AWS re:Invent

STG306-R

Deep dive on Amazon FSx for Windows File Server

Edward Naim
General Manager, Amazon FSx
Amazon Web Services

Prashanth Bungale
Product Manager, Amazon FSx
Amazon Web Services

Darryl Osborne
Solutions Architect, File Services
Amazon Web Services

AWS re:Invent

aws

Agenda

- What is Amazon FSx for Windows File Server?
- Deep dive
- Amazon FSx in action
- Summary and Q&A

What is Amazon FSx for Windows File Server?

What is Amazon FSx for Windows File Server?

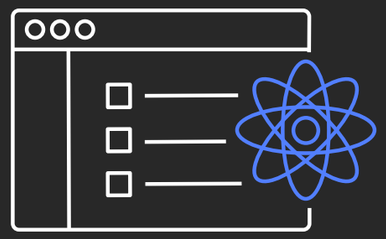

Fully managed native Windows file systems

Deeply integrated with AWS

Fully managed means you no longer need to…

Manage hardware

Plan capacity

Procure and purchase hardware

Set up storage servers and volumes

Detect and address hardware failures

Invest CapEx

Manage software

Install and configure server software

Set up and configure file systems

Apply Windows updates

Manage software licenses

Manage backups

Monitor security

So, how easy does Amazon FSx make it?

Creating an Amazon FSx file system

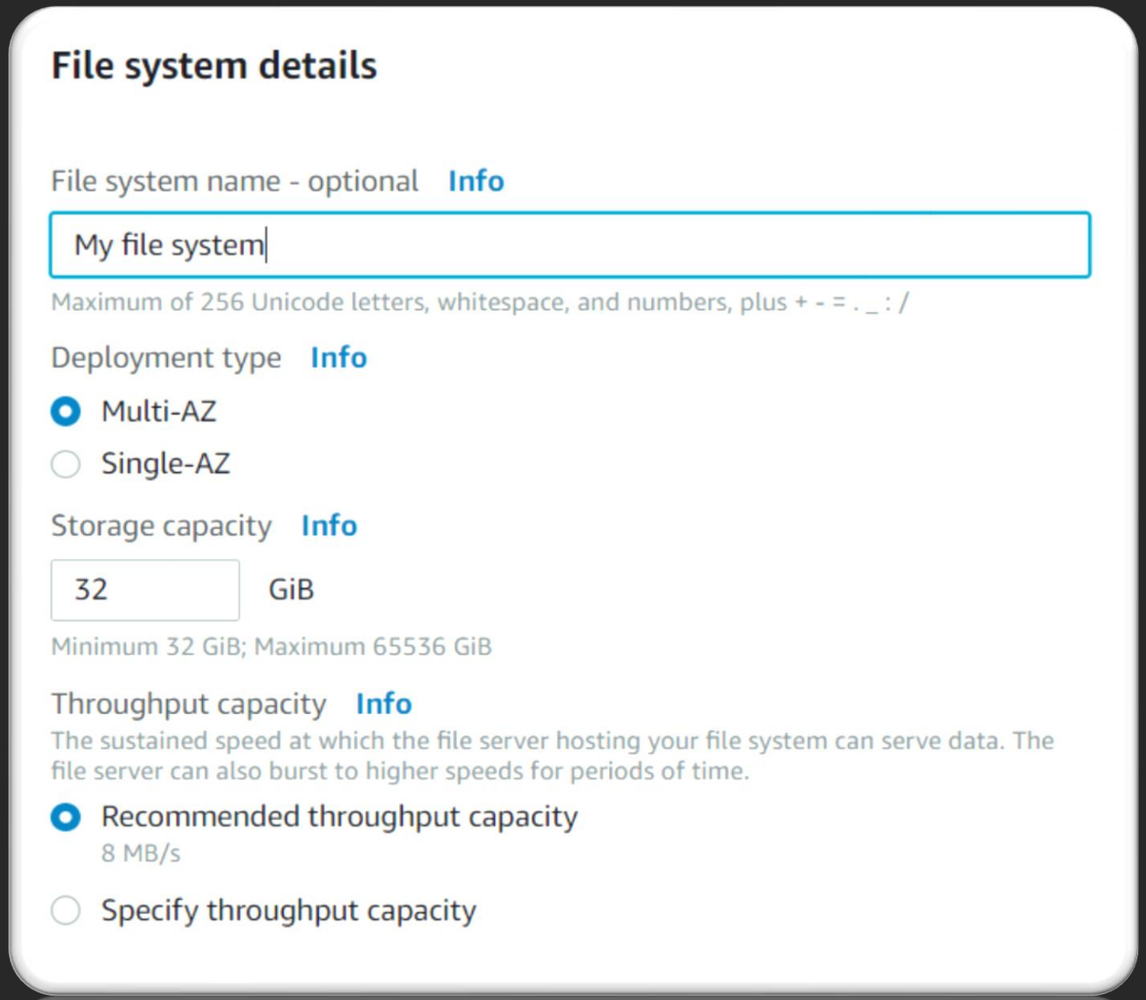

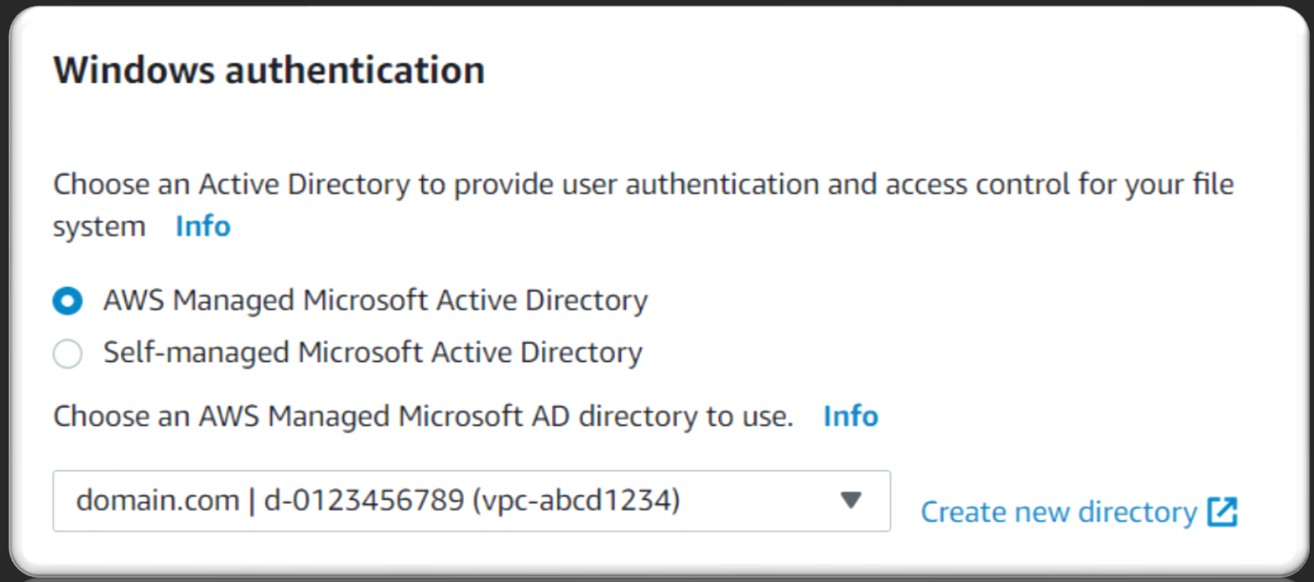

Rich administrative features, when you need them

AWS Management Console / AWS CLI

CloudWatch metrics

Amazon CloudWatch alarms

On-demand backups

Automatic backups schedule

Tags

Amazon FSx CLI on PowerShell

Shadow copies schedule

Data deduplication

Quotas

Enforcement of encryption in-transit

Managing active user sessions and open files

Broadly accessible

Microsoft Windows Server 2008+ and Windows 7+

Linux (SMB client)

