



W E B I N A R

# Achieve Cost Optimisation & Innovation While Accelerating Business Transformation with AWS for SAP

Patrick Flückiger

Senior Solutions Architect AWS Storage Technologies

Narendra Eskala

Senior Solutions Architect SAP



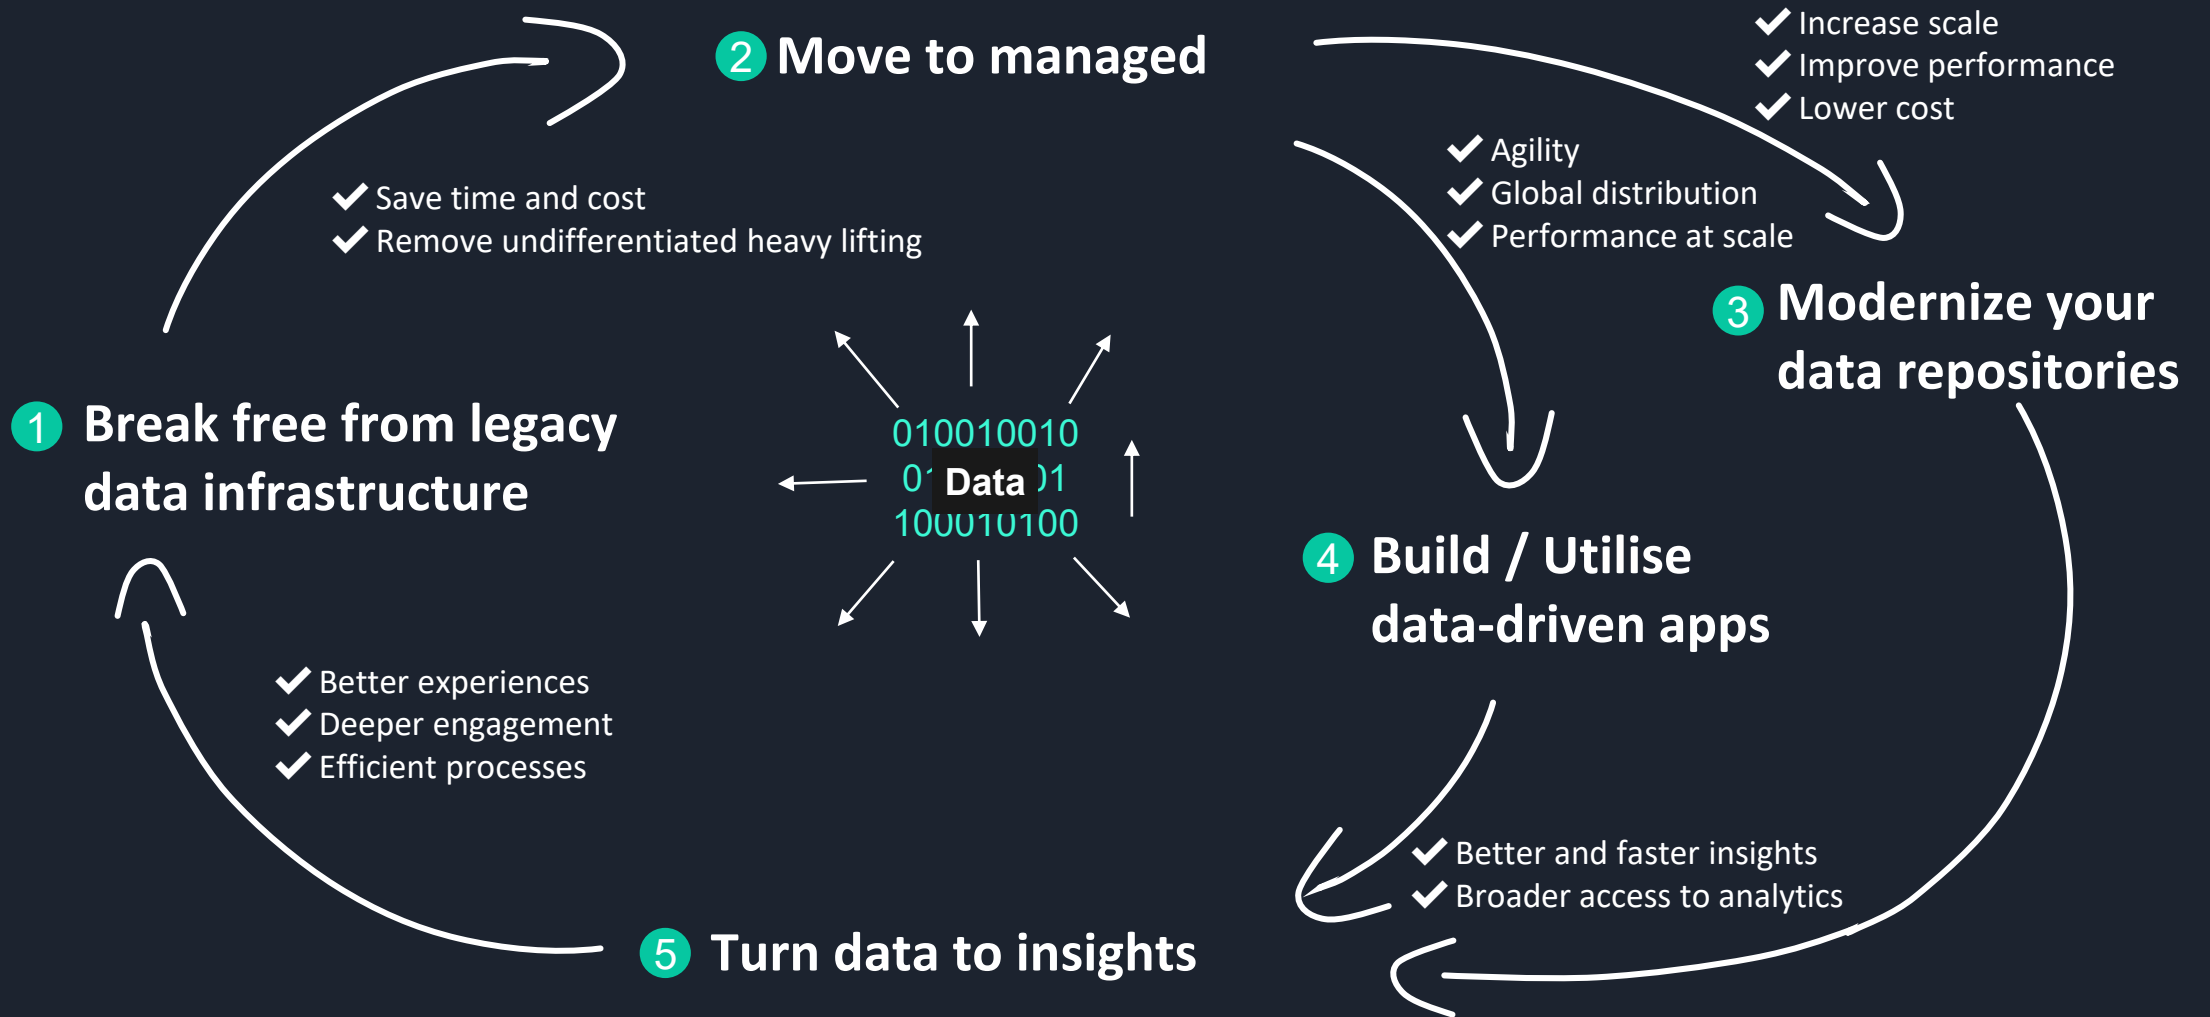
# Agenda

Introduction to Amazon Storage Services and FSxN  
(Patrick Flückiger)

