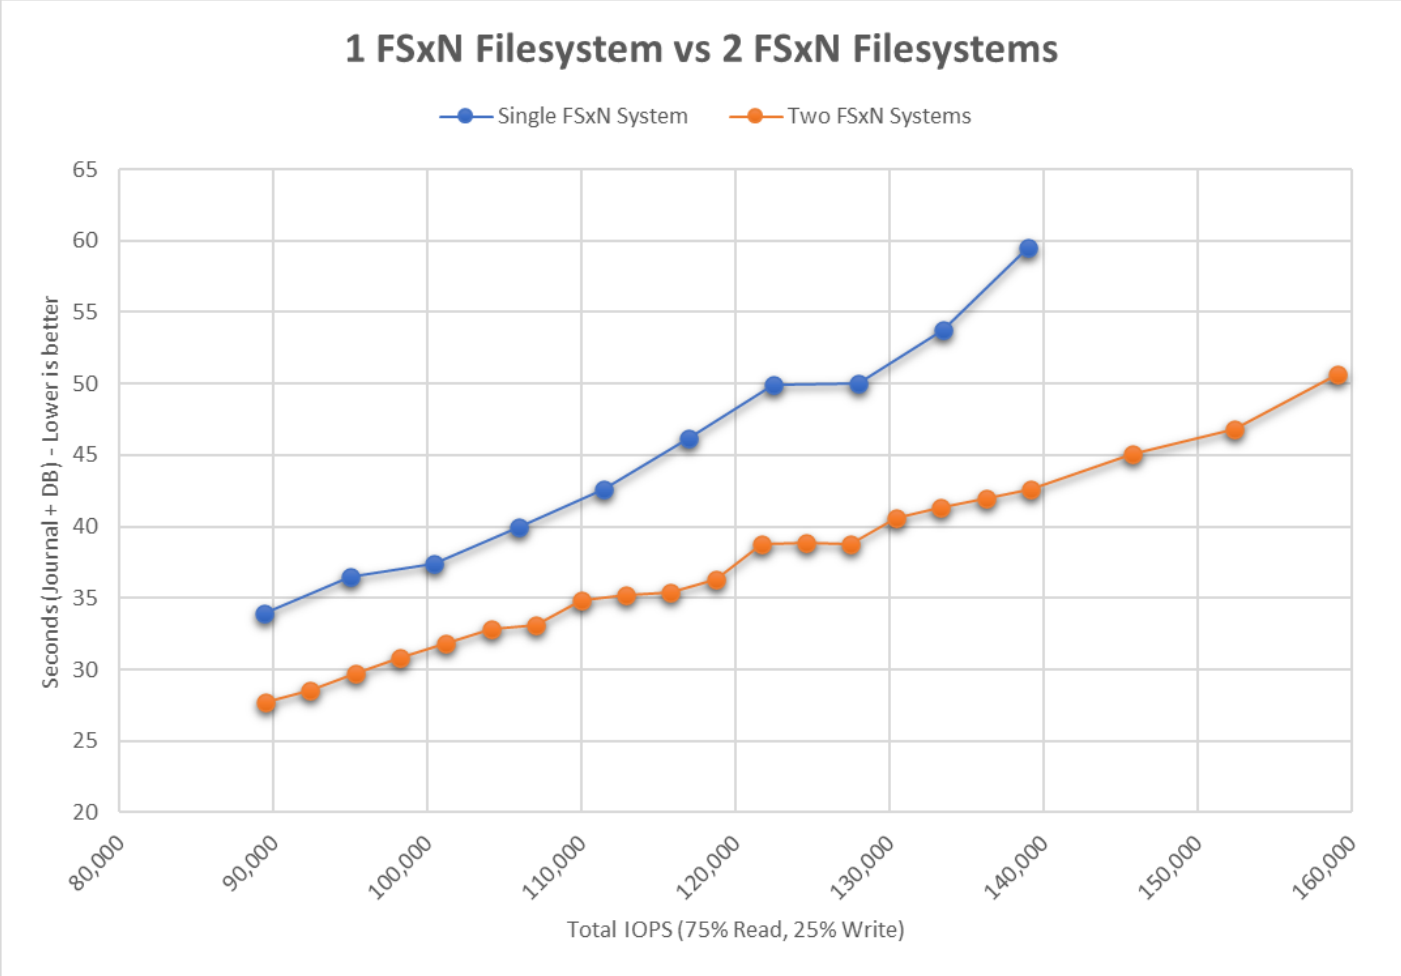


As can be seen in Figure 16, two systems consistently outperformed one system by at least 7 seconds across the full range of IOPS tested.

# Conclusions

NetApp can run as a first-party service in AWS with the FSX for ONTAP offering. Our testing demonstrates that production workloads can be run in AWS with FSX for ONTAP using key features built into ONTAP. Other common use cases are disaster recovery, test/dev, and file shares.

# Where to find additional information

To learn more about the information described in this document, refer to the following documents and/or websites:

- NetApp for SAN
  https://www.netapp.com/data-storage/san-storage-area-network/
- Epic on modern SAN
  https://www.netapp.com/media/27905-nva-1159-design.pdf

- NetApp best practices for Epic TR-3928
  https://www.netapp.com/media/17137-tr3928.pdf
- Business Continuity Guide
  https://www.netapp.com/cyber-resilience/data-protection/disaster-recovery/
- AWS FSx with NetApp
  https://www.netapp.com/aws/fsx-ontap/

# Version history

| Version | Date | Document version history |
|---|---|---|
| Version 3.0 | January 2023 | Updates to best practices/configurations and performance data. |
| Version 2.0 | November 2021 | Updates from AWS. |
| Version 1.0 | October 2021 | Initial release. |

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

**Copyright information**



**Trademark information**



TR-4937-0822

NetApp