MacOS

OS

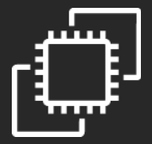

Amazon EC2

Amazon WorkSpaces

VMware Cloud on AWS

Amazon AppStream 2.0

On-premises compute instance

Compute instance

VPN

AWS Direct Connect

VPC Peering

AWS Transit Gateway

In-VPC access

Network connectivity

Fully featured and compatible

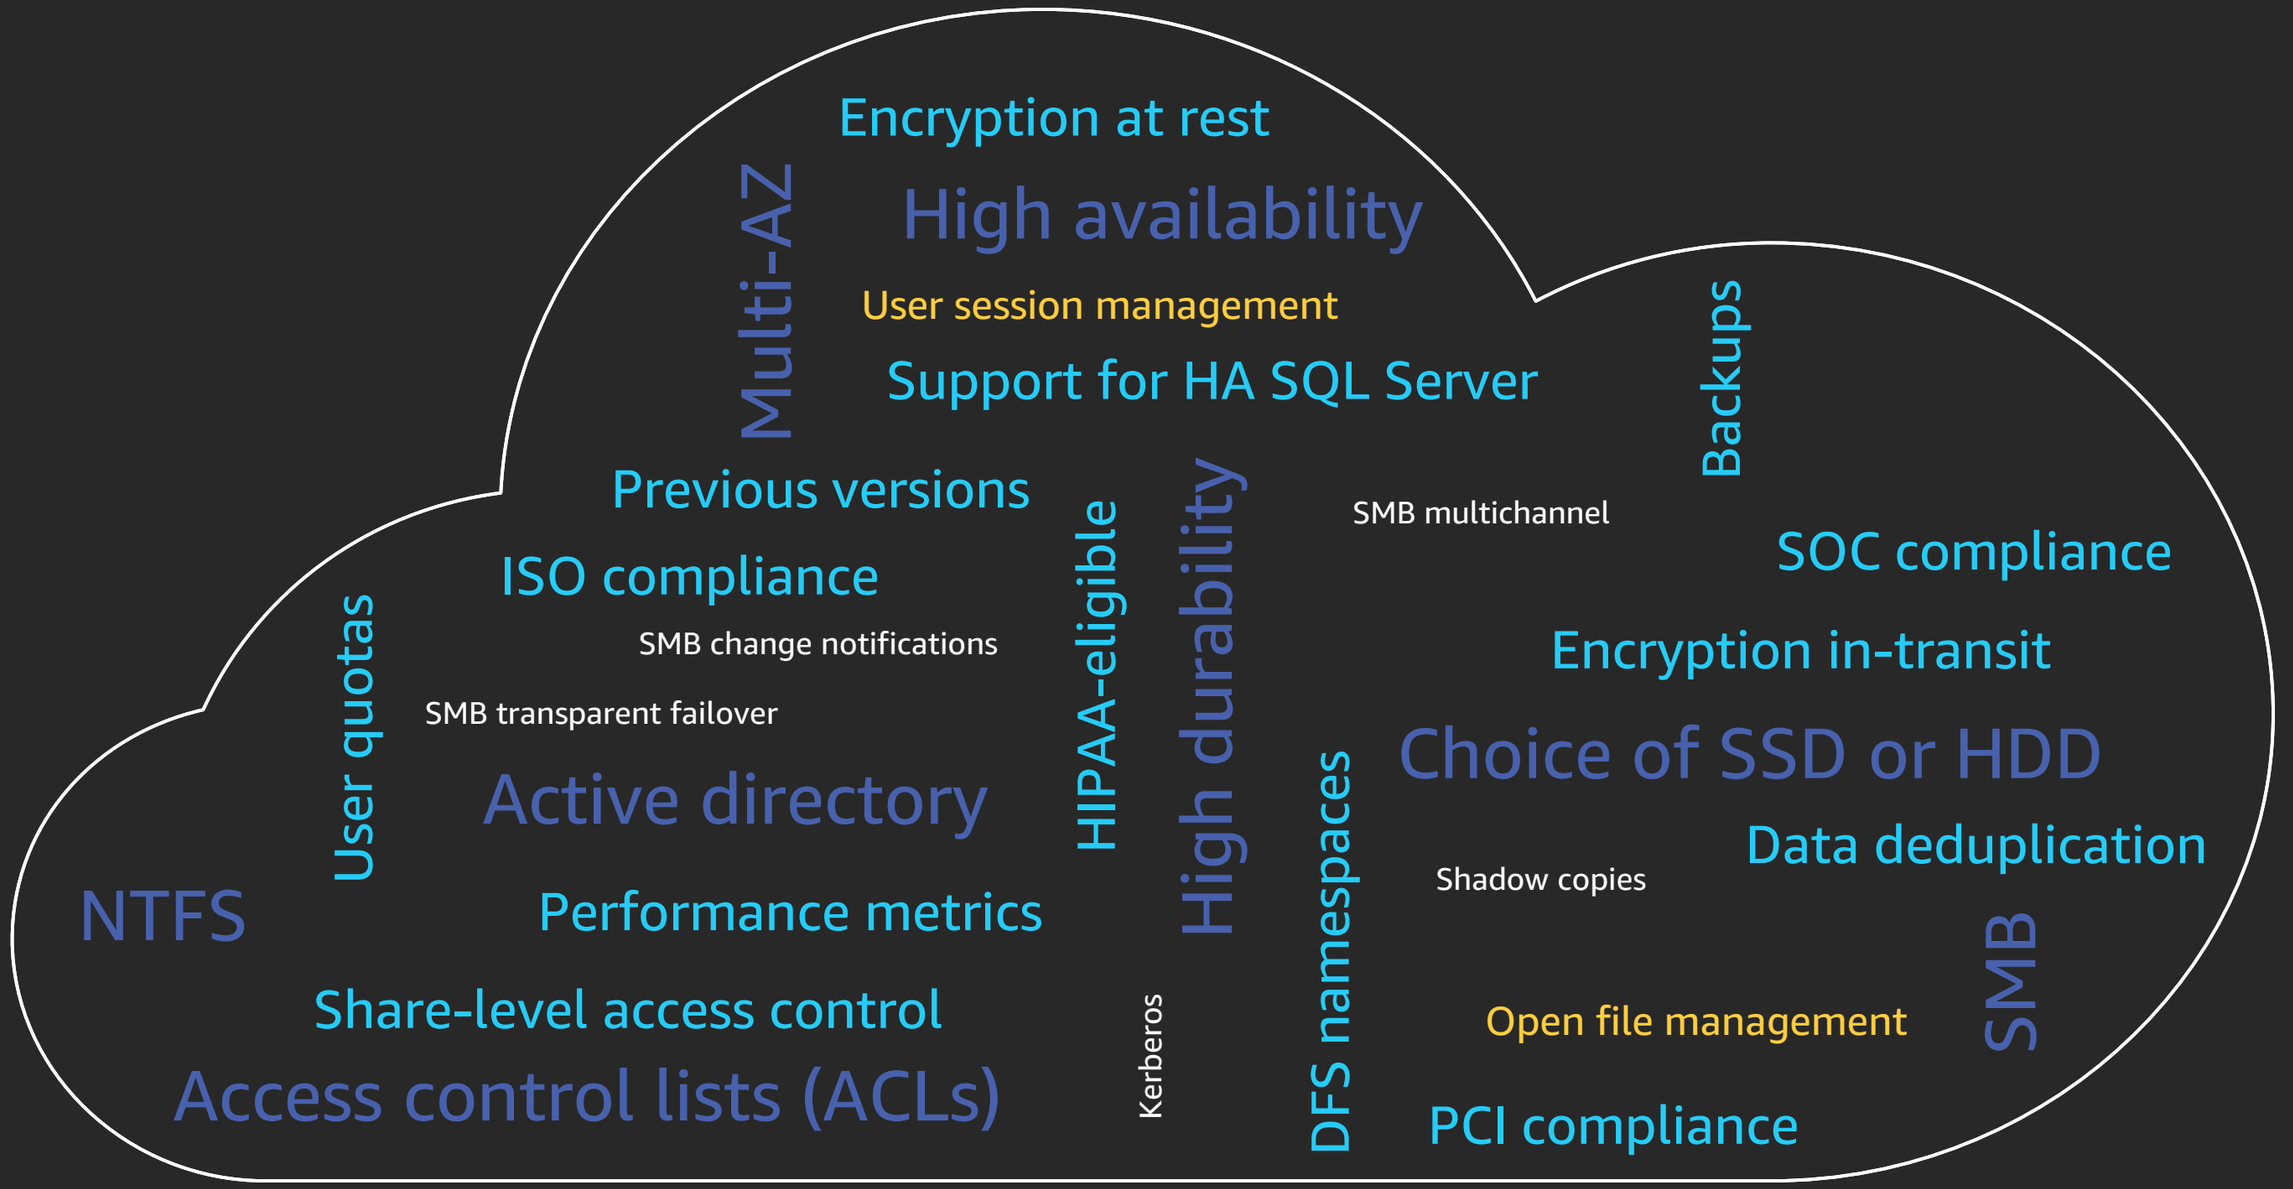

AWS integrations | Amazon FSx

Amazon FSx

Integrations


Amazon VPC


AWS IAM


AWS KMS


Amazon CloudWatch


AWS CloudTrail


AWS CloudFormation


AWS Direct Connect


AWS VPN


Amazon EC2


VMware Cloud on AWS


Amazon AppStream 2.0


Amazon WorkSpaces


Amazon ECS


AWS Management Console


AWS CLI


AWS Directory Service

What fully-managed, fully-featured, and compatible means for customers

*“With Amazon FSx we have a **managed solution** that eliminates the need for us to manage the tasks associated with operating file storage, including performing backups and applying software patches."* -Greg Bell, Senior Systems Administrator

"We were able move our Windows file-based application to AWS quickly and solved a critical performance issue. Based on our successful lift-and-shift experience, we’re looking to move more of our legacy applications requiring Windows native compatibility, to Amazon FSx.” -Hemanth Jayaraman, Director, Cloud Center of Excellence

*“We no longer need to worry about maintenance of the file system environment or patching of the underlying servers. This is allowing the team to **focus on more business-critical parts of our application structure**.”* -John Waller, Senior Engineer, DocXellent

Beyond simplicity and features…

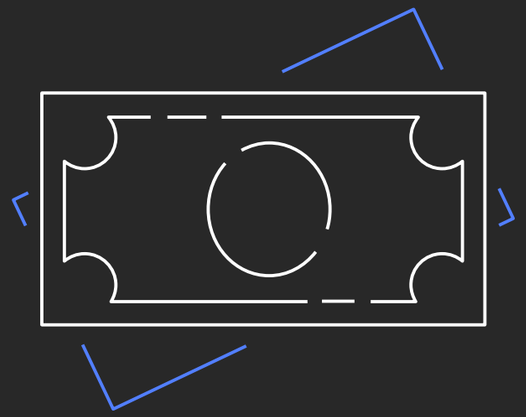

Storage pricing

(per GB-month)

	Single-AZ	Multi-AZ
SSD-based storage	13 cents	23 cents
HDD-based storage (COMING SOON!)	1.3 cents	2.5 cents

Note: Based on Amazon FSx for Windows File Server pricing in US East (N. Virginia)

Effective storage cost with data deduplication

(per GB-month)

	Single-AZ	Multi-AZ
SSD-based storage	6.5 cents	11.5 cents
HDD-based storage (COMING SOON!)	0.65 cents	1.25 cents

Typical savings from deduplication for general file shares is 50-60%

Note: Based on Amazon FSx for Windows File Server pricing in US East (N. Virginia)

High performance

< 1ms latencies

Multiple GB/s

Hundreds of thousands of IOPS

Designed for a wide spectrum of use cases…

NAS lift-and-shift

Home directories

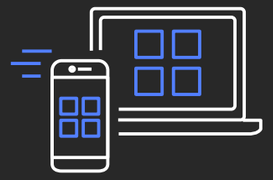

Line-of-business applications

Web serving and content management

Software development environments

Backup and disaster recovery

HA SQL Server databases

Designed for a wide spectrum of use cases…

Cloud-native workloads

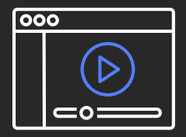

Media workflows

Analytics

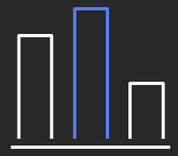

HPC

Dev and Test environments

SaaS applications

Amazon FSx availability

More coming soon!

US

US West (Oregon)
US West (N. California)
US East (N. Virginia)
US East (Ohio)

EU

Europe (Ireland)
Europe (Frankfurt
Europe (London)
Europe (Stockholm)

APAC

Asia Pacific (Sydney)
Asia Pacific (Singapore)
Asia Pacific (Tokyo)