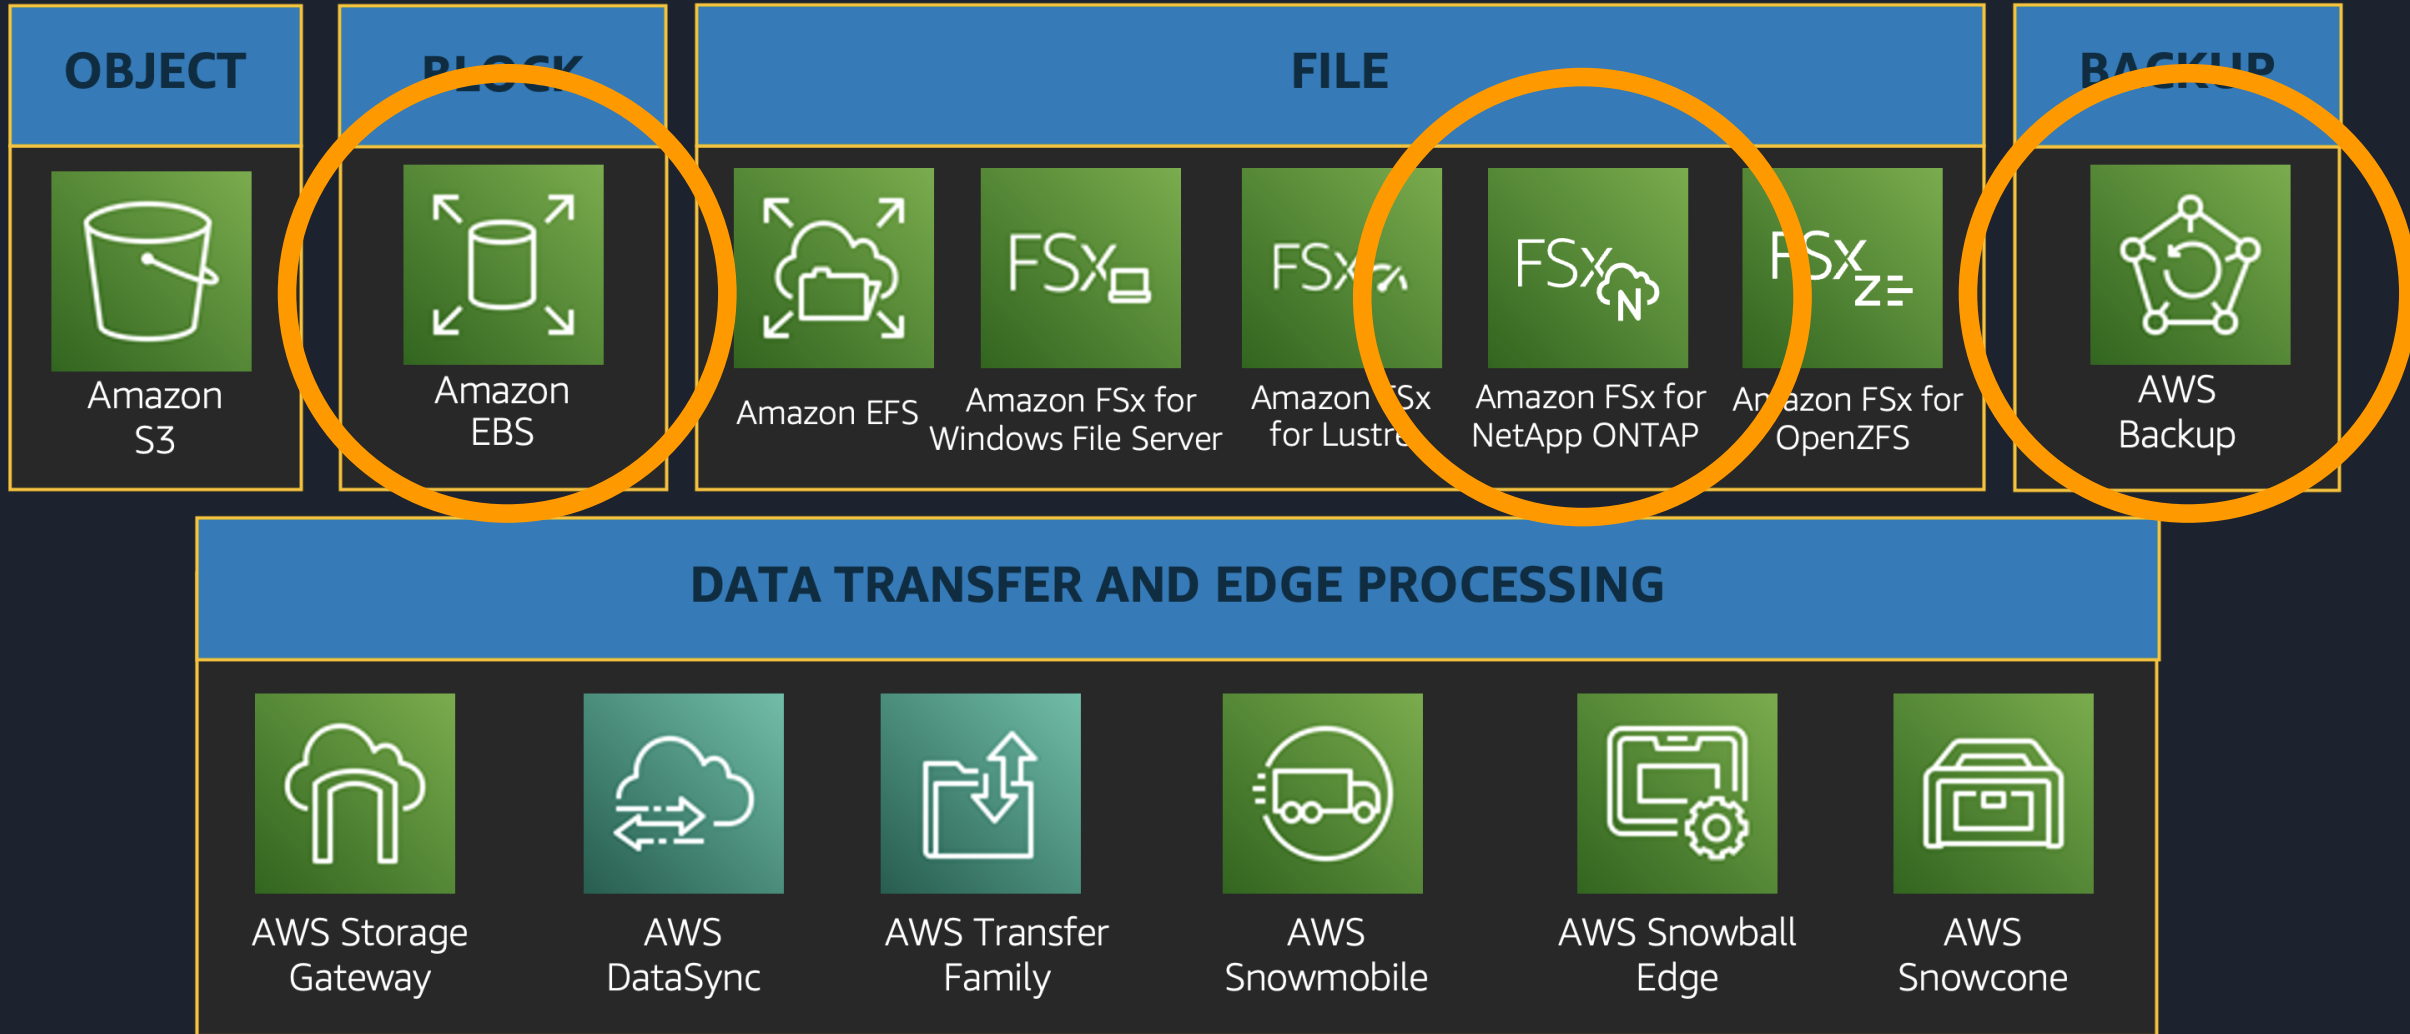
Customer Success & Value for SAP Deployments  
(Narendra Eskala)