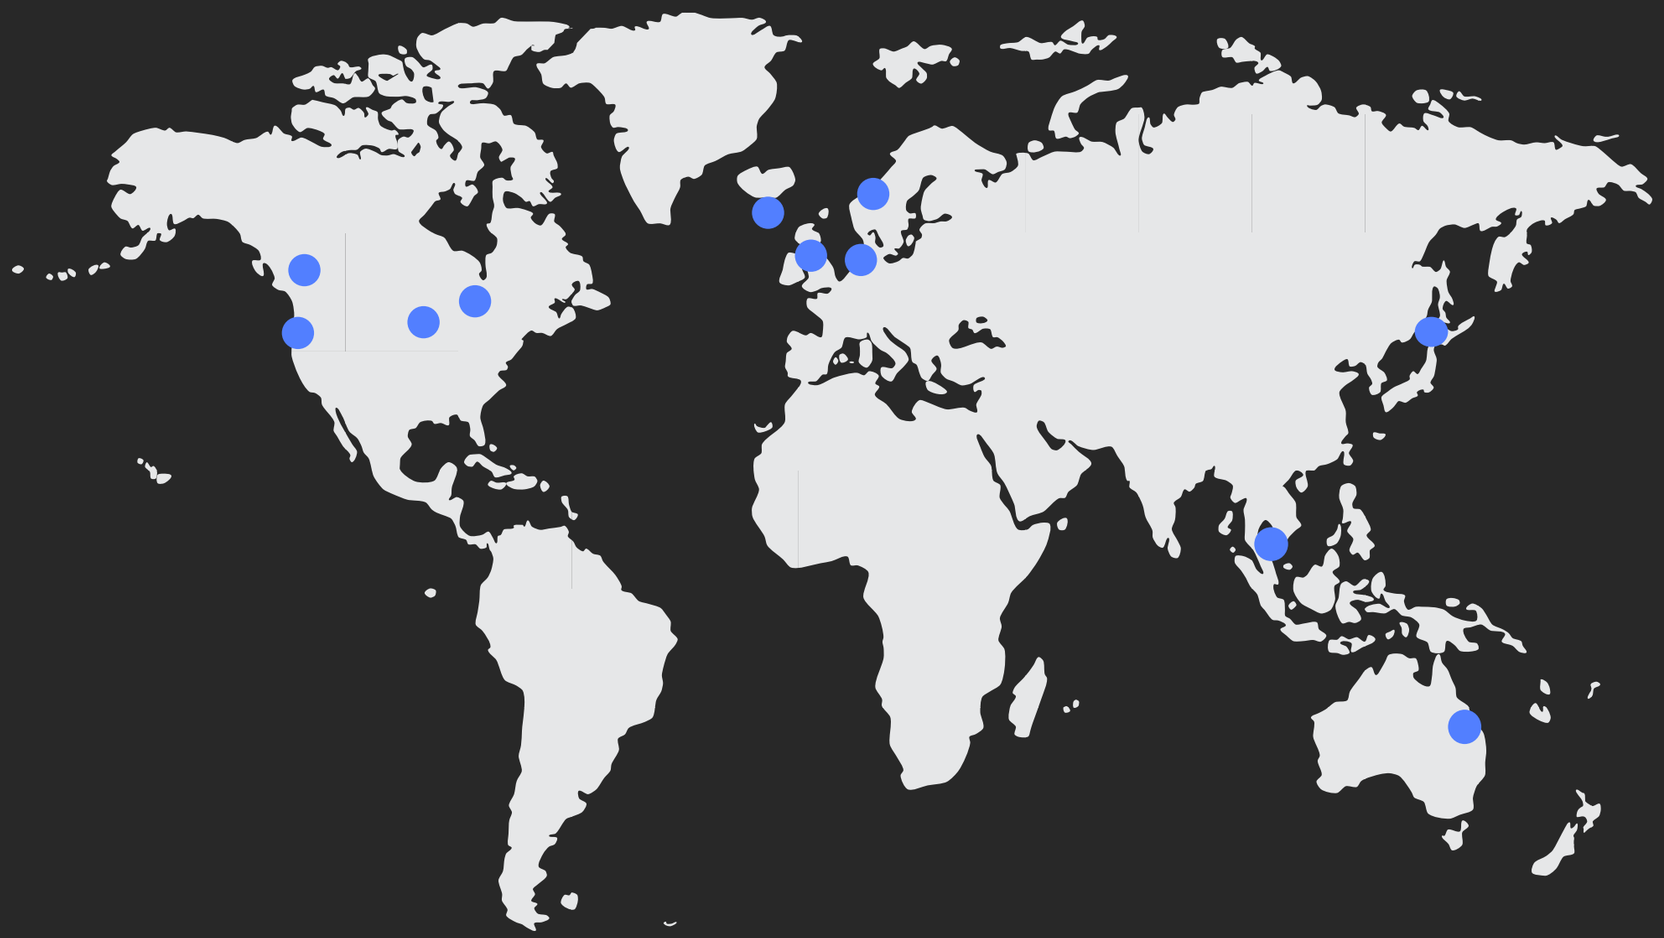

FSx for Windows File Server Pace of Innovation

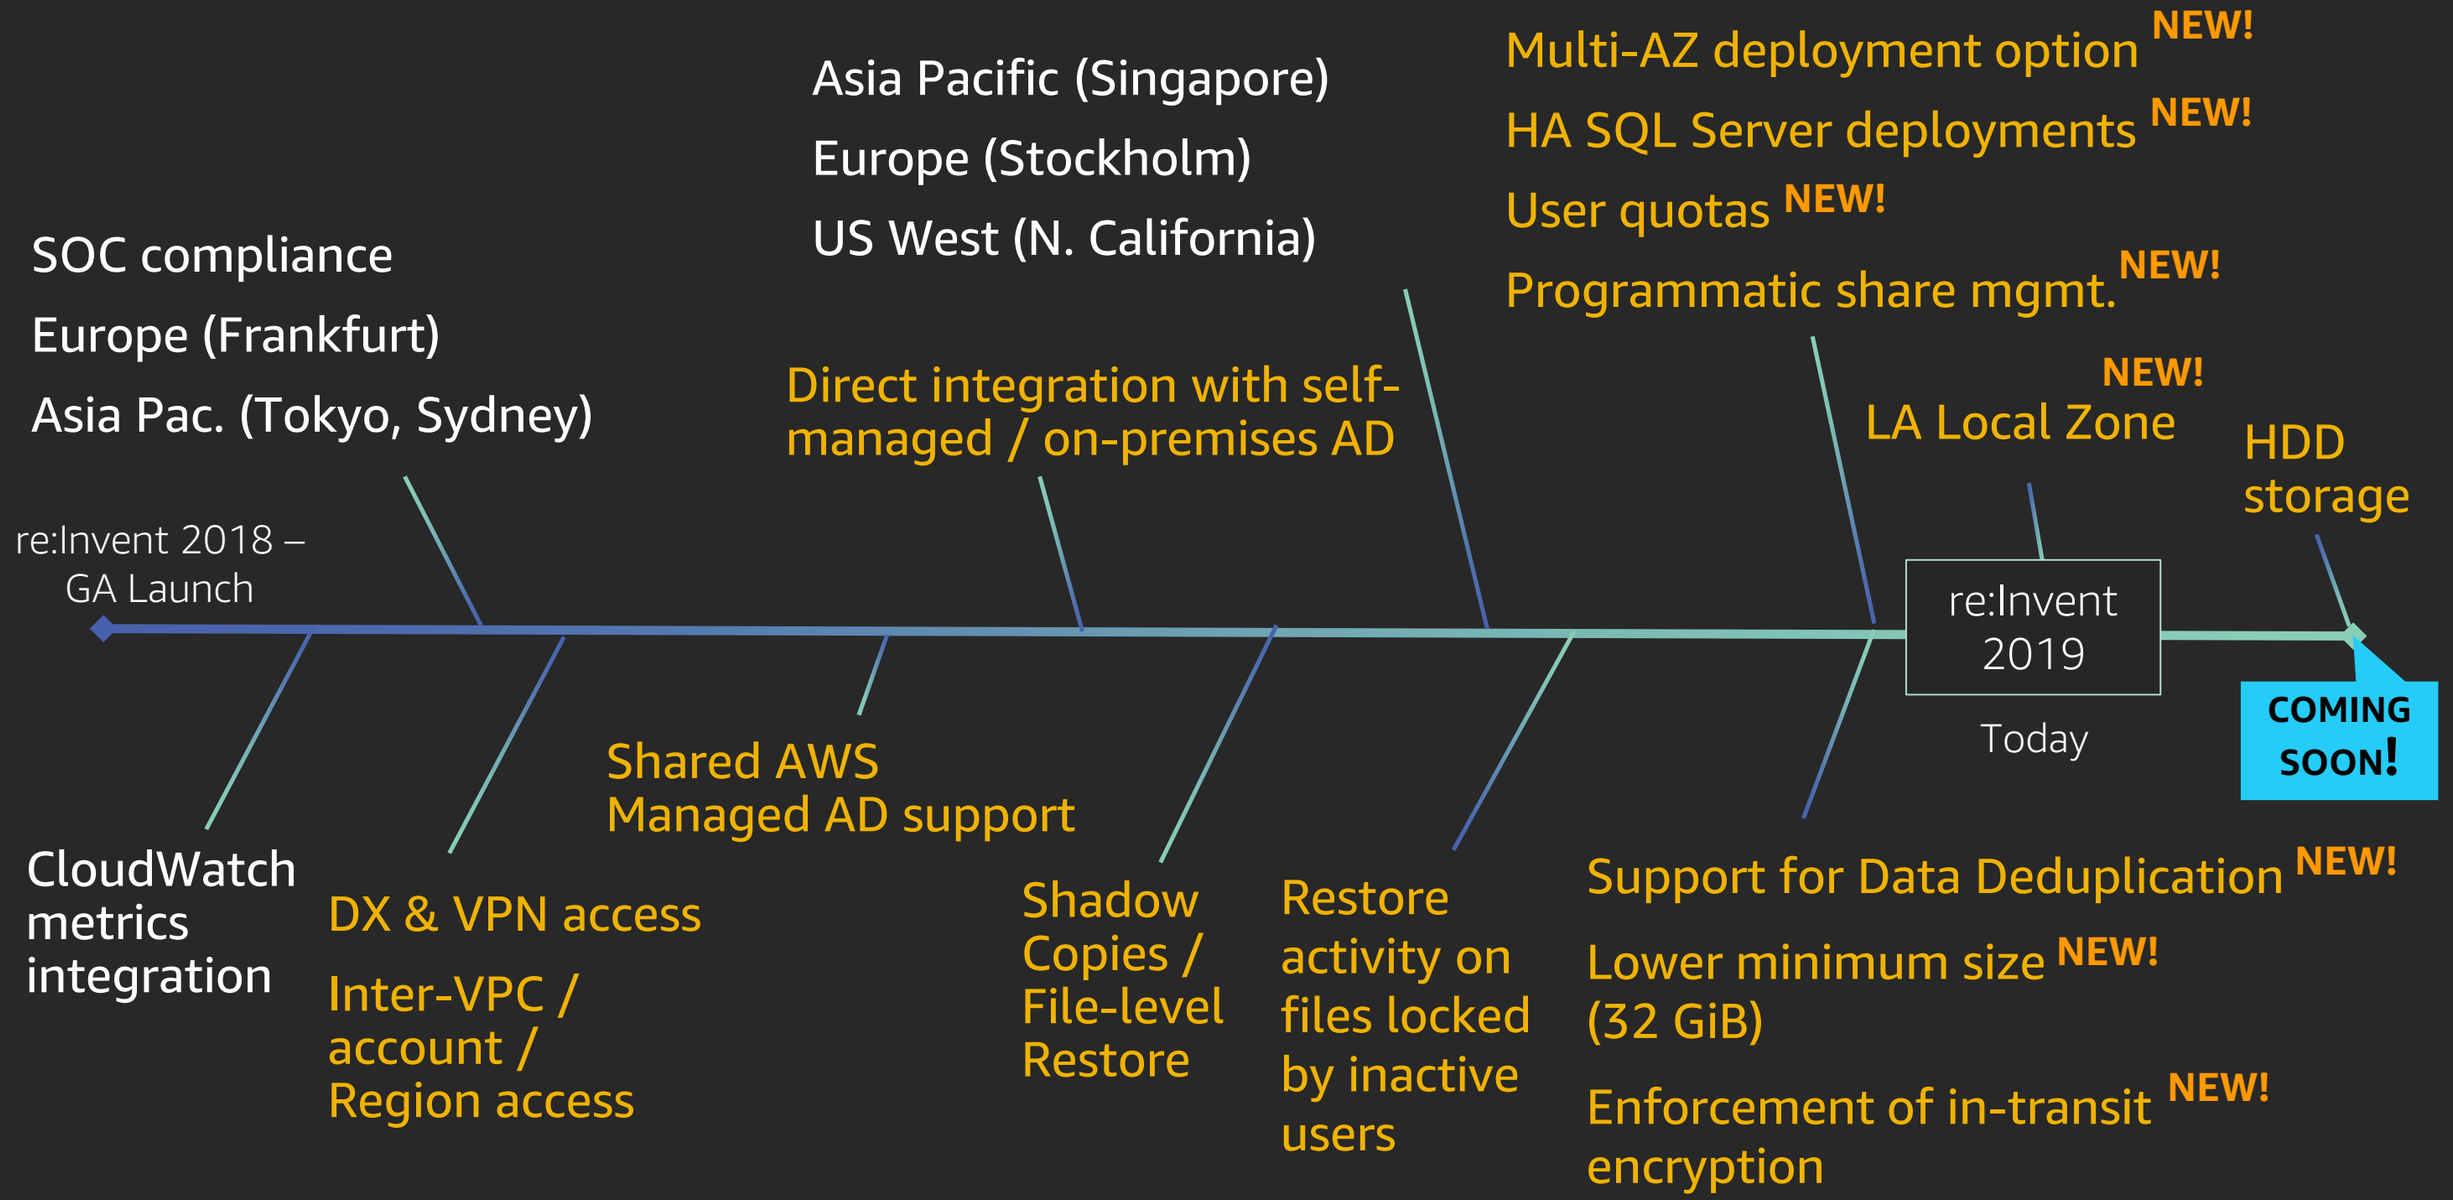

A glimpse of customers using Amazon FSx

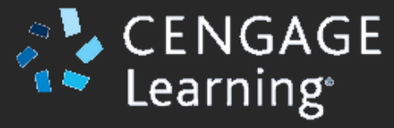

Amazon FSx partners

Deep dive

Six areas we will dive deep on today

Availability and durability

Performance and scale

Cost-effectiveness

Security and data protection

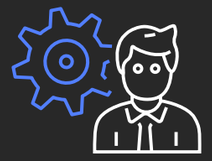

Administrative tools

Migration

Availability and durability

File system deployment options

Single-AZ

Continually monitors and addresses hardware failures

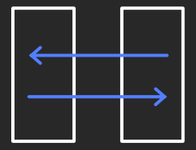

Replicates data within Availability Zone

Multi-AZ

NEW!