# Introduction to Amazon Storage Services and FSxN

# How do you maximise the value of data?



# AWS delivers broadest storage portfolio in industry



# Amazon file storage services



Build cloud-native apps across cloud resources



Run NAS workloads in the cloud



Accelerate compute-heavy workloads



Amazon EFS



Amazon FSx for Windows File Server



Amazon FSx for NetApp ONTAP



Amazon FSx for OpenZFS



Amazon FSx for Lustre



Amazon File Cache

# Amazon FSxN is the easiest path to run NetApp ONTAP workloads



Migrate to like-for-like storage with the same capabilities and APIs



Re-architect



Re-invent



Re-certify

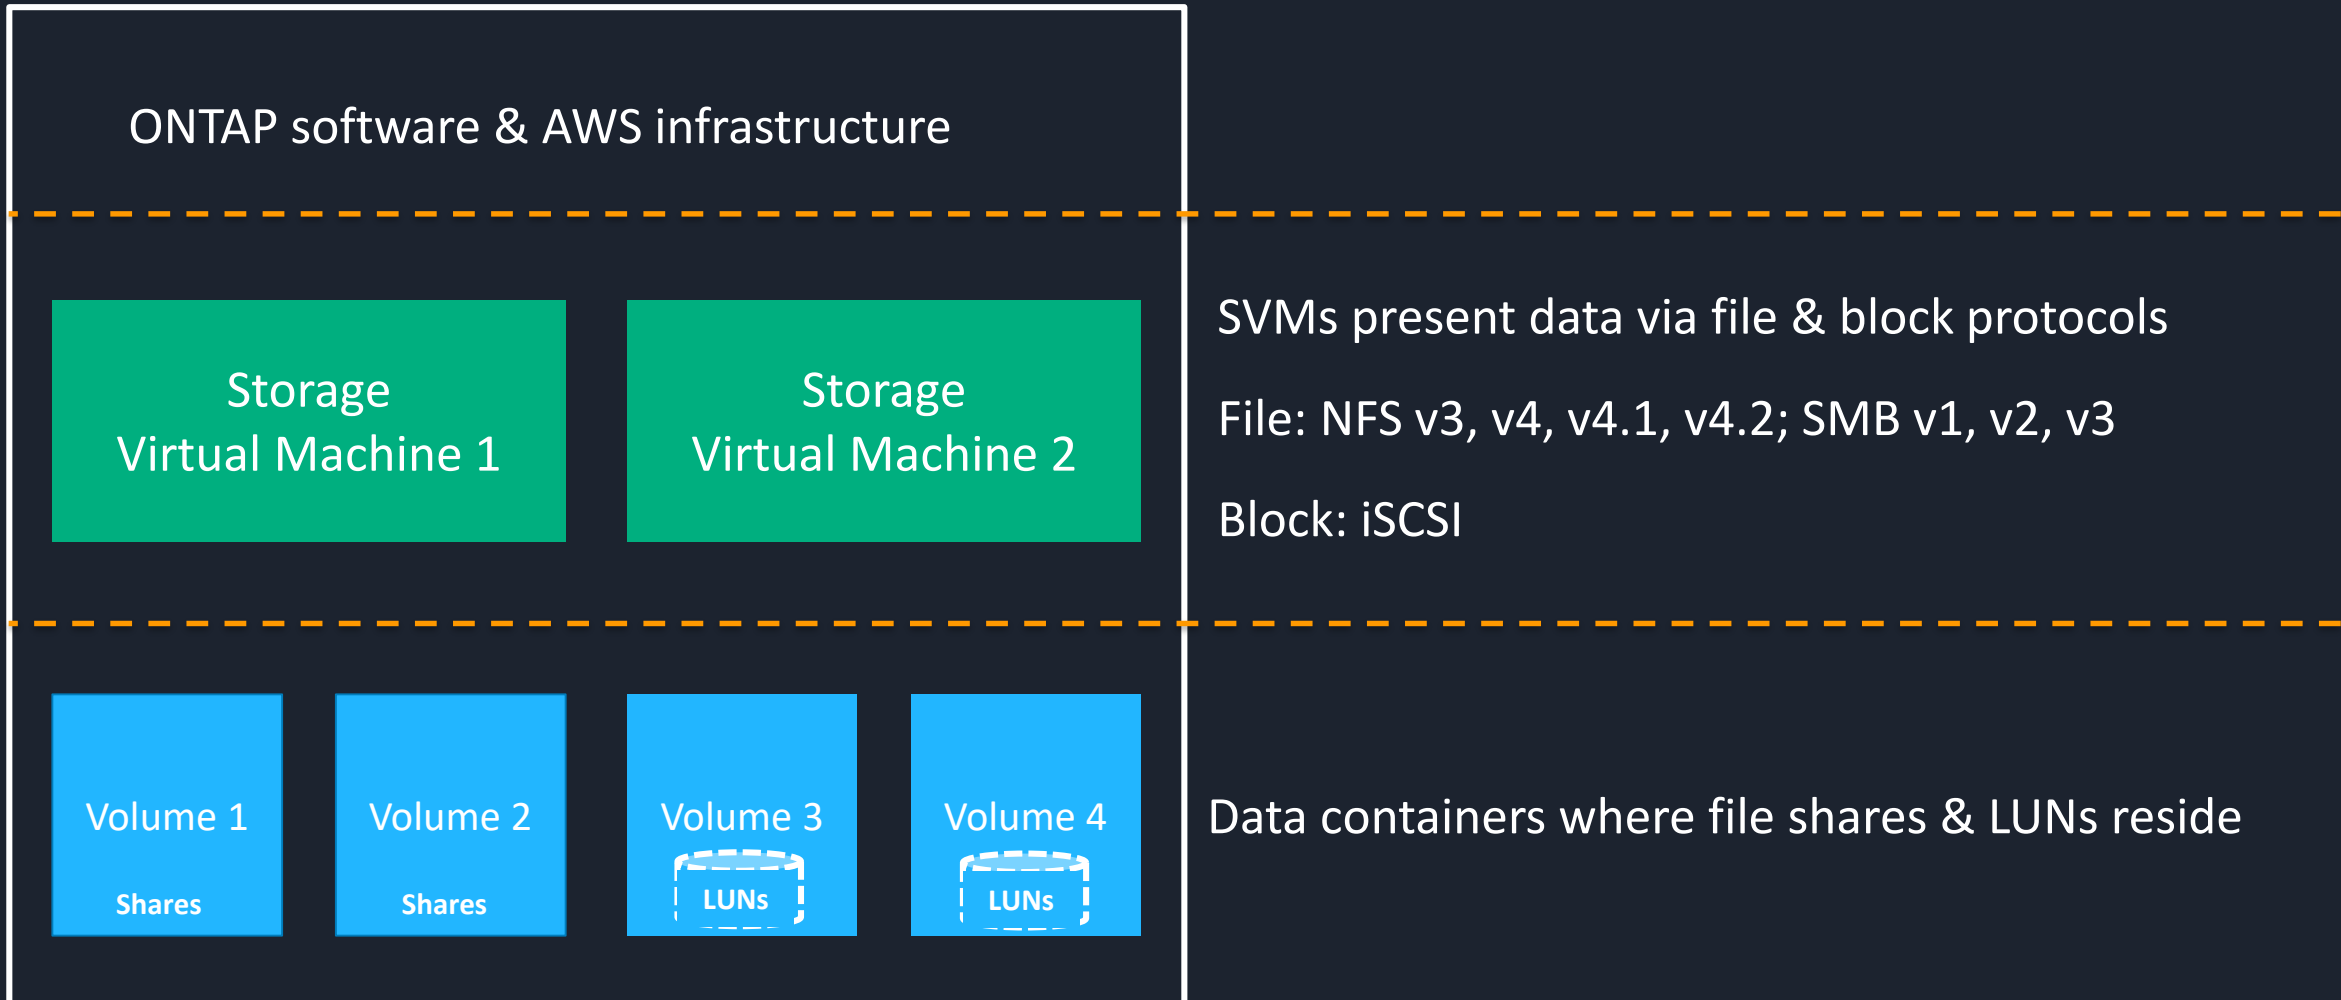


Re-train



Build and run SAP workloads leveraging the capabilities and performance of popular a NetApp ONTAP on-premises system

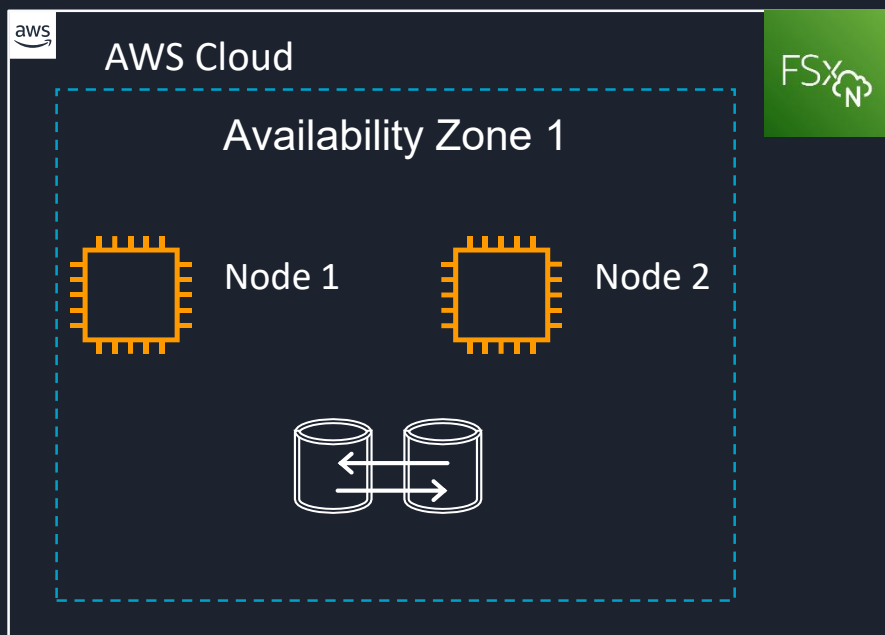
# Amazon FSx for NetApp ONTAP architecture