Continually monitors and addresses hardware failures

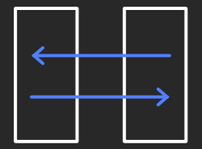

Replicates data within Availability Zone

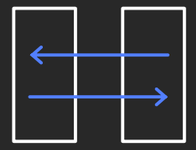

Replicates data across Availability Zones

Automatically fails over across Availability Zones

Single-AZ file system architecture

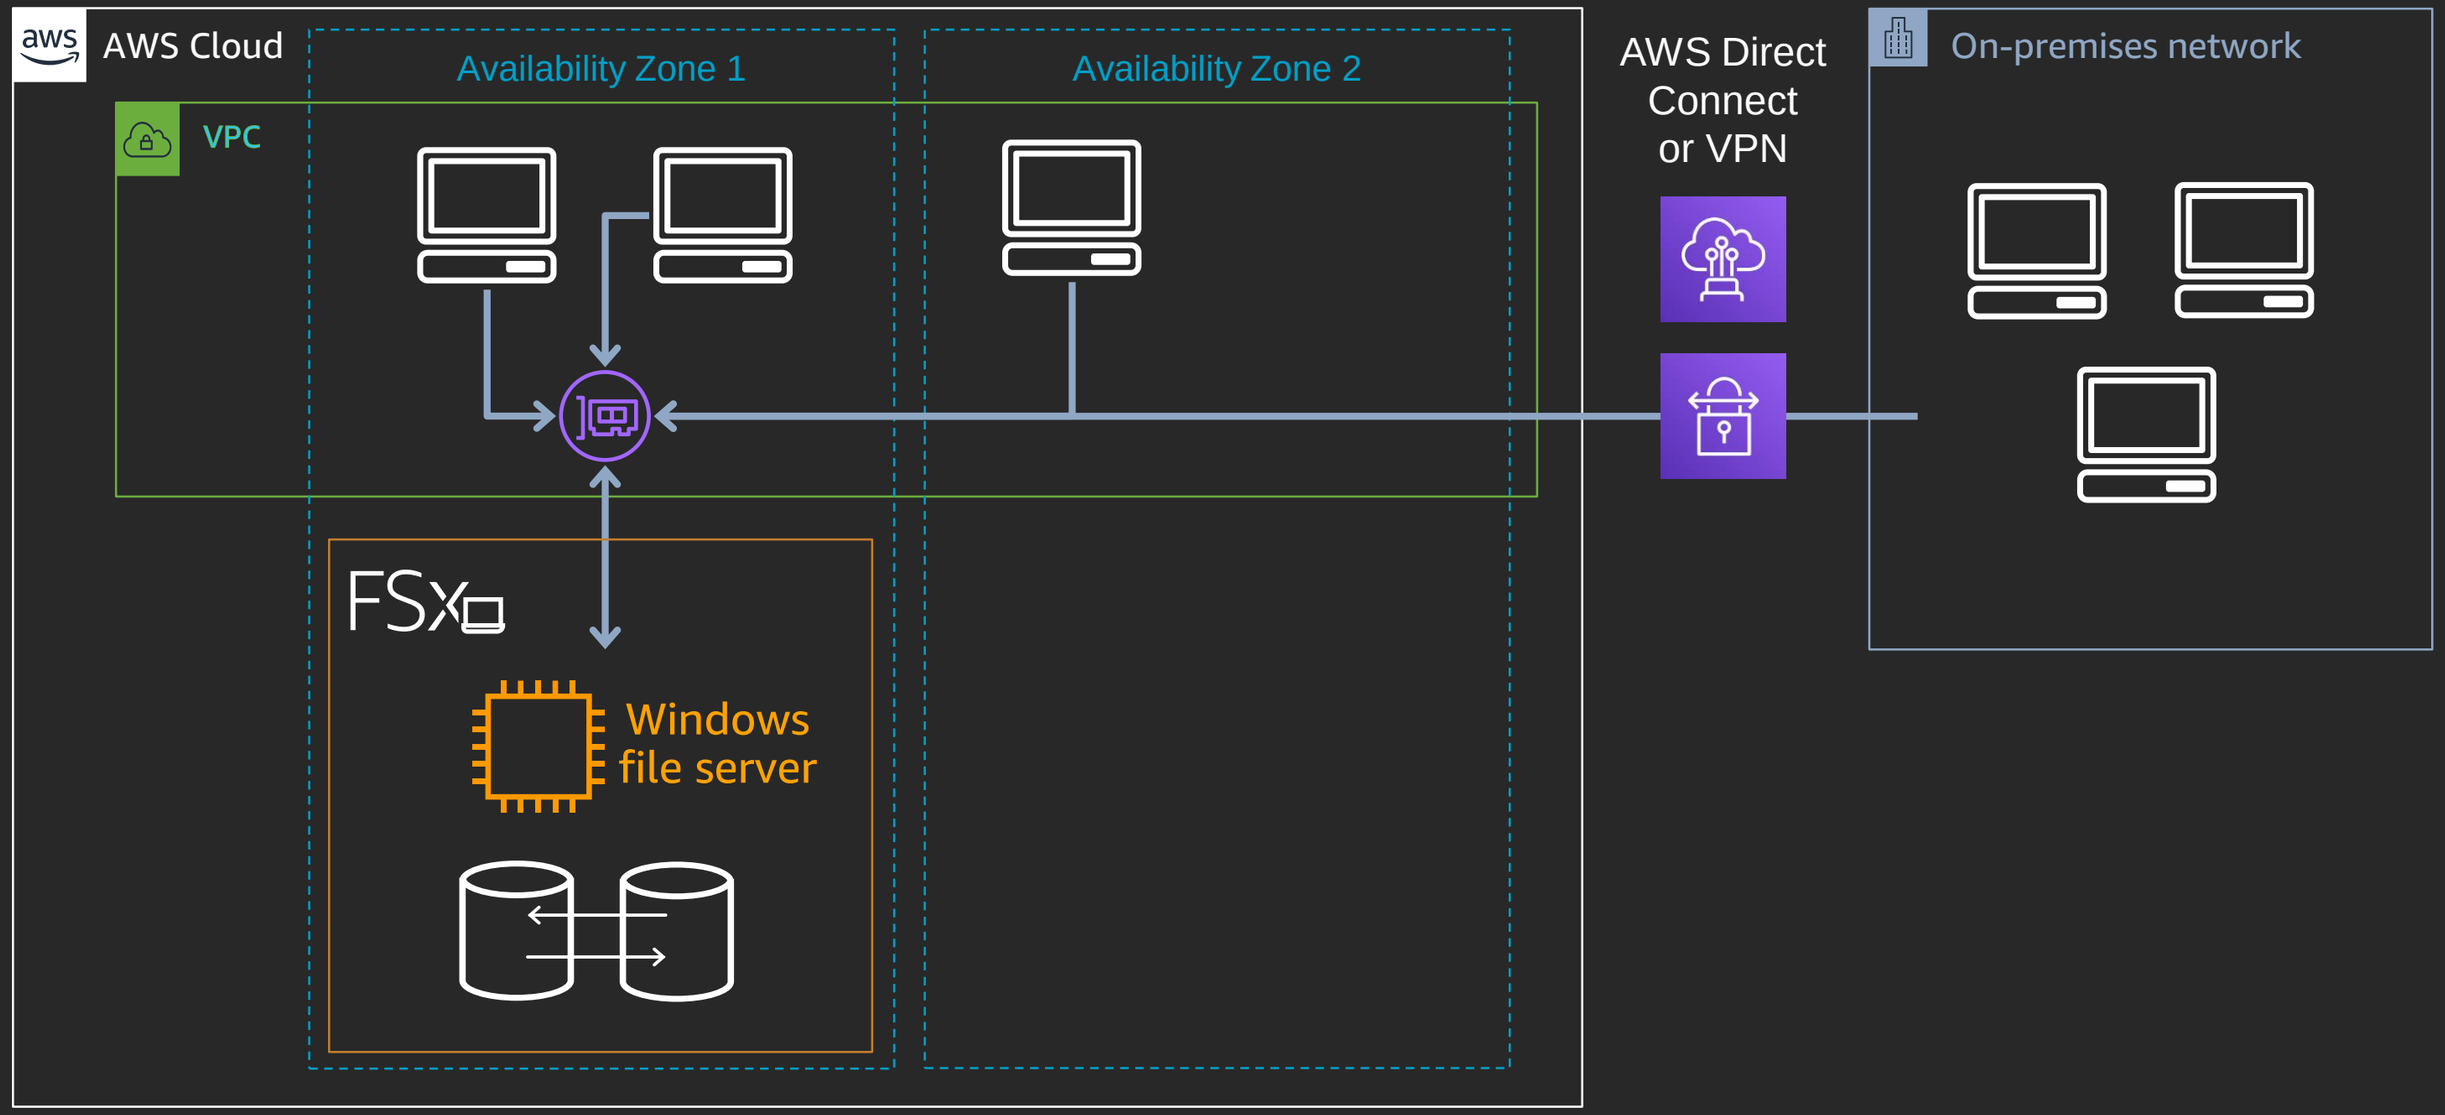

Multi-AZ file system architecture

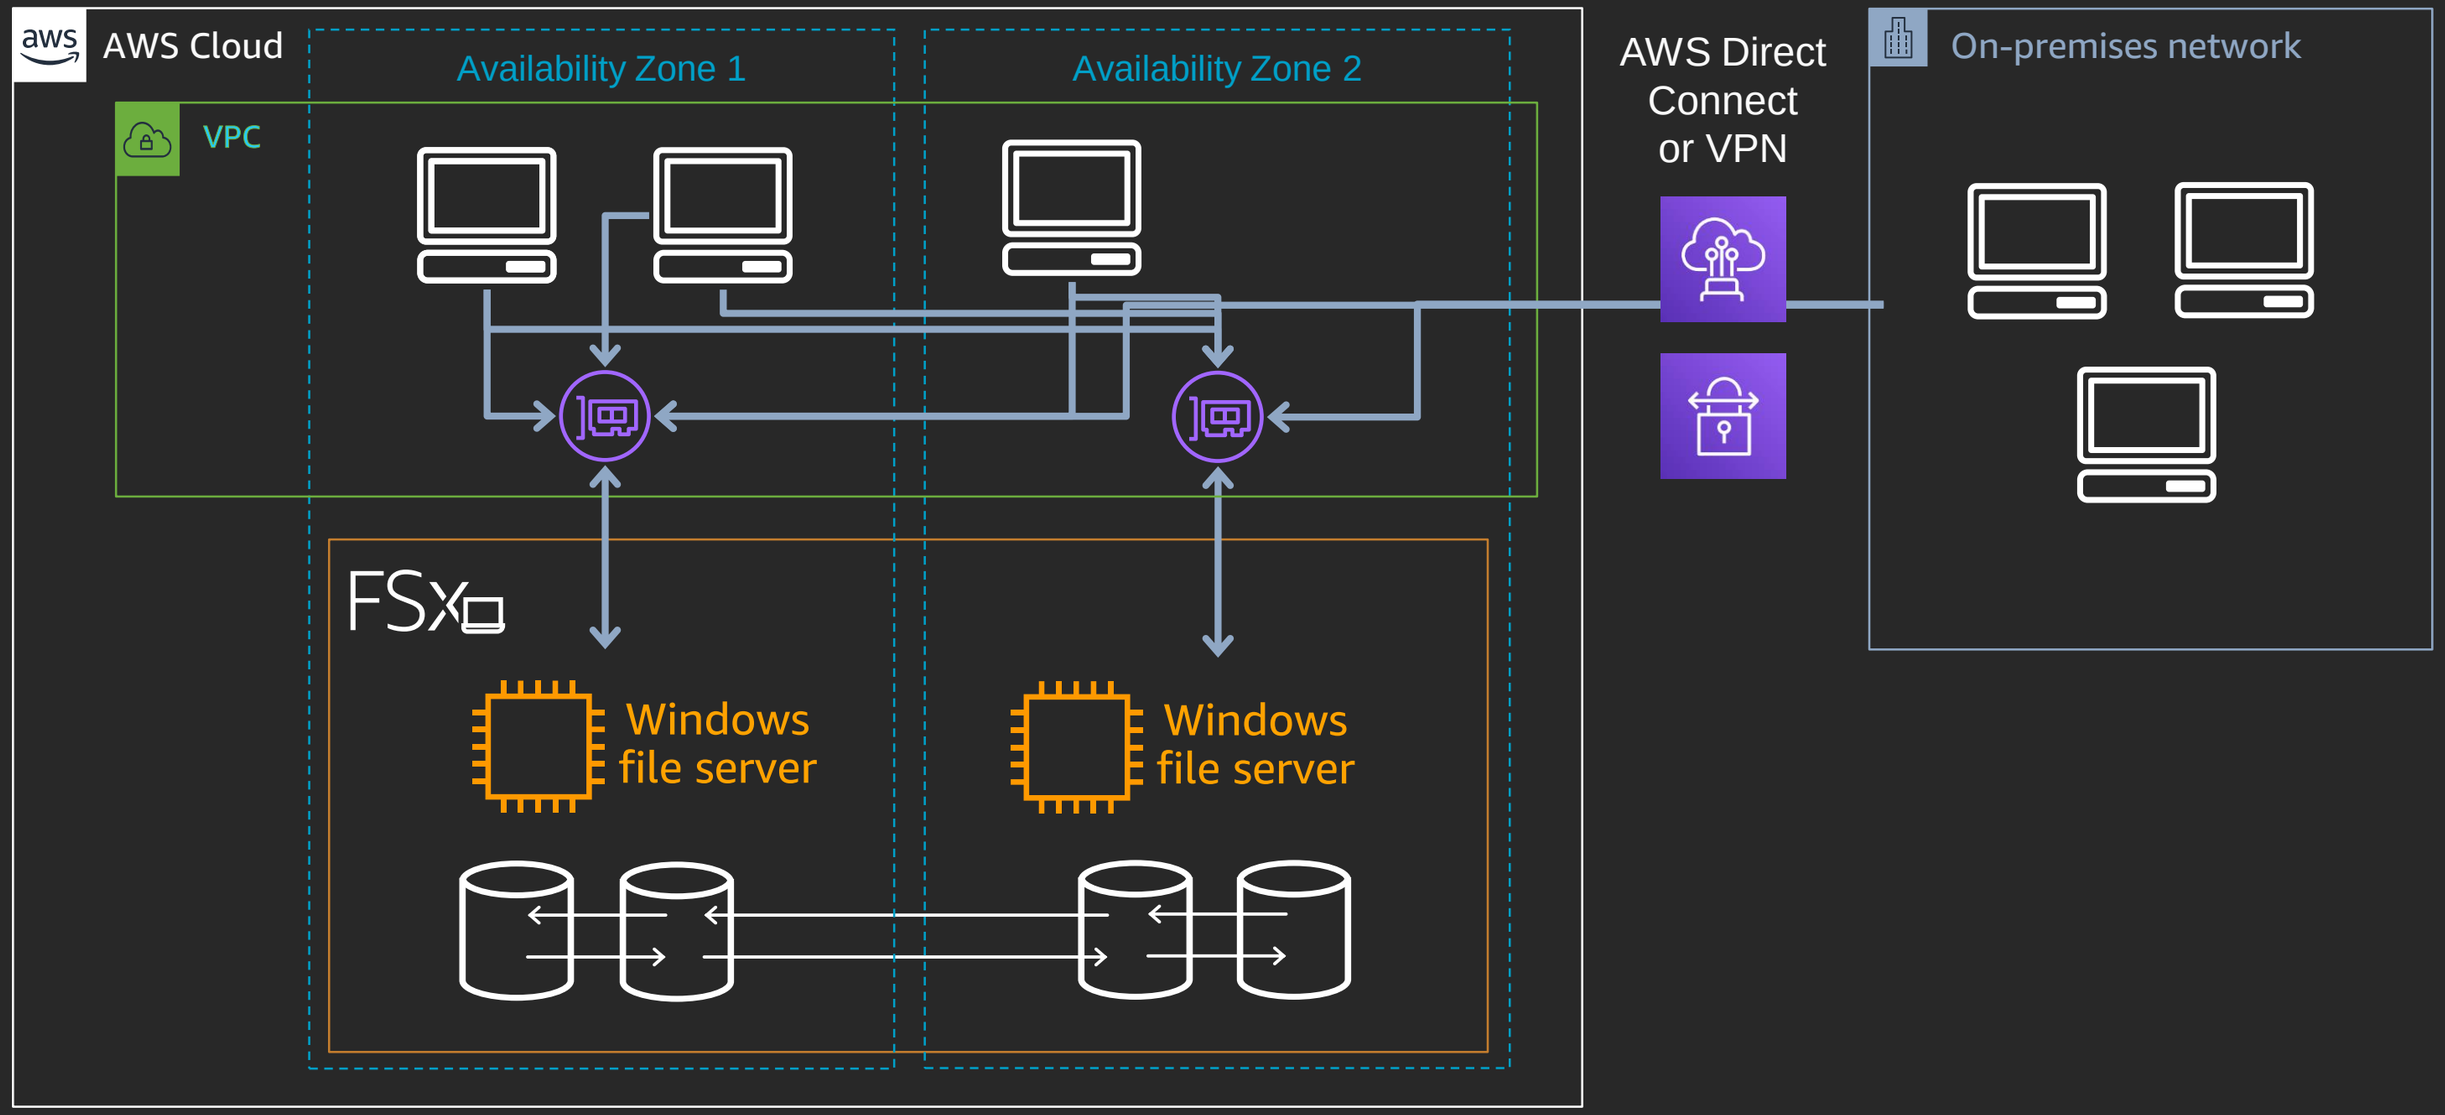

Events that trigger a Multi-AZ failover

- Loss of availability to preferred file server
  - Unavailability of preferred file server’s Availability Zone
  - Unavailability or failure of preferred file server
- Planned maintenance

Support for SQL Server HA deployments

- Supports SMB Transparent Failover (aka Continuously Available shares)
- Use Amazon FSx to store databases and logs for SQL Server Always On Failover Cluster Instance (FCI) deployments
- No need to deploy, manage, and pay license fees for storage replication software solutions