# Amazon FSx for NetApp ONTAP architecture

## Single-AZ

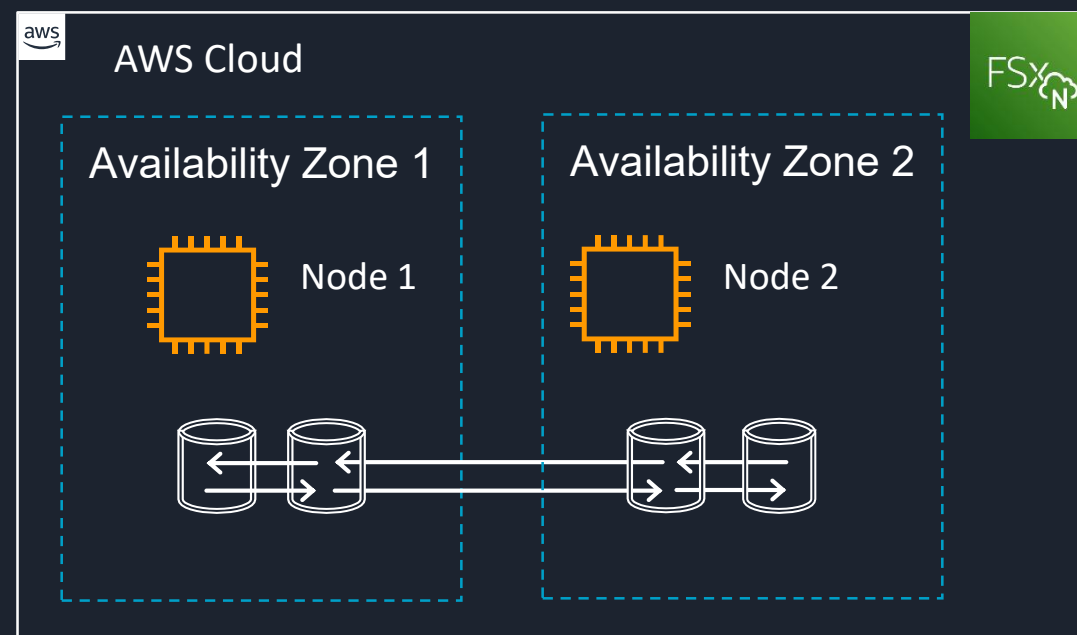


Replication within an AZ

(11) 9's durability for capacity pool and backups

Equal or better availability & durability than a single datacenter implementation

## Multi-AZ



Replication across two AZs

(11) 9's durability for capacity pool and backups

Equal or better availability & durability than a multi-datacenter implementation

# Amazon FSx for NetApp ONTAP key capabilities



Administration and Management



Performance and Scale



Availability and Data Protection



Cost Optimization

## Fully managed, built on ONTAP

- Use AWS and NetApp tools
- Multi-protocol (NFS, SMB, iSCSI)
- Snapshot, clones, replication, caching

## Optimize price and performance

- Multiple GB/s throughput, 100K+ IOPS, sub-ms latencies
- Automatic tiering to low cost, elastic storage
- Deduplication, compression, compaction, thin provisioning

## AWS Integrations

- AWS CloudFormation (ONTAP-as-code)
- Logging with AWS CloudTrail
- EC2, EKS, Workspaces, Appstream 2.0, VMware Cloud
- AWS security, compliance
- AWS Backup (backint)

# Automatic performance and cost optimization with FSxN

## Intelligent policy-based data movement between tiers

### Primary Tier

SSD

Up to 192 TB

Optimized for performance



~20%



Bi-directional data movement  
based on access patterns (hot/cold)

### Automated Tiering Policies

- Snapshot-only (default)
- None
- Auto
- All

### Capacity Pool Tier

Elastic

Unlimited capacity (PB+ file systems)

Cost-optimized for less accessed files



~80%

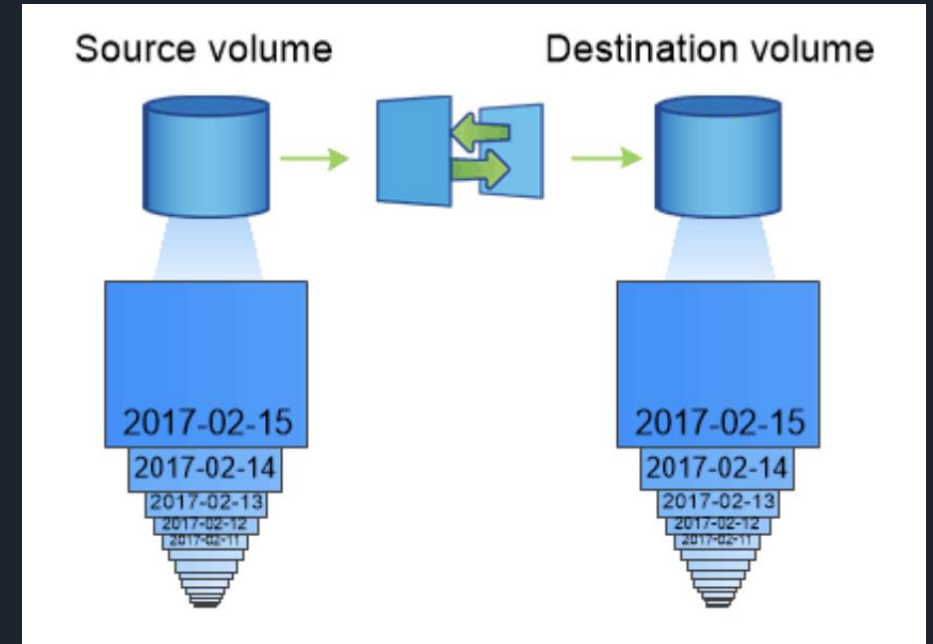


# Snapshots: FSxN / NetApp ONTAP's great strength

FSx ONTAP creates inline snapshots that are **space-efficient**, create no performance impact, and are **independent of database size**.

- Stored on the primary volume. Can be replicated to a volume in a another AZ or Region (SnapMirror).
- Snapshots are application-consistent (additional steps on application side may be necessary).
- Near-instantaneous create and restore.

SAP HANA database backups (e.g., SAP backint or AWS Backup) are stored in Amazon S3. Creation times are dependent on database size.



# SAP + NetApp ONTAP: 20+ year history; now on AWS!

- [Amazon FSx for NetApp ONTAP](#) is a complete implementation of NetApp ONTAP, with the simplicity, agility, and scalability of an AWS service.
- Customers have operationalized SAP on NetApp ONTAP for [20+ years on-prem](#).
- SAP [certified FSxN](#) for HANA (scale up, scale out) in 2022.



# Start fast, start now with Amazon FSxN



- Bi-directional, high-speed data management/replication between availability zones or regions with [SnapMirror](#)



- [Seamless expansion/migration](#) of your existing on-prem NetApp footprint into the Cloud



- Consume additional capacity/systems as required. [Fully integrated with AWS services](#)



- [SnapCenter](#) enables near-instantaneous backups regardless of the size of your SAP databases.

# Start fast, start now with Amazon FSx for NetApp ONTAP



- Administration & Automation



- Point-in-time, instantaneous cloning



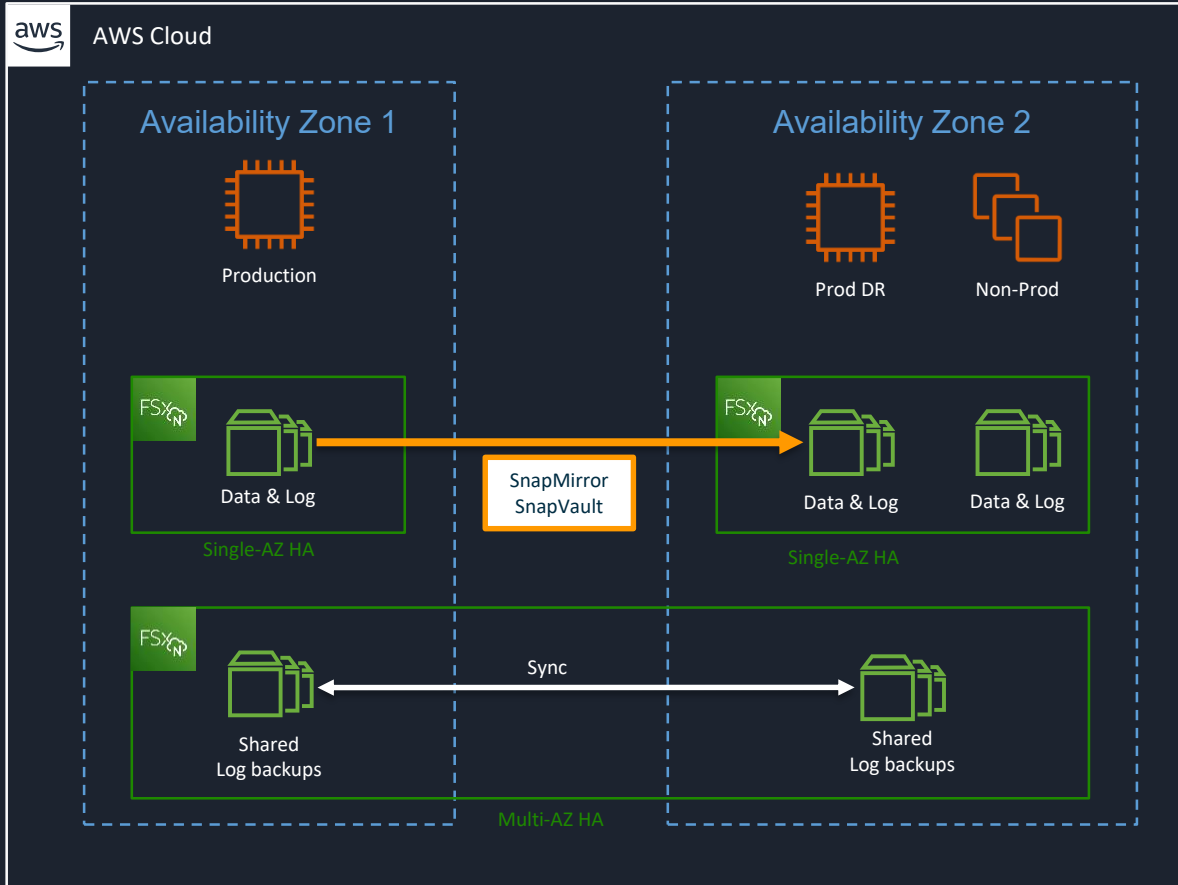
- Low-latency access



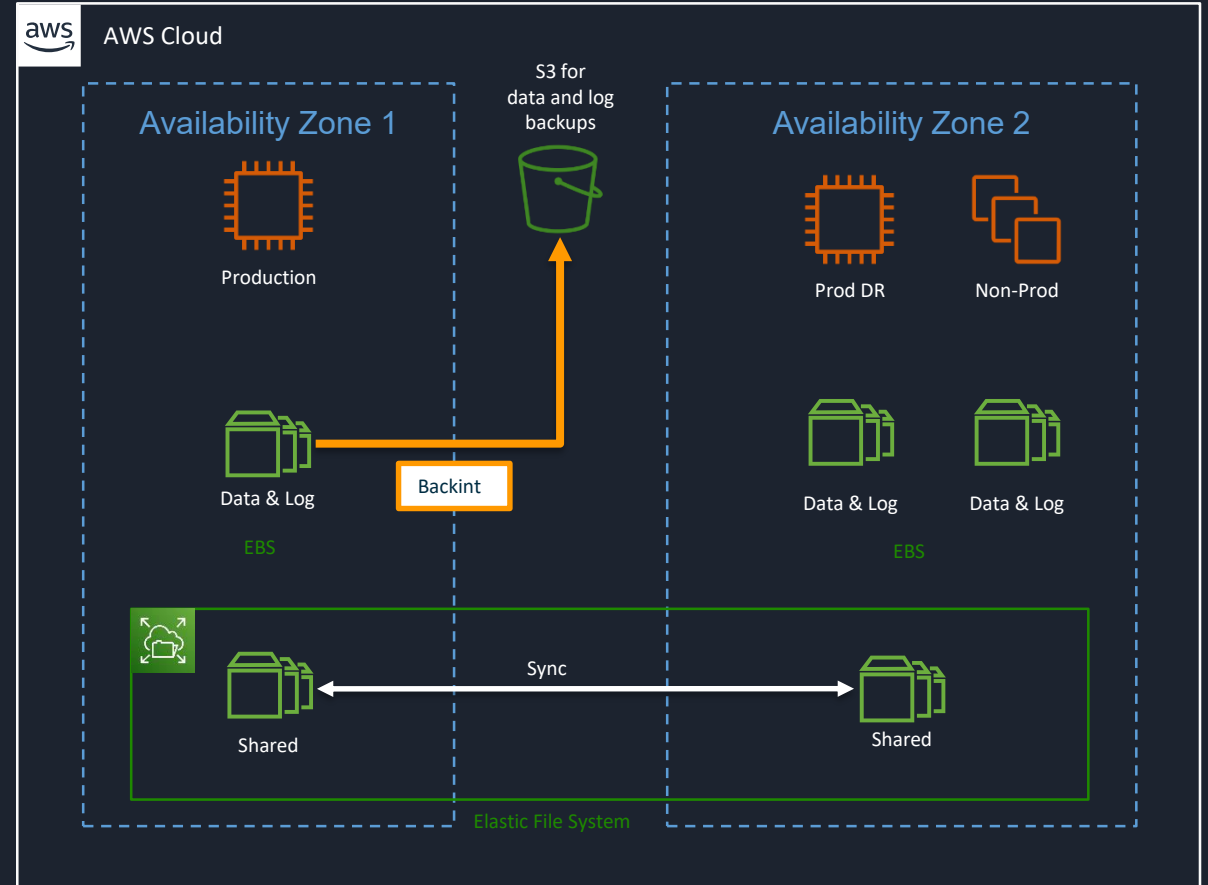
- Multi Availability Zone & Single Availability Zones

# FSxN vs. EBS architecture differences

SAP HANA with FSxN  
Snapshot-based backup/restore



SAP HANA with EBS, EFS  
Backups are stored in S3 via Backint/AWS Backup





# Customer S/4HANA requirements

- **Modernize** their legacy SAP ECC (ERP Central Component) deployment on AWS, **consolidate** three ERP silos to S/4HANA in AWS.
- **Complete** restore / recovery of their 5.8 TB HANA database in **10 minutes**.
- As they scale their database to 48 TB, their requirement is to **complete** HANA restore / recovery in **<60 minutes**.
- Cross-AZ and cross-region **resilience**.

# Why customers chose FSxN for their S/4HANA deployment

- Ability to complete restore and recovery of HANA in **10 mins**.
- Ability to create point-in-time database-consistent snapshots of their 5.8 TB database in **<2 mins**.
- Ability to provision refresh and repair copies in **<1 min**.
- Ability to **optimize architecture** by:
  - Tiering snapshots to lower-cost capacity pool storage.
  - Using the DR site for multiple environments (Dev, Test, QA, Sandbox) without disrupting DR replication from Production.

# Customer quotes

*“Utilizing FSx for NetApp ONTAP for our SAP environment enables us to leverage NetApp ONTAP, the same technology as our on-prem deployment. This was a key reason for selecting AWS over Azure for this workload.”*

Biotech Company

*“Our stats are better than we imagined. With FSxN, system backups of our 9TB S4 database instance take less than two minutes, we can restore & recover the instance in 10 minutes, and we can restart our entire environment (DB + app servers) in 20 minutes.”*

Apparel company

*“This changes everything. Until now, we were cloud agnostic in all our decisions. Not anymore, for now it will be AWS first!”*

Apparel company

# Agenda

Introduction to Amazon Storage Services and FSxN  
(Patrick Flückiger)

Customer Success & Value for SAP Deployments  
(Narendra Eskala)