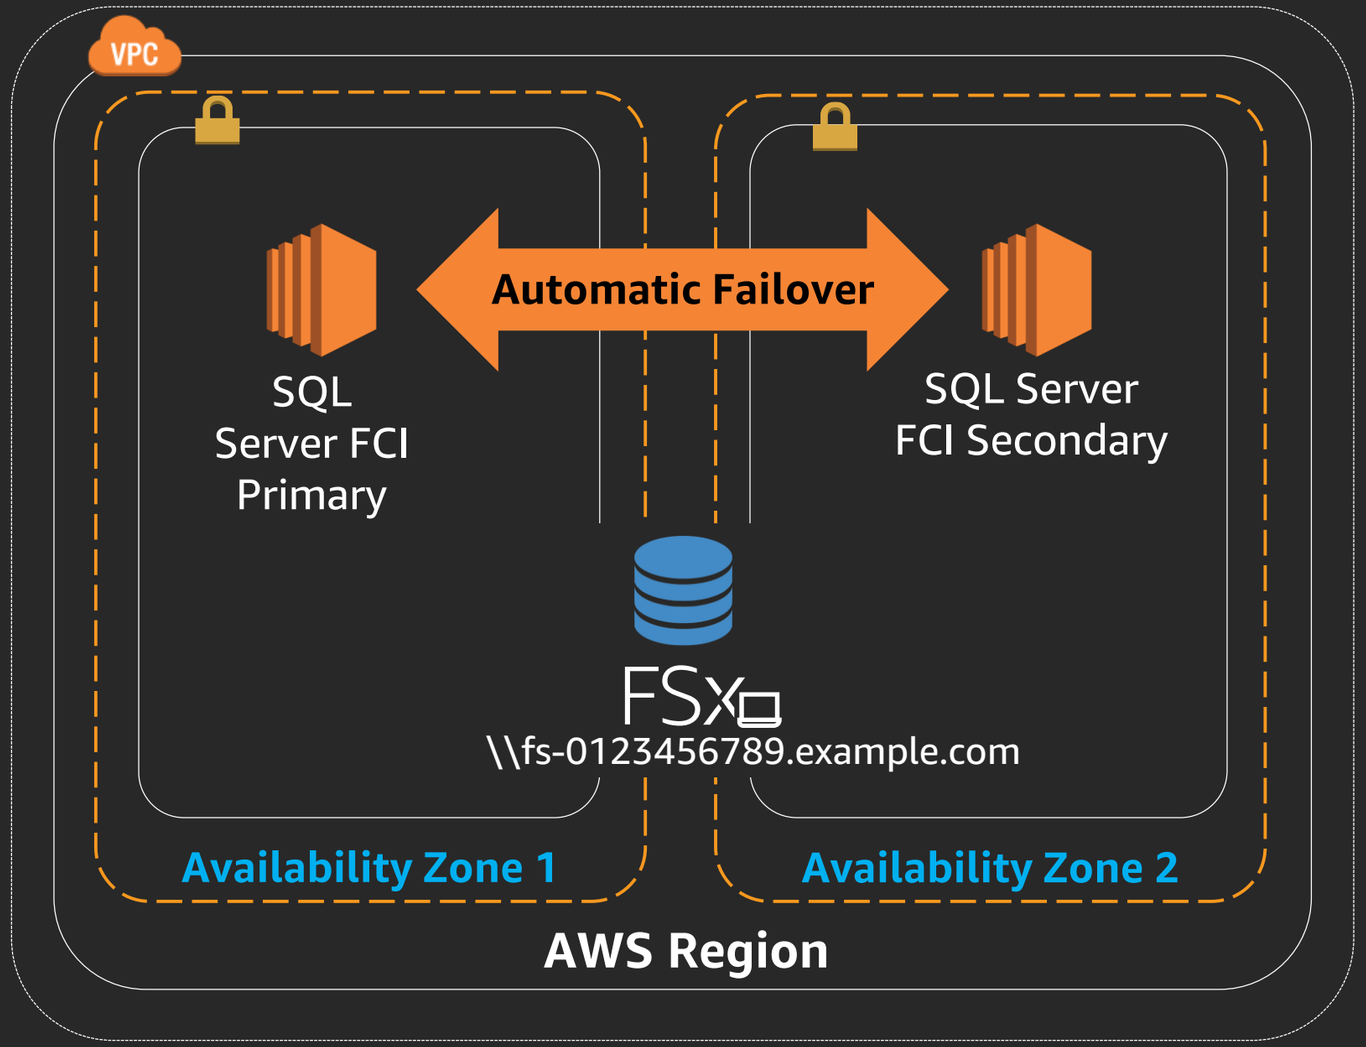

Backups

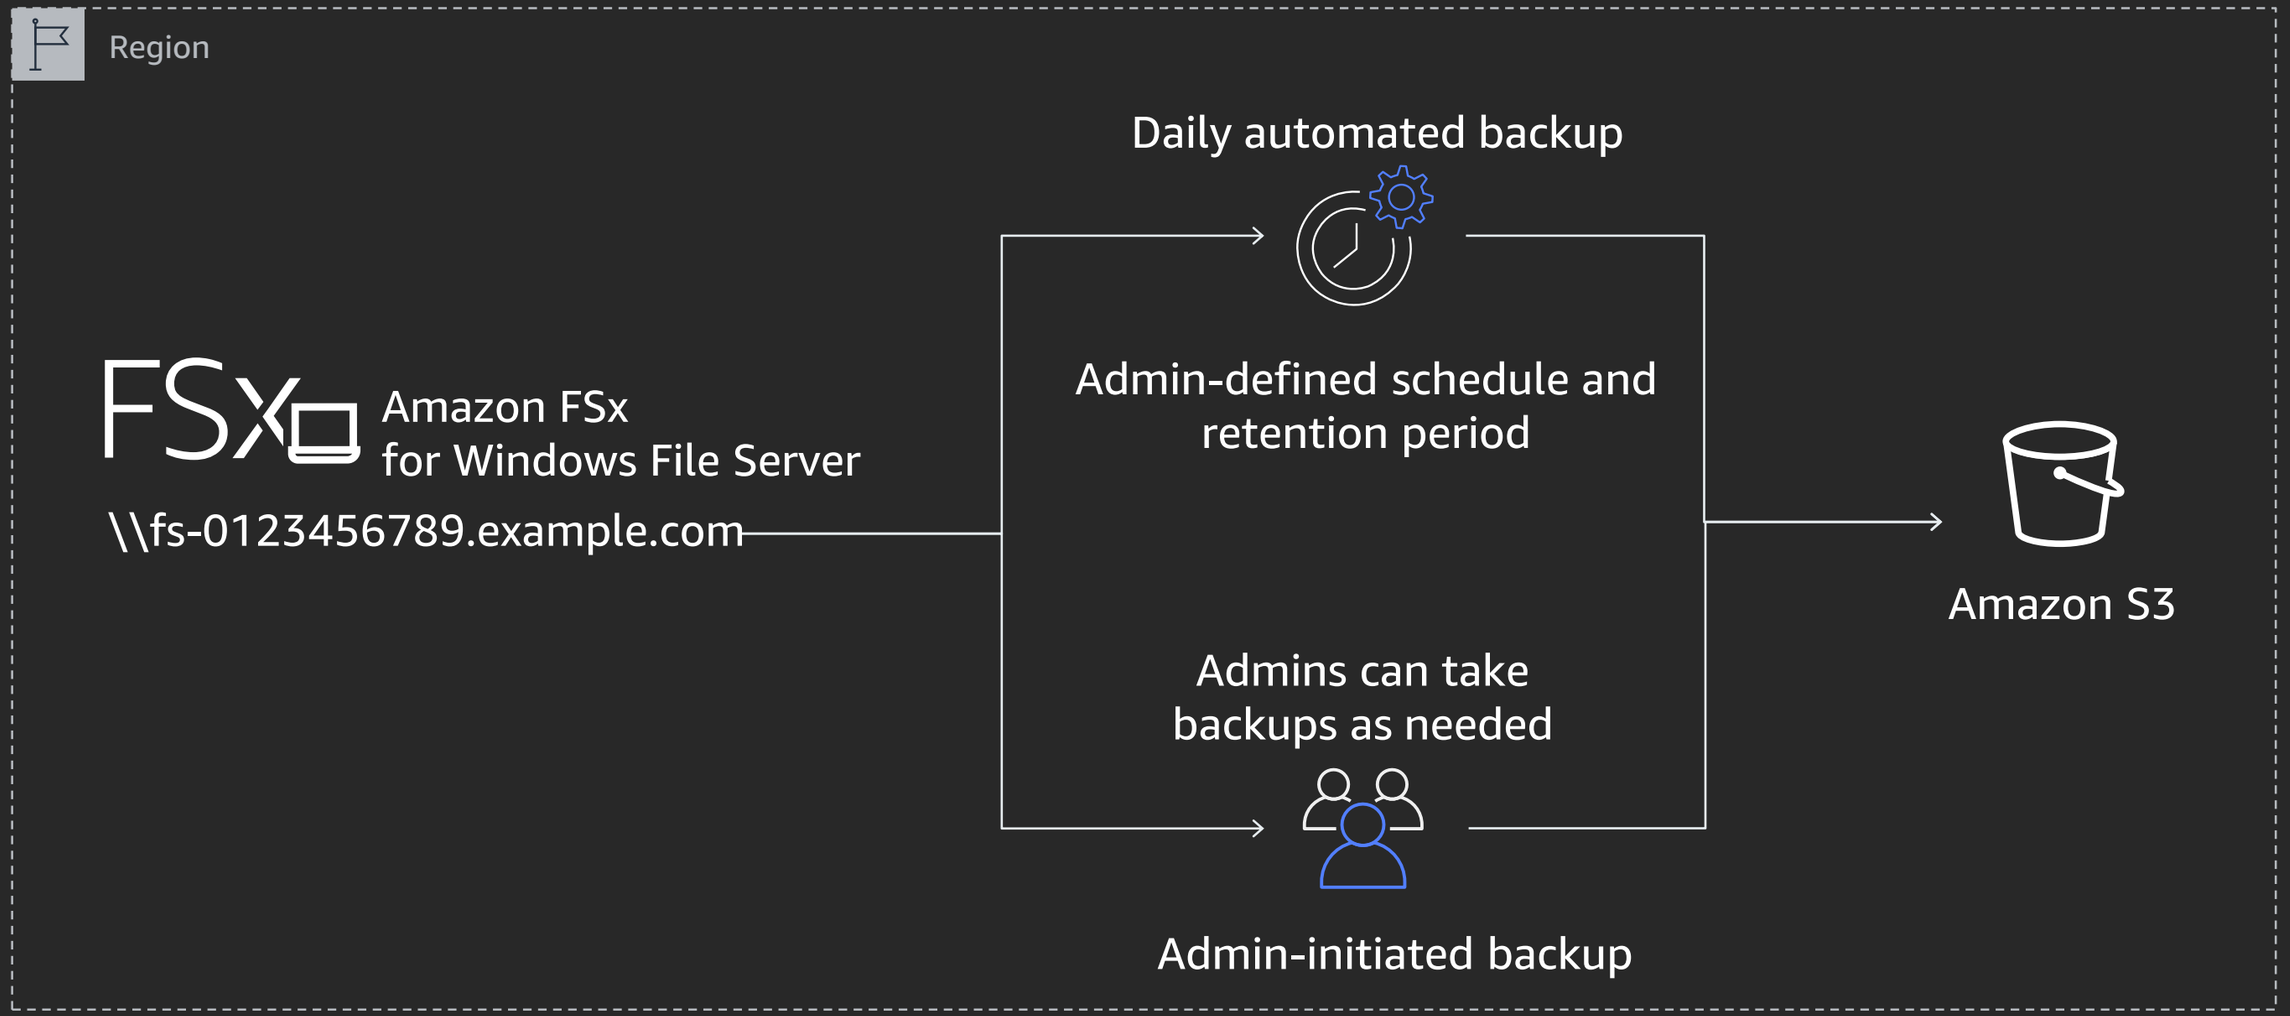
Region
Daily automated backup
Amazon FSx
for Windows File Server
\\fs-0123456789.example.com
Admin-defined schedule and
retention period
Amazon S3
Admins can take
backups as needed
Admin-initiated backup