# Customer Success & Value for SAP Deployments

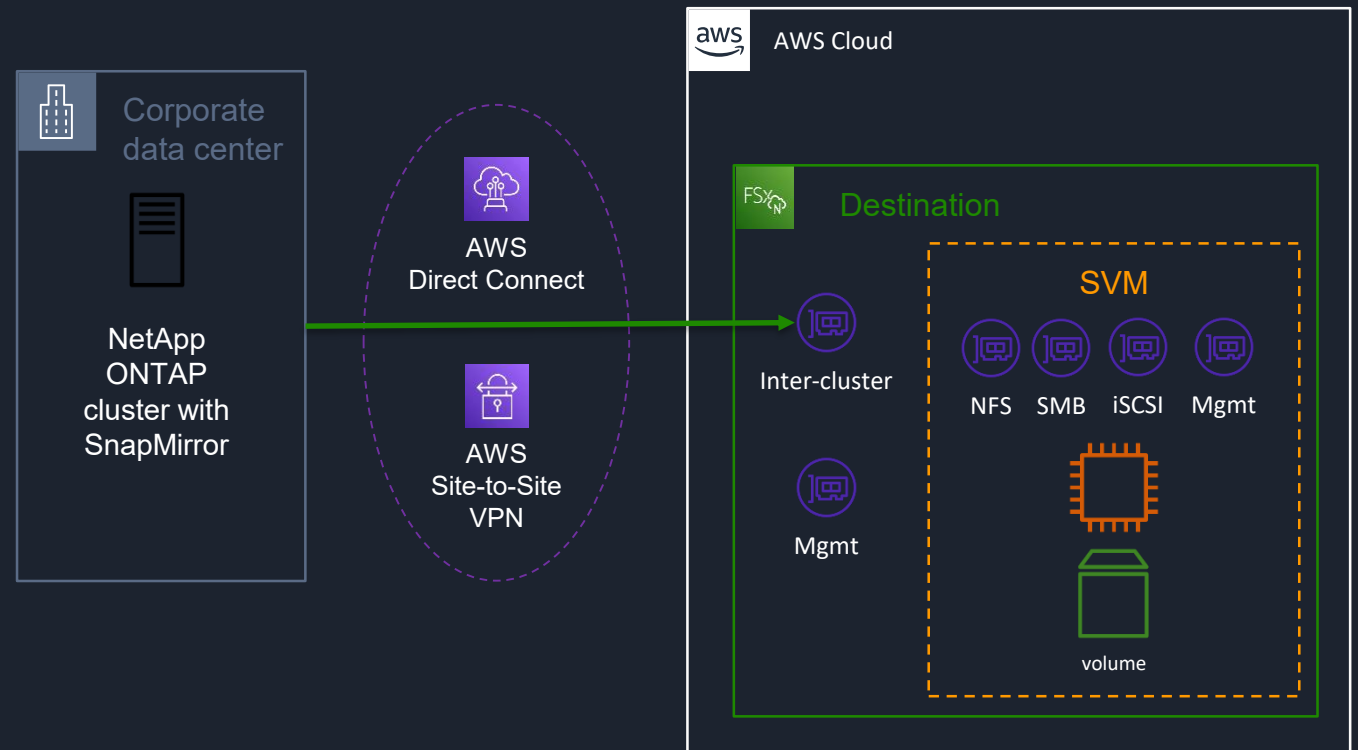
# Migrating from on-premises to AWS

Customers that already use NetApp ONTAP in their own data centers can use SnapMirror to migrate to Amazon FSx for NetApp ONTAP.

- Deduplicated and compressed data remains in those states, which reduces transfer times and reduces the amount of bandwidth required for migration.
- Snapshots that exist on the source ONTAP volumes are preserved when migrated to the destination volumes.

Additional advantage:

- Important processes, such as backup, restore and DR, can remain the same as on-premises.
- Existing automation remains intact.



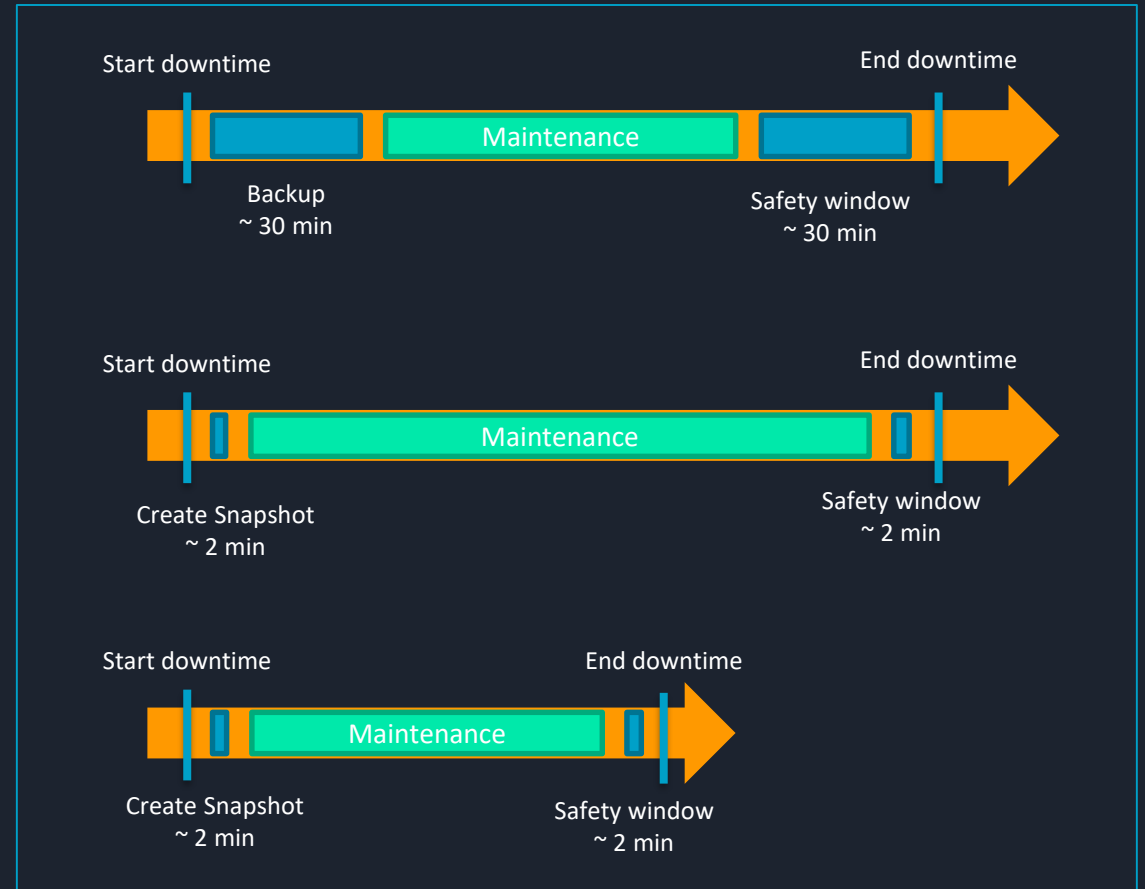
# Shorten downtime during maintenance

Maintenance activities often require a system downtime. A database backup is fundamental in case the system needs to be restored to the previous state. The time needed for creating the backup and a safety window, in case a restore is necessary, needs to be factored into the downtime window.

**With FSx ONTAP you can create and restore snapshots within minutes.**

Reducing the time of both, backup and restore procedures, leads to

- Longer time window for maintenance tasks, or
- Shorter downtime.



# Addressing application-level corruptions (logical)

Application-level corruptions are logical corruptions, caused, for example, by

- Hardware failure,
- Accidental deletion,
- Import of transports containing table or dictionary changes, or
- Failure during maintenance activities.

FSx ONTAP allows creating copies or clones of the corrupted system within minutes. This repair system can be used for investigation and correction of the corruption.

The fix can then be applied on the corrupted system and the repair system is decommissioned.

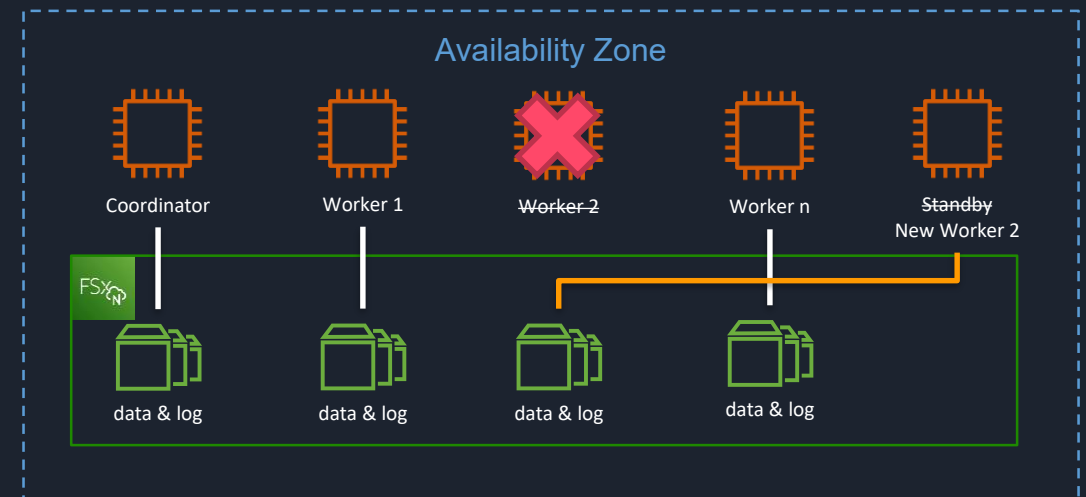




# SAP HANA Host Auto-Failover

FSx ONTAP supports SAP HANA Host Auto-Failover

- **High-Availability** within the same AZ for SAP HANA scale-out landscapes.
- Standby instances can take-over data and log volumes of a failed instances and assume their role.
- Take-over is fully automated without human intervention.
- All volumes are mounted on all instances at all times.
- Each volume is accessed by only one instance at any time.
- Multiple standby instances possible.
- RPO = 0
- RTO = Time to discover failure + HANA data load