Creating a custom backup schedule

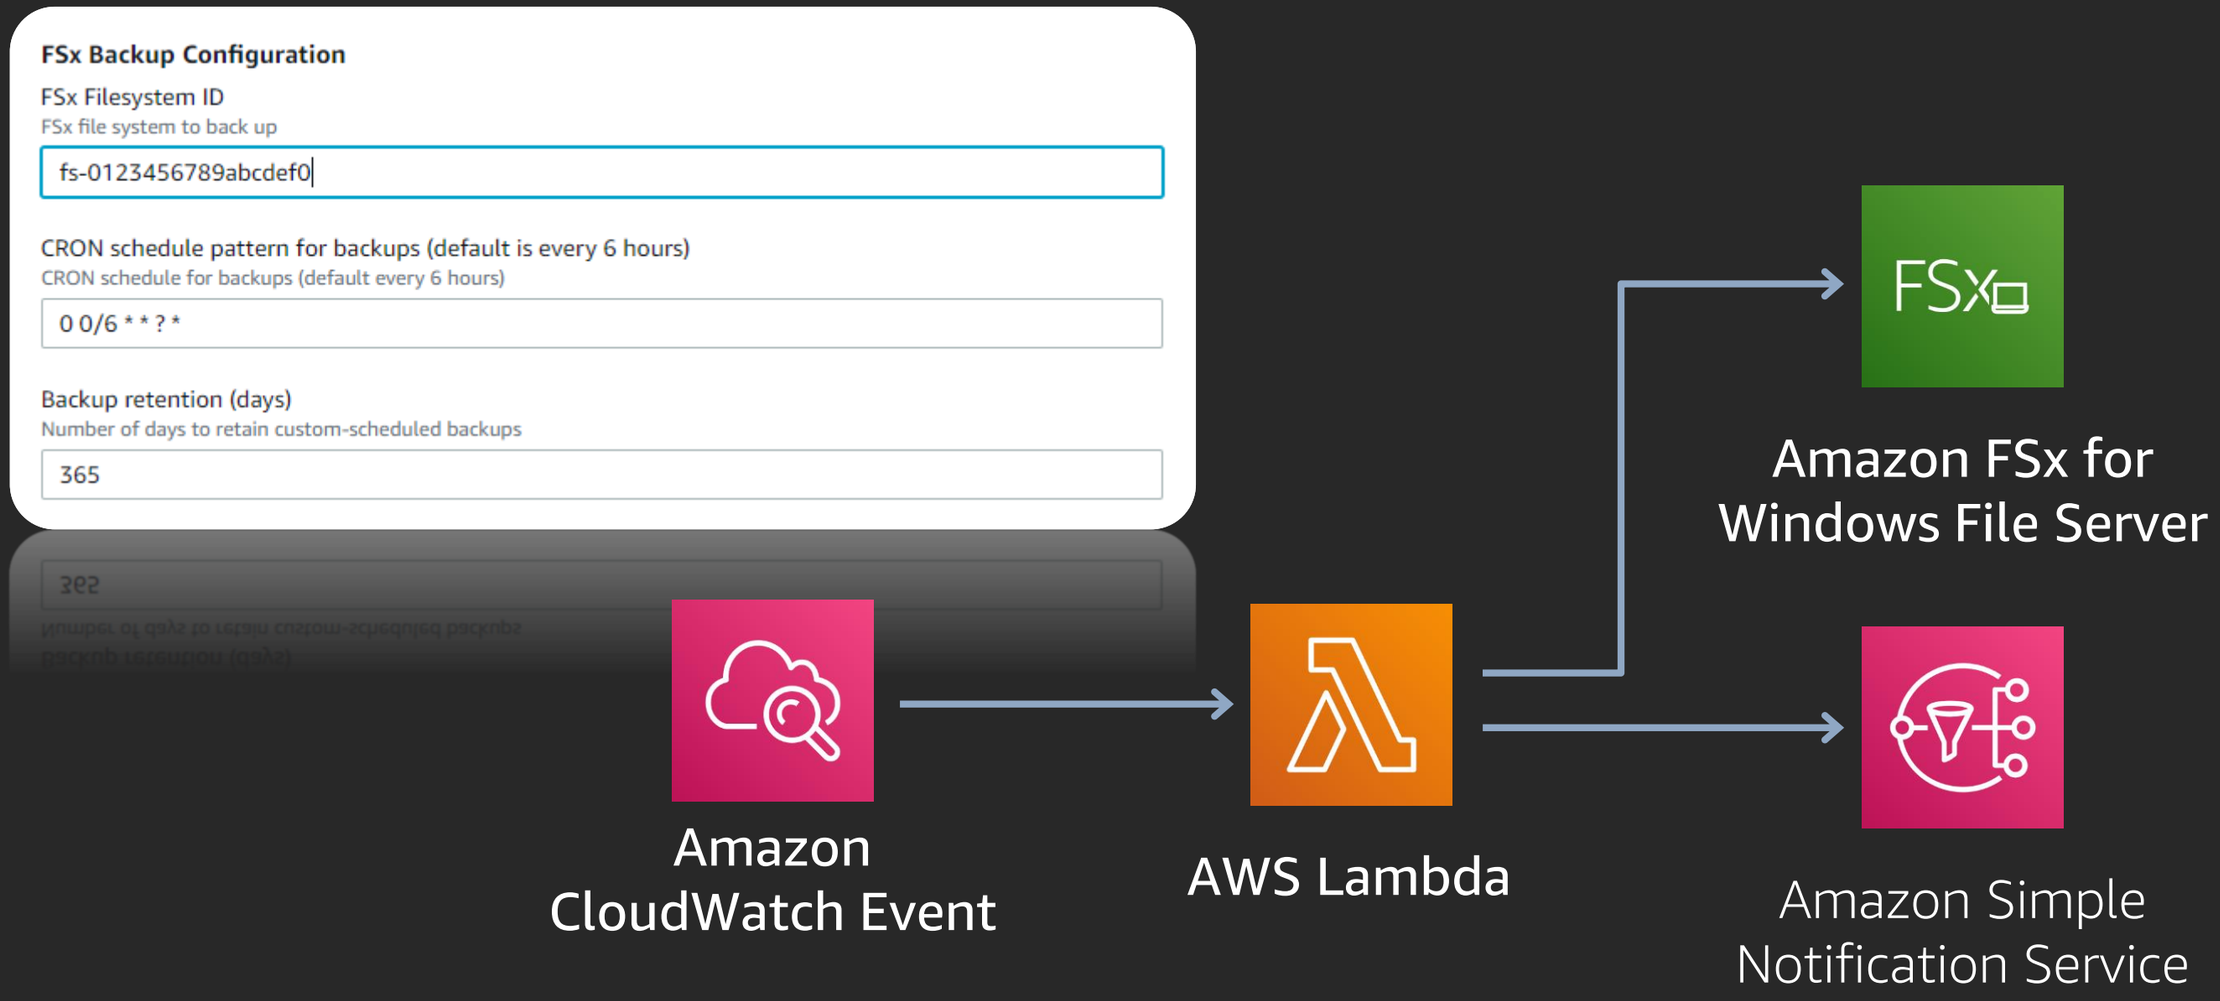

Performance and scale

Performance and scale

Latency

Sub-millisecond latencies with SSD

Throughput and IOPS

Direct file server access: up to 3 GB/s of throughput and hundreds of 1000s of IOPS per file system

With client-side caching: up to 10+ GB/s of throughput and millions of IOPS per file system

Single-client performance

With SMB Multichannel, a single client can drive up to the full throughput/IOPS of a file system

Throughput capacity

- Determines speed at which the file server hosting your file system can serve file data
- Higher levels of throughput -> higher levels of IOPS and more memory for caching

Throughput capacity / Baseline throughput (MBps)	**Burst throughput** (MBps)
8	192
16	192
32	192
64	256
128	438
256	438
512	N/A
1,024	N/A
2,048	N/A

You get even higher throughput with in-memory caching on the file server (600 MBps – 3 GBps)

How do you pick the right level of throughput for your file system?...

Throughput capacity is automatically picked for you – sufficient for vast majority of apps

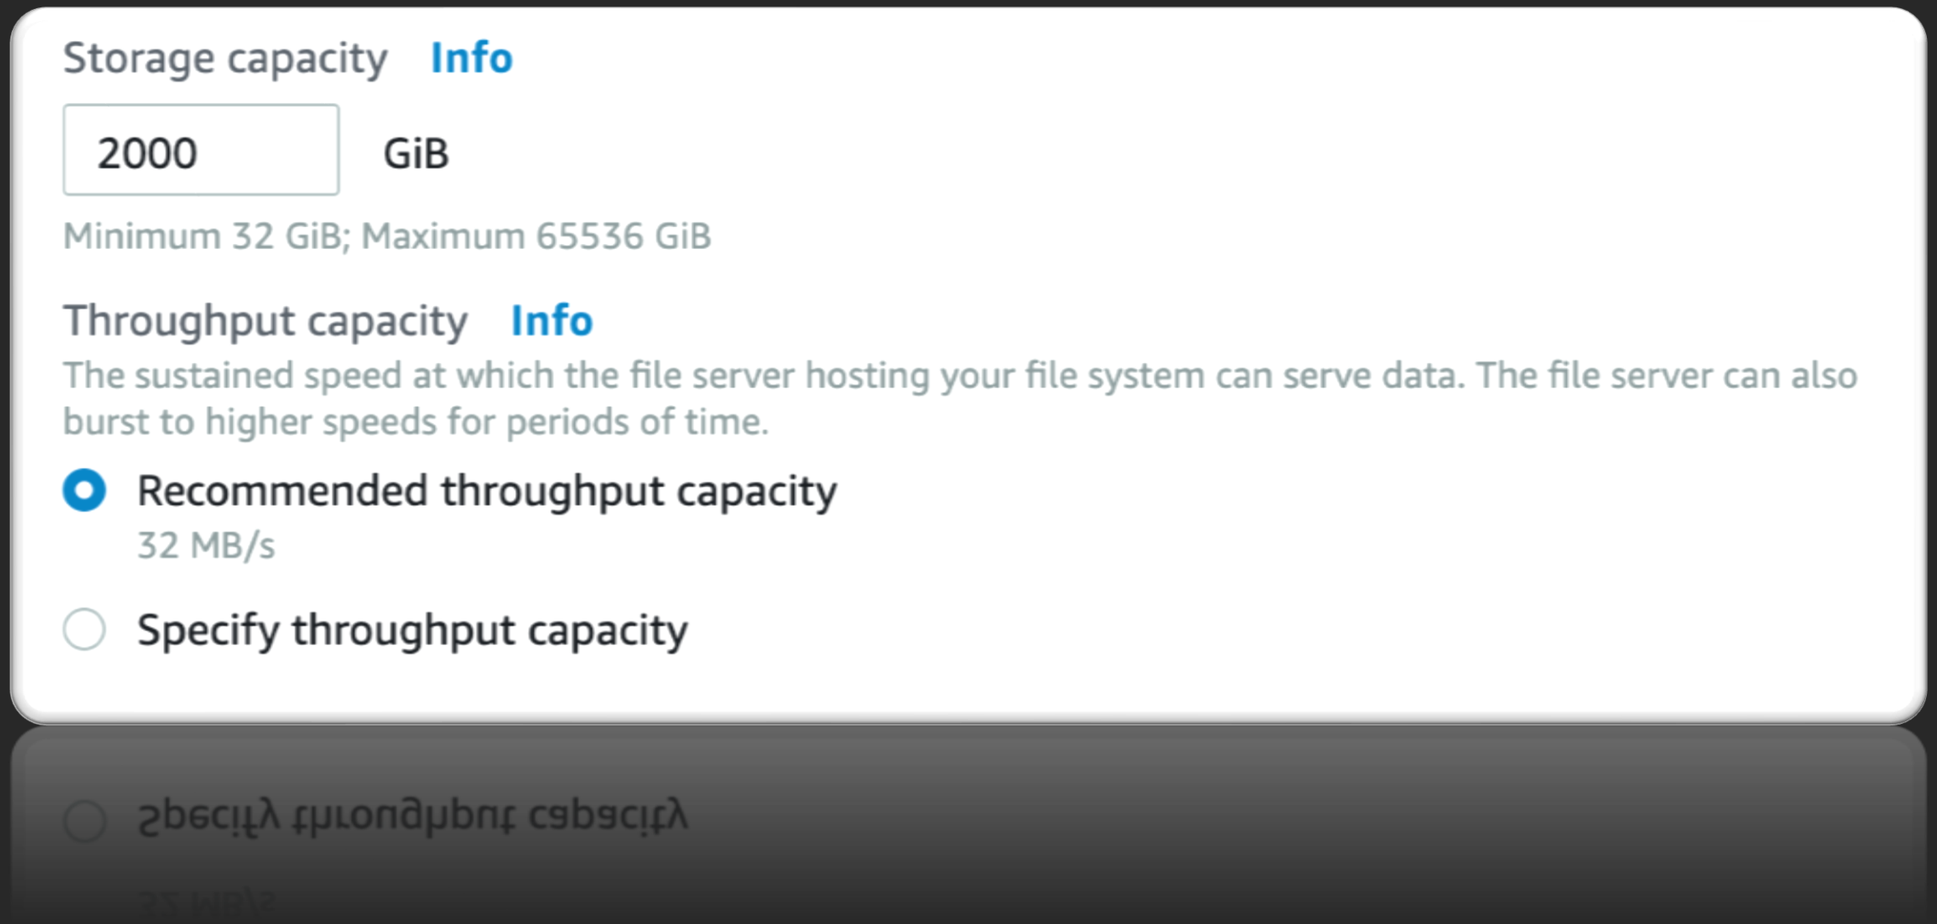

If you need more than the default…

1. Test your workload with Amazon FSx
2. Look at CloudWatch metrics for your Amazon FSx file system
3. If the total throughput is approaching the throughput capacity limit, pick a higher level

CloudWatch metrics

Every minute, Amazon FSx emits metrics to Amazon CloudWatch:

- Free storage capacity
- File system throughput (read & write)
- File system operations (read, write, and metadata)

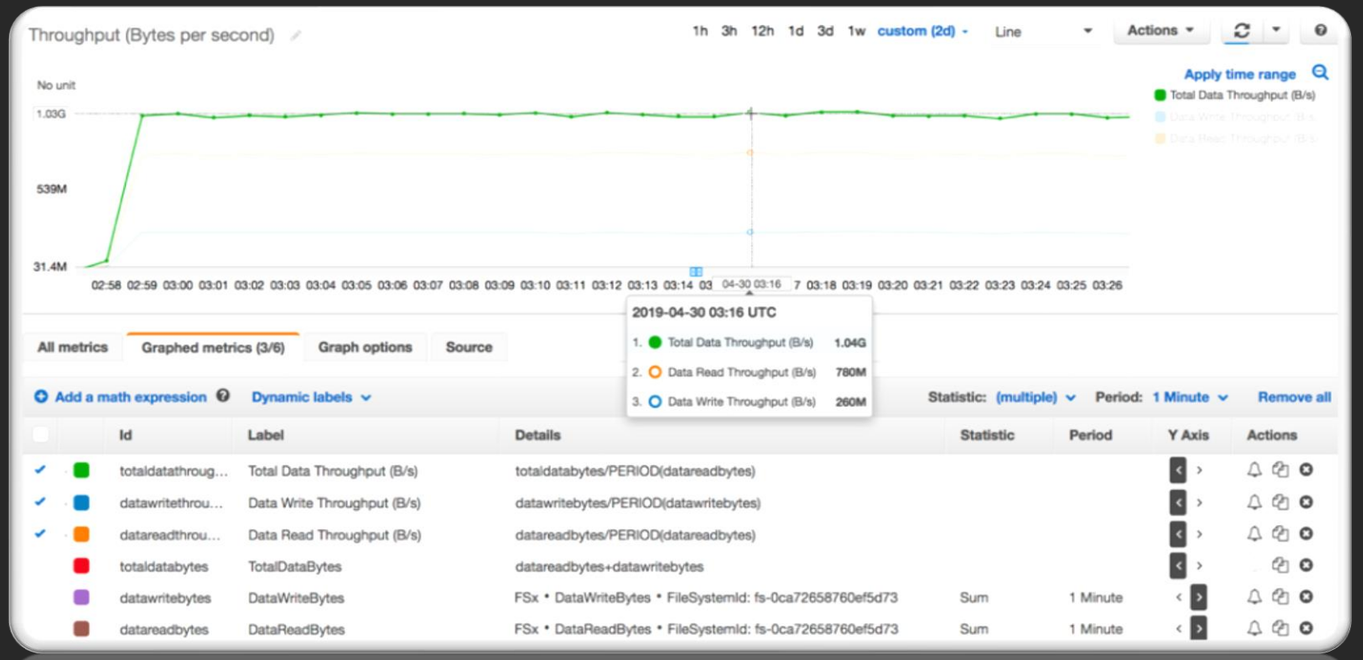

You can use Metric Math to simplify near real-time monitoring of your file system

Scaling out storage & performance with DFSN

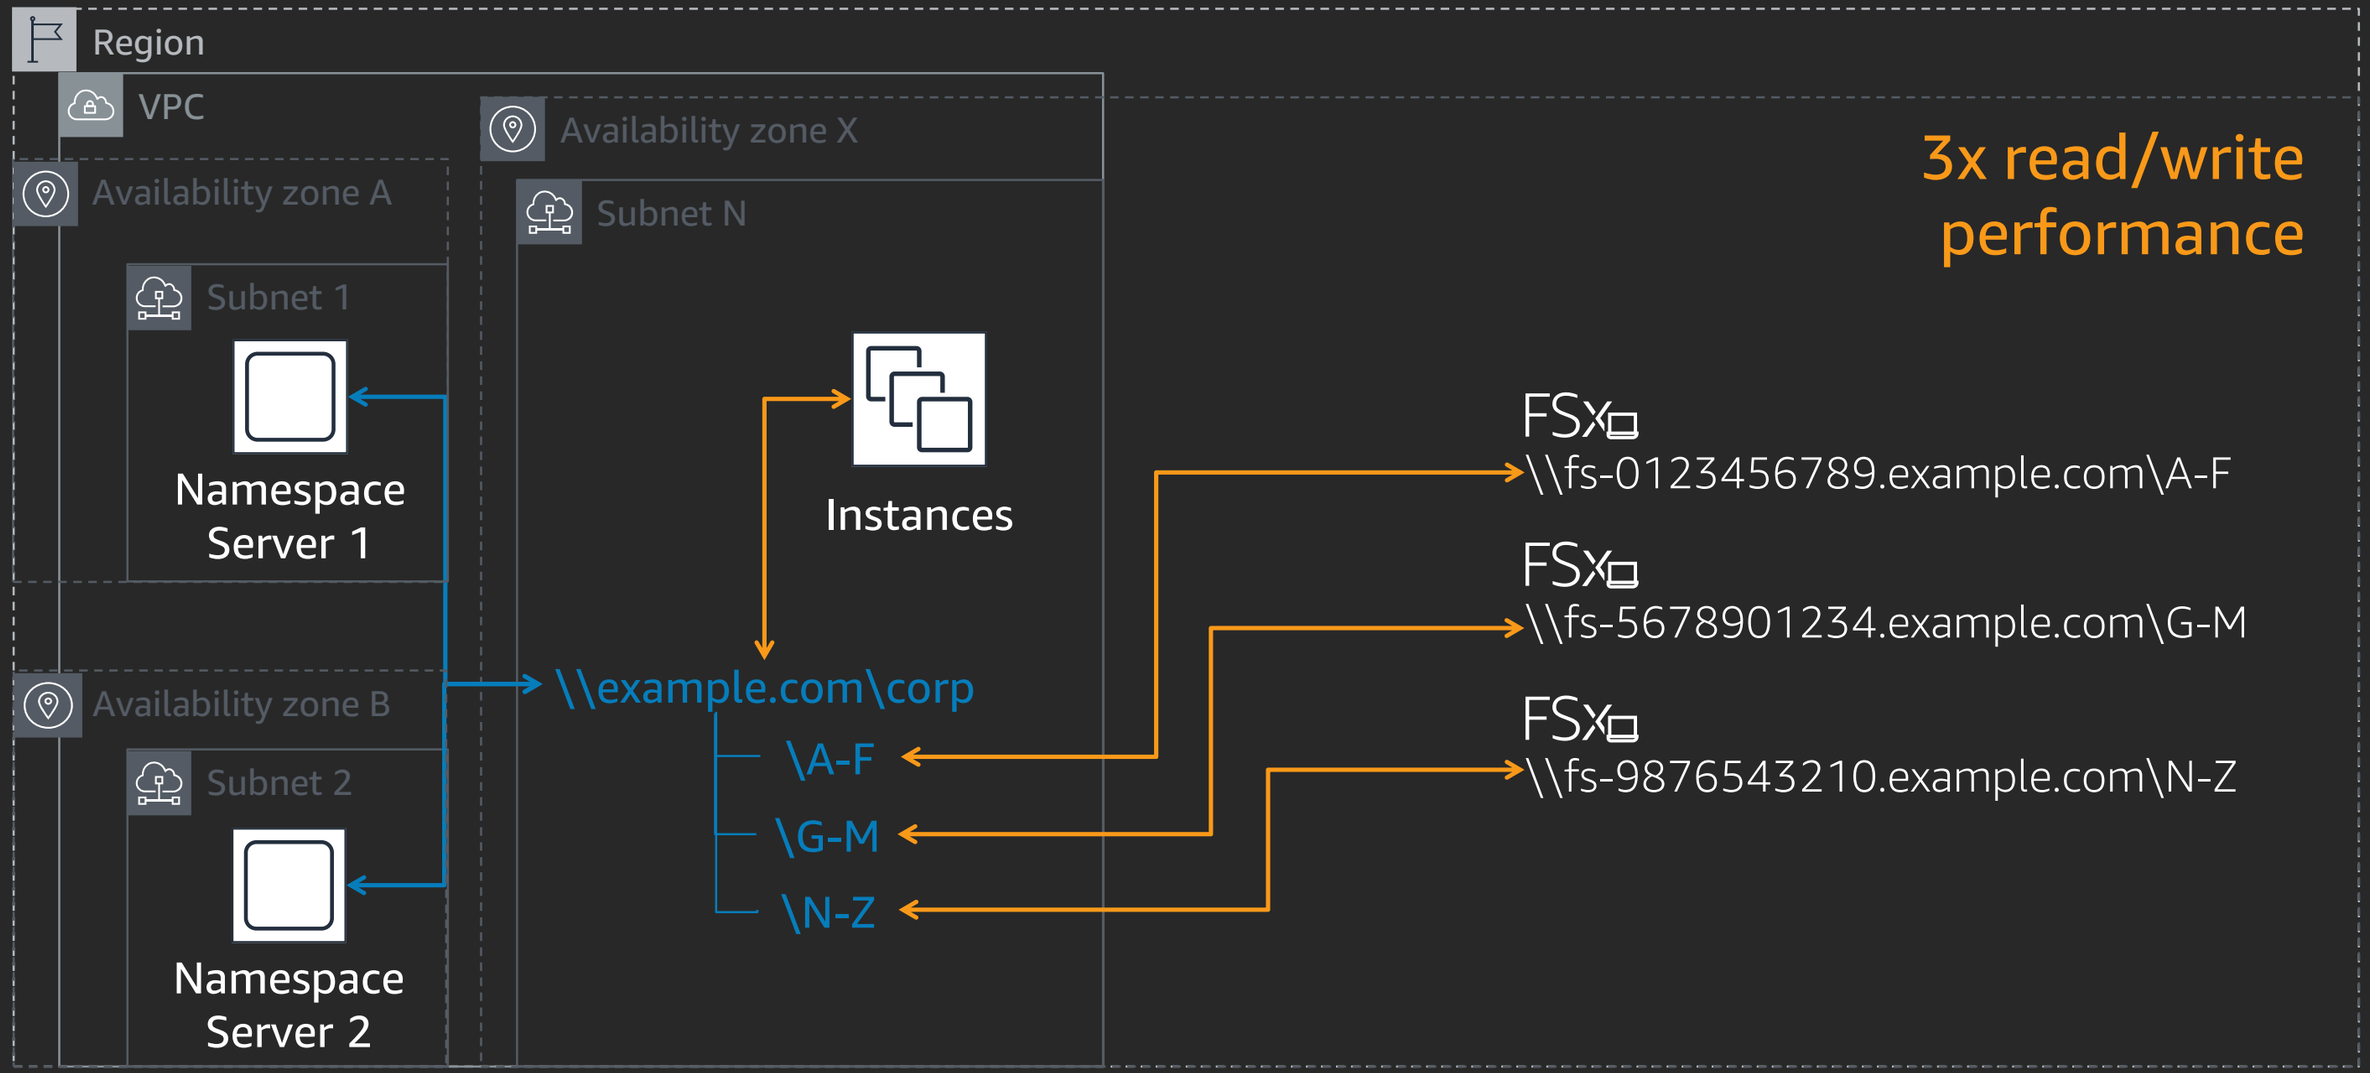

Cost-effectiveness

Choose the ideal storage type

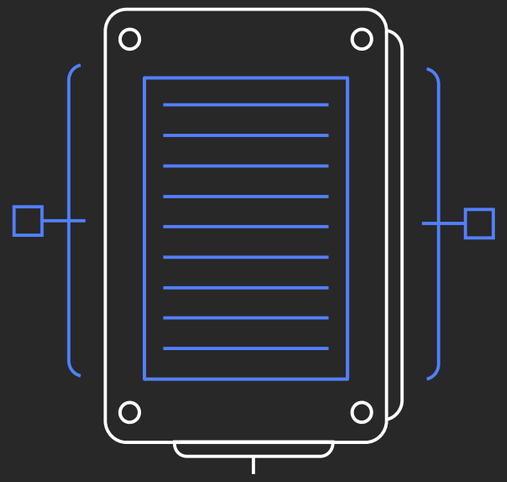

SSD

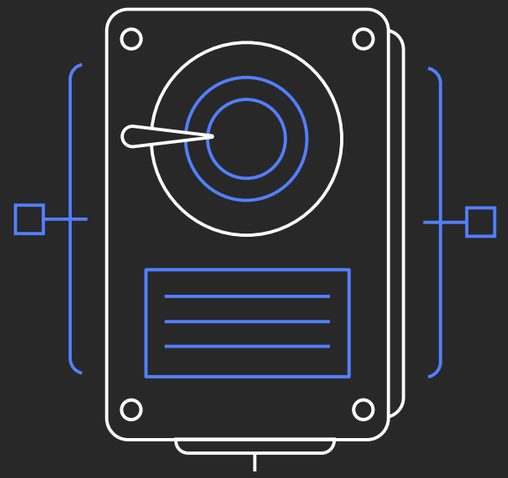

HDD

Data Deduplication

Large datasets often have a lot of duplication, which increases storage costs

User shares (home directories)
Multiple users have many copies or versions of a file

Software dev shares
Most portions of binaries remain unchanged from build to build

Use Data Deduplication to reduce costs associated with duplicated data

Scenario	Content	Typical space savings
User documents	Office documents, photos, music, and videos	30-50%
Software dev shares	Software binaries, build files, and program symbols	70-80%
General file shares	Mix of the above	50-60%

Data Deduplication

Works at the sub-file level

Uses post-processing optimization to minimize performance impact

Removes duplicated content and compresses common content

Use remote management PowerShell CLI on your file system to…

- Enable/disable Data Deduplication
- Customize schedule for deduplication jobs
- Monitor how much savings you're achieving with deduplication

```
PS C:\Windows\system32> Invoke-Command -ComputerName amznfsx        .fsx-ad.com -ConfigurationName FsxRemoteAdmin -ScriptBlock {
>> Get-FSxDedupStatus } | select OptimizedFilesCount,OptimizedFilesSize,SavedSpace,OptimizedFilesSavingsRate

OptimizedFilesCount OptimizedFilesSize SavedSpace OptimizedFilesSavingsRate
------------------- ------------------ ---------- -------------------------
                779         1666284377  929007343                        55
```

User Storage Quotas

- Monitor and control user-level storage space consumption on file systems
- Use remote management PowerShell CLI on your file system to…
  - Enable/disable user quotas
  - Configure quota levels for the default and for specific users or groups
  - Choose between *track* and *enforce* modes
  - List current quota violations

```
PS C:\Users\Admin>
PS C:\Users\Admin> Invoke-Command -ComputerName $PS_ENDPOINT -ConfigurationName FSxRemoteAdmin -ScriptBlock {
>>     Get-FSxUserQuotaEntries
>> } | Select User, Status, DiskSpaceUsed, Limit

User                                          Status   DiskSpaceUsed Limit
----                                          ------   ------------- -----
Win32_Account (Name = "Admin", Domain = "dom") Exceeded 22412234752   2147483648
```

Example TCO

Storage requirements

- 10 TB of storage
- With deduplication, 50% of storage needed
- Deployment type: Multi-AZ
- Storage type: HDD

Throughput requirements

- 16 MB/s sustained, 100 MB/s burst

Backup requirements

- Expected backup storage usage: 1x of storage capacity

File system component	Total cost
Storage (Multi-AZ, HDD, 5 TB @ $0.025/GB-mo)	$128
Throughput capacity (16 MB/s @ $4.50/MBps-mo)	$72
Total cost (excl. backups)	$200/month (or $0.02/GB-mo)
Backups (5 TB @ $0.05/GB-mo)	$256
Total cost (incl. backups)	$456/month (or $0.04/GB-mo)

Security and data protection

Security and compliances

Data encrypted
at-rest and in-transit

Option to enforce
encryption in-transit

NEW!