# Disaster Recovery (cross-AZ)

Failover to second availability zone within minutes

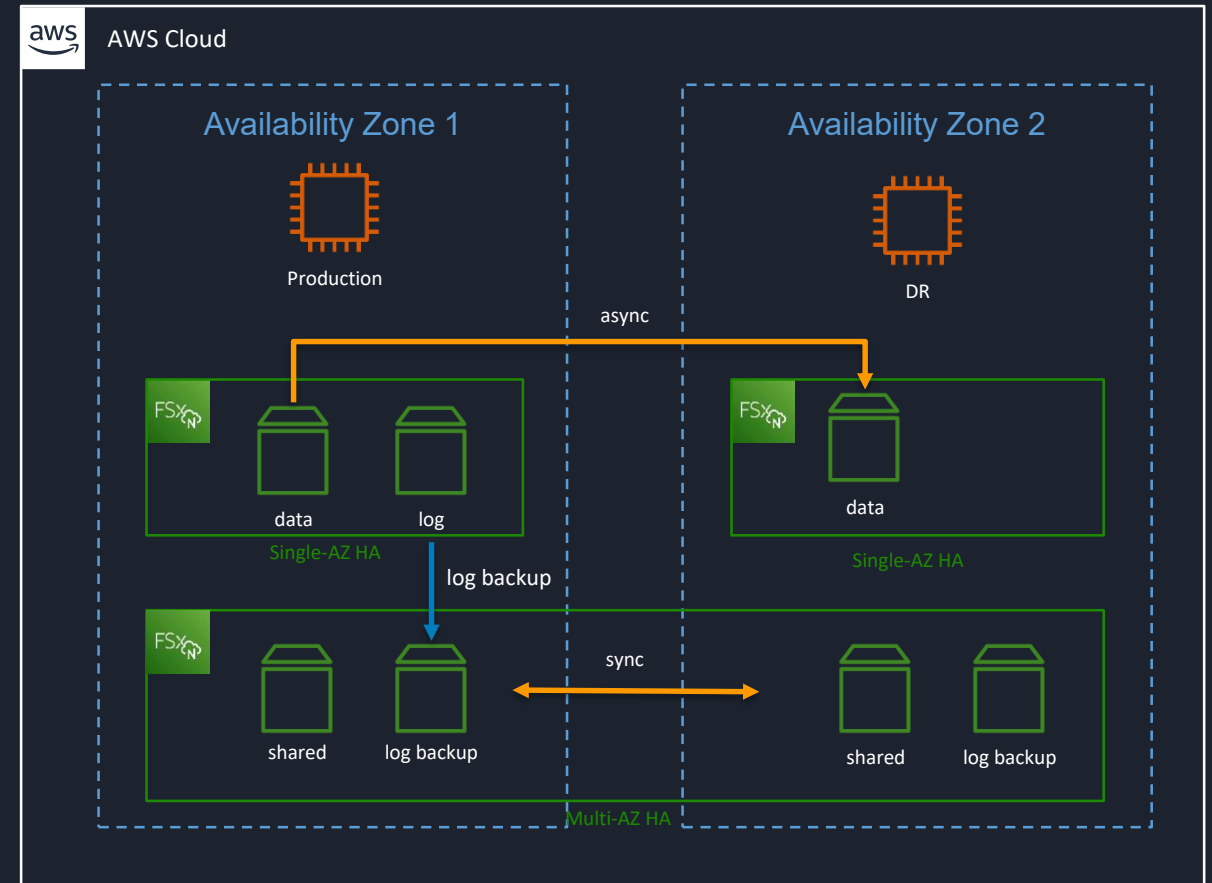
FSx ONTAP offers rapid disaster recovery from replicated snapshots.

- Data & log volumes on **Single-AZ HA** file system.
- Data volume snapshots are copied to second AZ by the file system, using **SnapMirror**.
- **Log backup & shared** volumes are available from both AZs due to **Multi-AZ HA** file system.
- EC2 instance for DR can be used otherwise until failover.

In the event of a disaster, SAP HANA will

1. recover from the most recent data volume snapshot, and
2. roll-forward from synchronously replicated log backups,

Leading to an **RPO of ~15 minutes**. (based on HANA default settings for log backup interval)



# Disaster Recovery (cross-Region)

## Failover to second region within minutes

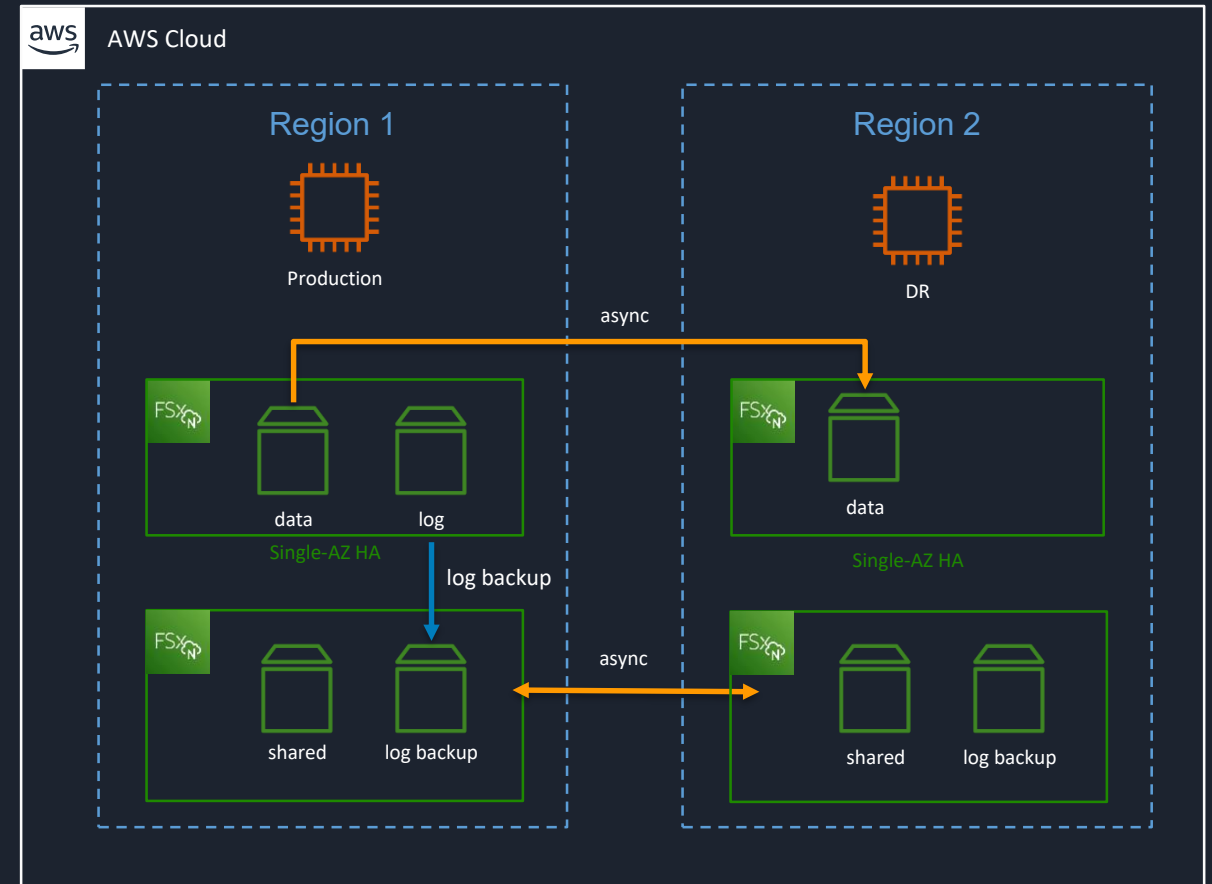
FSx ONTAP offers rapid disaster recovery from replicated snapshots.

- Data & log volumes on **Single-AZ HA** file system.
- Data volume snapshots, log backups and shared volumes are copied to second region by the file system, using **SnapMirror**.
- EC2 instance for DR can be used otherwise until failover.

In the event of a disaster, SAP HANA will

1. recover from the most recent data volume snapshot, and
2. roll-forward from asynchronously replicated log backups,

Leading to an **RPO of ~15 minutes**. (based on HANA default settings for log backup interval)



# SAP Certification Scope

FSx ONTAP is SAP certified for

- SAP HANA
  - Single-AZ HA for data and log volumes.
  - Sizing based on throughput requirements.
- SAP NetWeaver
- Databases, in combination with SAP applications
  - Oracle DB, IBM Db2, SAP ASE, SAP MaxDB (incl. LiveCache)

Details in SAP Note 1656250 - SAP on AWS: Support prerequisites.

<https://me.sap.com/notes/0001656250>

## Amazon FSx for NetApp ONTAP

[Amazon FSx for NetApp ONTAP \(FSx for ONTAP\)](#) is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on NetApp's popular ONTAP file system.

Amazon FSx for NetApp ONTAP is **supported** for

- SAP NetWeaver file systems, like /sapmnt, /usr/sap/<SID>, and /usr/sap/trans.
- SAP HANA 2.0
  - shared file systems (e.g., /hana/shared, /usr/sap),
  - database files (data and log volumes).
- SAP ASE
- SAP MaxDB and liveCache (versions 7.9.10 and higher)
- The following databases in combination with SAP applications
  - SAP HANA
  - SAP ASE
  - SAP MaxDB and liveCache (versions 7.9.10 and higher)
  - IBM Db2
  - Oracle Database (see SAP Note [2358420](#) for further information)

# SAP Note 1656250 - SAP on AWS: Support prerequisites

<https://me.sap.com/notes/0001656250>

## Supported instances

- r6i, x2idn, x2iedn, u-\*tb1.
- Upcoming generations of commonly used instances for SAP HANA will be certified as well.
- Scale-out for all instances that are scale-out certified in general.
- OLTP and OLAP workloads.

## Requirements

- Only Single-AZ HA for data and log volumes.
- Multi-AZ HA is allowed for shared volumes.
- NFS only.

Performance-based sizing allows right-sizing, based on throughput requirements. This includes

- Sharing one file system between multiple SAP HANA nodes.
  - QoS policies are available to reduce noisy neighbor effects.
- Scale to meet throughput requirements, by
  - Increase throughput of file system.
  - Add/remove file systems and use SAP HANA data volume partitioning.
  - Both are online operations.

## Amazon FSx for NetApp ONTAP for SAP HANA Database Files

FSx for ONTAP is supported for SAP HANA data and log volumes on selected Amazon EC2 instance types. Support statements are available in the "Restrictions & Comments" section of the respective EC2 instance type in the [SAP HANA Hardware Directory](#).

Customers are responsible for provisioning FSx for ONTAP with sufficient capacity and performance to handle their SAP HANA workloads. A dedicated FSx for ONTAP file system per SAP HANA node is recommended. Sharing a file system between multiple SAP HANA nodes is supported, if the file system meets the combined requirements of all SAP HANA nodes at any time and under all workload conditions. Additional file systems may be required for higher performance demands (e.g., separate file systems for data and log volumes). Workloads from other applications are not allowed on this file system. In case of performance issues on a shared file system, SAP support may require you to use dedicated FSx for ONTAP file systems per SAP HANA node as per the recommendation.

FSx for ONTAP file systems for SAP HANA data and log volumes are only supported with the single-AZ [deployment type](#) and must be mounted via NFS. Scale-out is generally supported with each EC2 instance type that supports scale-out. SAP HANA Host Auto-Failover (HAF) is supported with FSx for NetApp ONTAP.

For further details, please refer to section [FSx for ONTAP for SAP HANA](#) in the SAP on AWS Technical Documentation.



# Setup and administration

## AWS Tools & Services



AWS Console




AWS CLI




Amazon FSx API/SDK



Amazon CloudFormation



Amazon CloudWatch

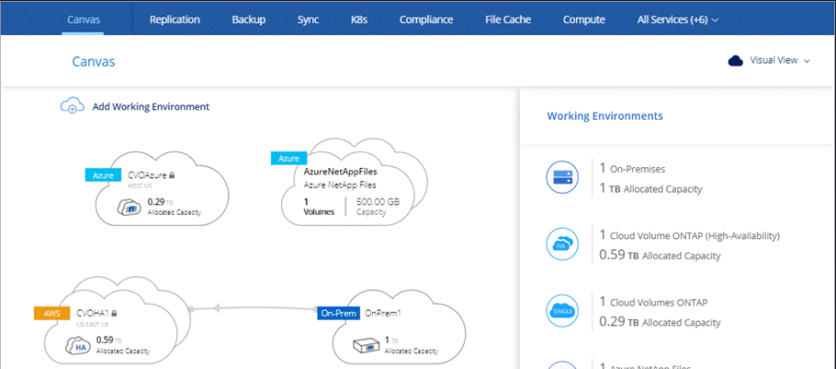


Amazon CloudTrail



## NetApp Tools

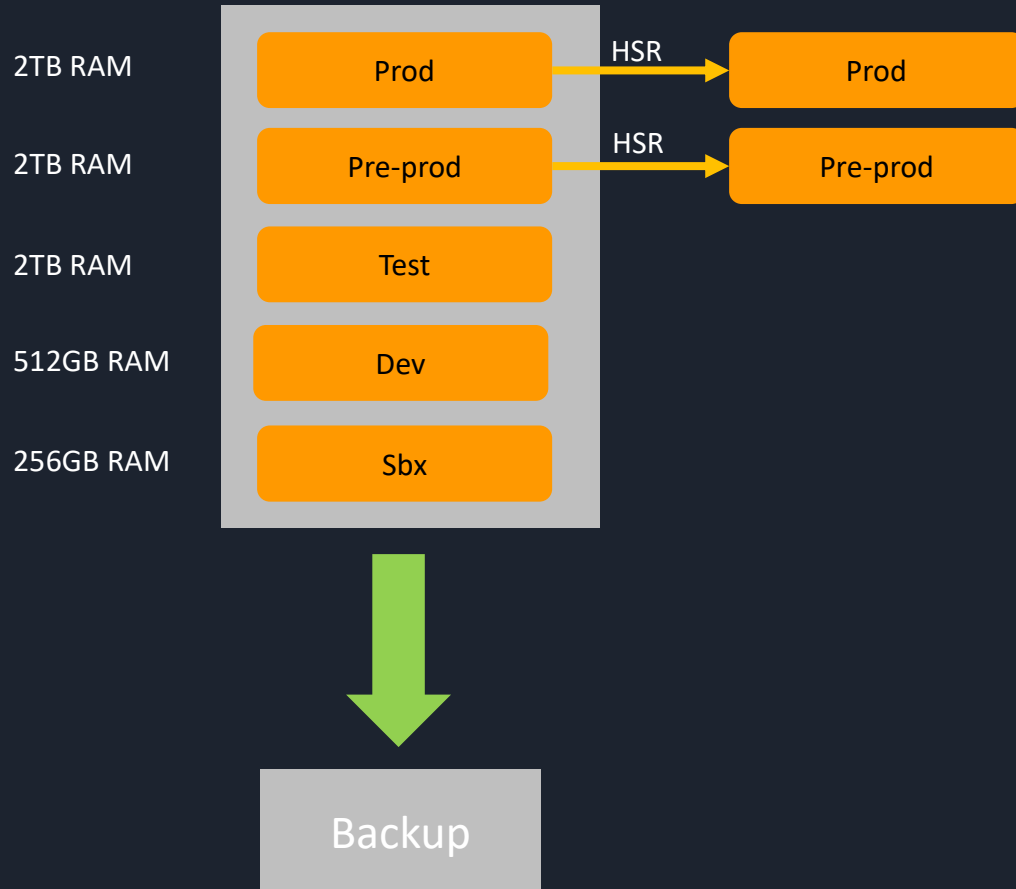
### BlueXP



### ONTAP CLI and API

```
fsxId05365829bb93cf758:~> volume show
Vserver Volume Aggregate State Type Size Available Used%
-----
fsx fsx_root aggr1 online RW 1GB 972.3MB 0%
svm01 svm01_root aggr1 online RW 1GB 972.5MB 0%
2 entries were displayed.
fsxId05365829bb93cf758:~>
```

# TCO example



Assumptions	FSx	EBS (gp3)
HANA production DB size (GiB)	2,048	2,048
Total size of landscape (GiB)	6,912	6,912
Throughput capacity (MB/s)	2 x 1,024	1,500 per data volume
Storage configuration	Efficiencies enabled on binary volumes	
Backup configuration	Snapshots using SnapCenter	Backint to S3
Backup retention (Data and Log backups)	30 days <ul style="list-style-type: none"> <li>3 days on SSD (\$0.125/GB-mo)</li> <li>27 days on capacity pool (\$0.022/GB-mo) ~82% lower cost</li> </ul>	30 days
Backup tiering	<ul style="list-style-type: none"> <li>3 days on SSD</li> <li>30 days on capacity pool</li> </ul>	-
Data change rate (daily)	<b>20%</b>	<b>20%</b>

# Benefits of Amazon FSx ONTAP for SAP HANA

## Bring Existing NetApps ONTAP Developments to AWS

- Seamlessly bring existing NetApp ONTAP based SAP applications to AWS using **SnapMirror**
- **Existing automation remain intact** with this approach
- Backup and DR procedures can be retained

## Minimizing Backup window to Minutes

- Any backups that need to be taken in downtime only requires **couple of minutes**
- Shorter backup windows help in **optimizing overall maintenance** window

## Rapid System Cloning

- FSx ONTAP allows creating copies or clones of the corrupted system **within minutes**
- This will help in addressing application-level corruptions

## Optimize Disaster Recovery RTO

- Bring down **RTO time to ~30** mins by leveraging SnapMirror feature
- Applicable for systems that do not have a DR solution using HSR

## Reduce Database refresh window

- **Accelerates system refreshes** of QA/training and project type systems by reducing time to restore the database

## Seamless failover of SAP HANA in Scale-out scenario

- **High-Availability** within the same AZ for SAP HANA scale-out landscapes.
- Achieve **RPO = 0** with FSx ONTAP for SAP HANA database





Thank you!