Integrates with
your organization's
AD and supports
Windows ACLs

Network traffic access
control using Amazon
VPC security groups

Admin API
access control
using AWS IAM

Monitor and log
API calls using
AWS CloudTrail

PCI-DSS + ISO-
+ SOC + GDPR compliant
and HIPAA eligible

Use Amazon FSx with your organization's AD

Directly integrate your Amazon FSx file systems with your organization's Active Directory (on-premises or in-cloud)

- **Authentication**: Your users continue to access file shares by authenticating with their existing AD user credentials
- **Authorization**: You can migrate and use your existing file and folder ACLs, and your share-level access controls as is, without any modifications needed

Supports two AD integration options:

- AWS Managed Microsoft AD
- Self-managed Microsoft AD (on-premises or in-cloud)

Use Amazon FSx with your self-managed AD

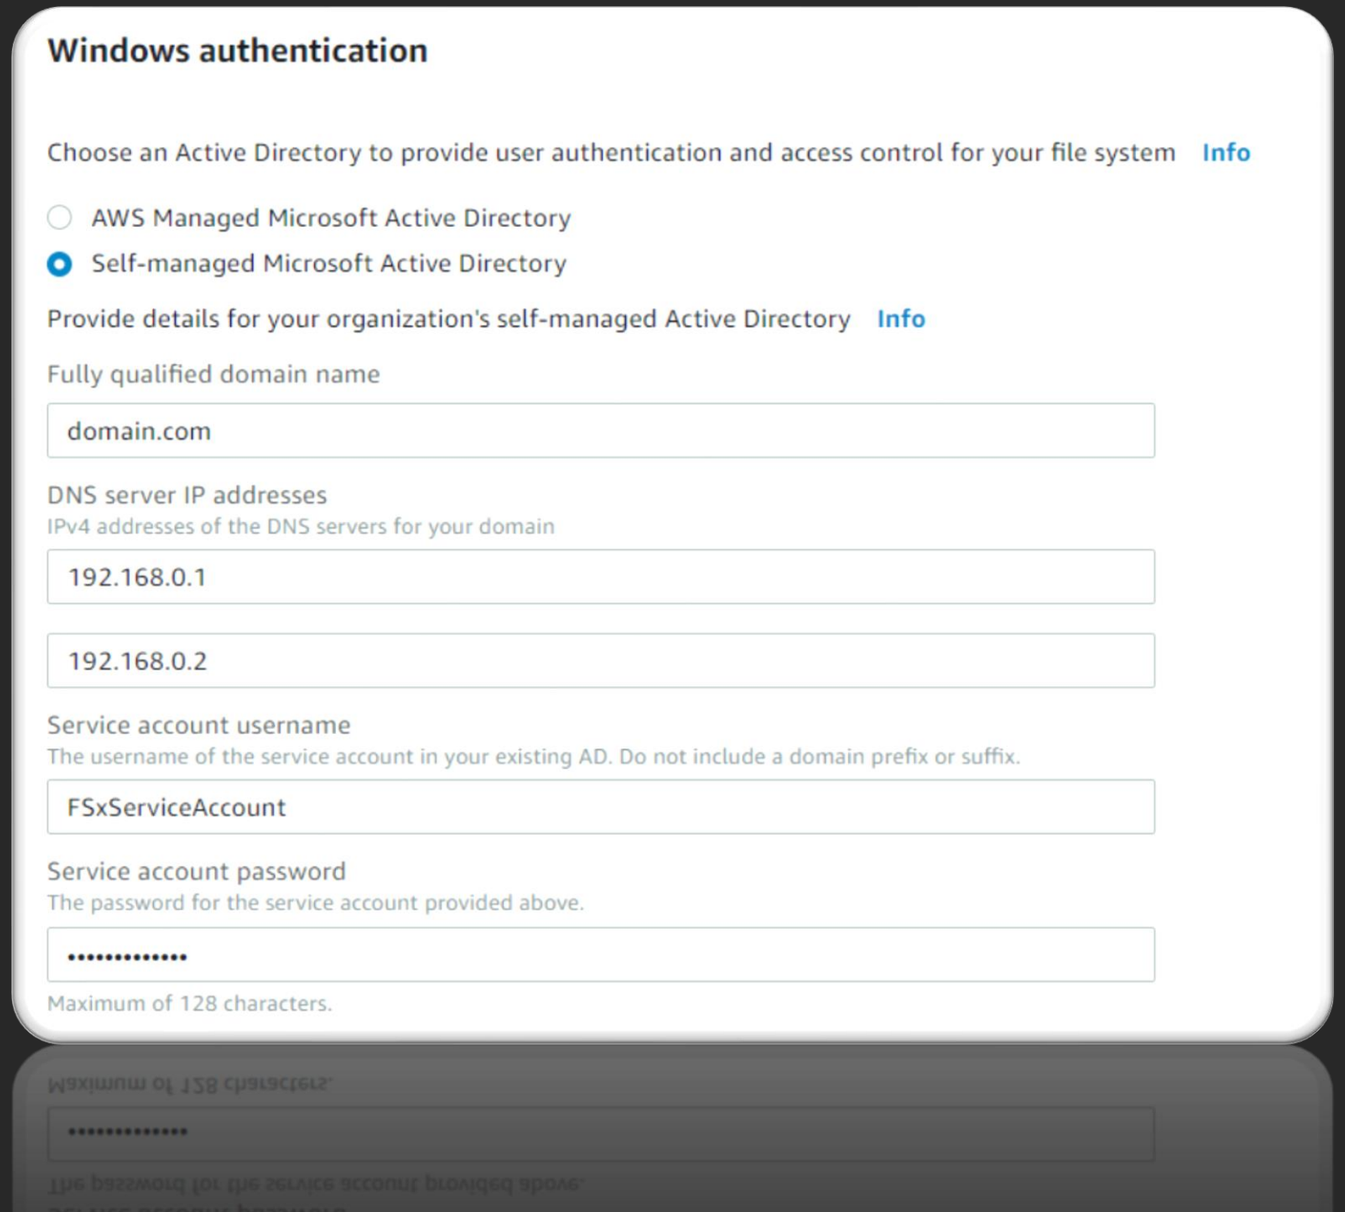

Backups

Highly durable	Highly durable (11 nines) – stored in Amazon S3
File system consistent	Capture and restore a point-in-time view of file system Ensures file system-consistency using VSS
Incremental	Only changes after your most recent backup use backup storage
Fully managed	Automatic daily backups, with retention policy Admin-initiated backups via API/Console

File-level restore by end-users (Shadow copies)

Undo changes to individual files
No need to restore the entire file system

Self-service
No more filing a ticket to admin

Compare previous file versions

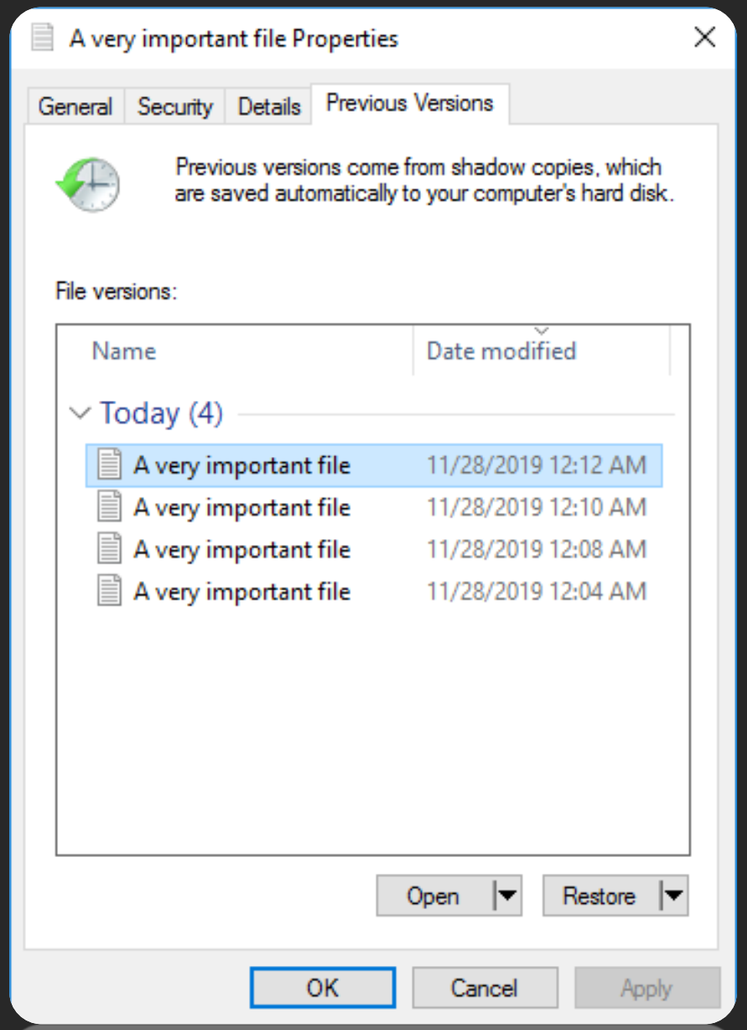

Administration

Administering AWS resource properties

- CloudWatch metrics
- CloudWatch alarms
- Maintenance window
- Automatic backups
- Tags
- Active Directory configuration

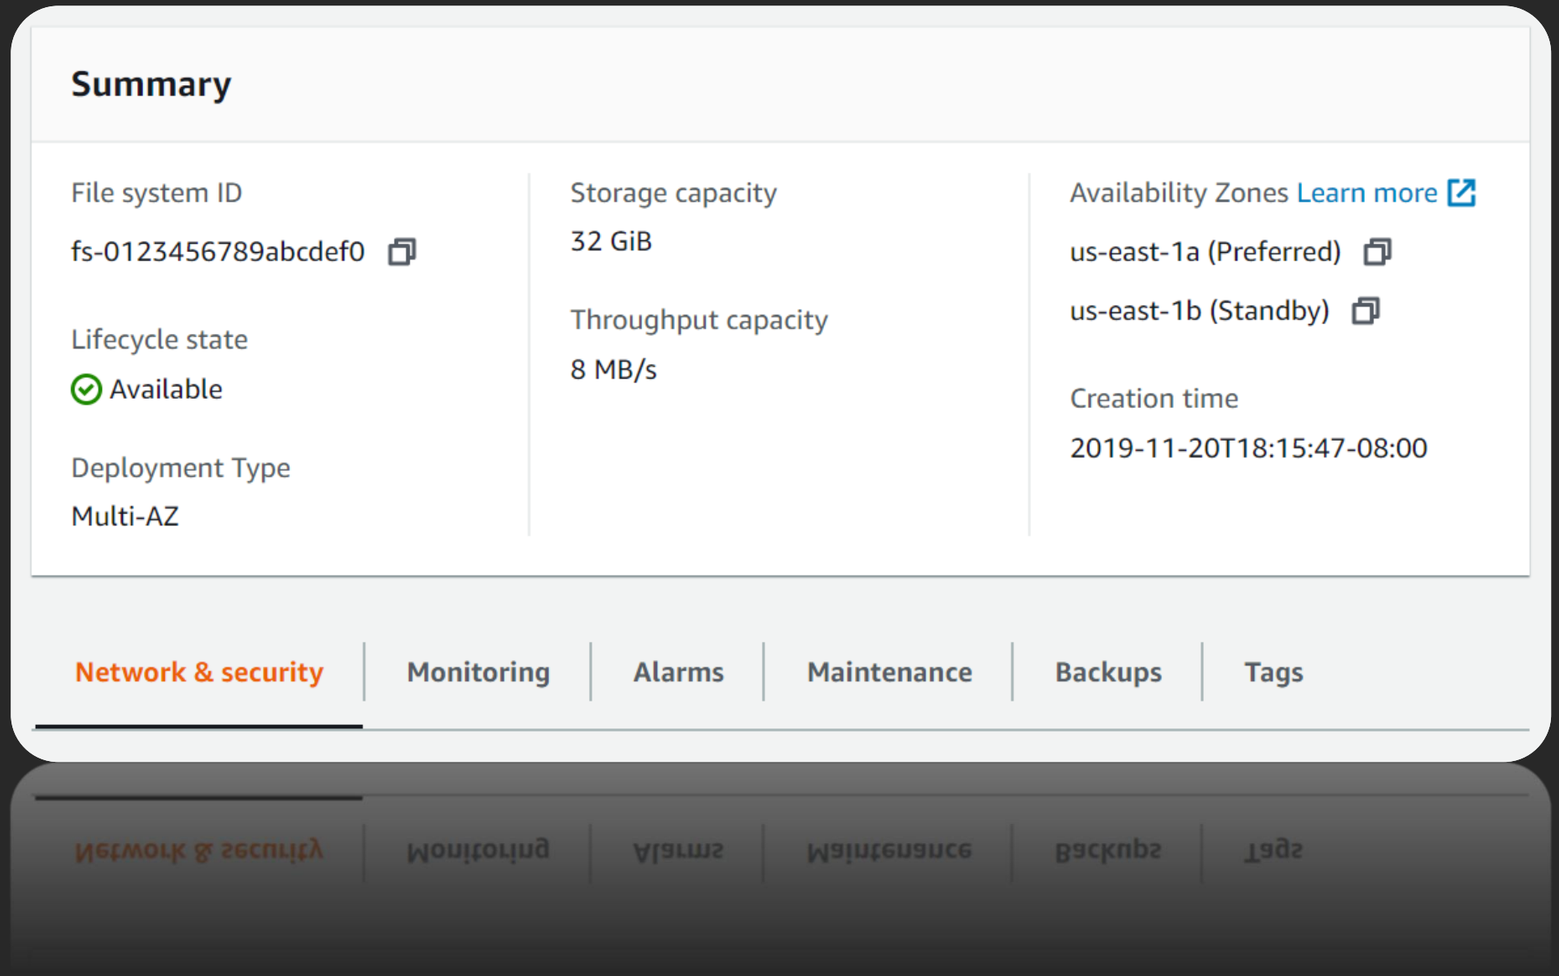

Administering file system features

What you can configure using remote PowerShell:

- SMB file shares
- Shadow copies
- Data deduplication
- User quotas
- Open sessions / files
- Enforcement of encryption in transit

Administering file system features

```
PS C:\Users\Admin> Enter-PsSession -ComputerName $PS_ENDPOINT -ConfigurationName FSxRemoteAdmin
[amznfsx45oujbsl.dom.com]: PS>Get-FSxSmbShare

Name  ScopeName        Path     Description
----  ---------        ----     -----------
D$    AMZNFSXTNS9I2P4 D:\      Default Share
IPC$  *                        Remote IPC
share AMZNFSXTNS9I2P4 D:\share


[amznfsx45oujbsl.dom.com]: PS>Exit-PsSession
PS C:\Users\Admin>
```

```
PS C:\Users\Admin> Invoke-Command -ComputerName $PS_ENDPOINT -ConfigurationName FSxRemoteAdmin -ScriptBlock {
>>     Get-FSxSmbShare
>> } | Select Name, Path, ScopeName, Description

Name  Path     ScopeName        Description
----  ----     ---------        -----------
D$    D:\      AMZNFSXTNS9I2P4 Default Share
IPC$           *                Remote IPC
share D:\share AMZNFSXTNS9I2P4
```

Migration

Migration

Simple and seamless migration

Migrate files using Windows-native tools like RoboCopy

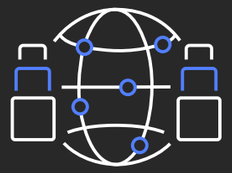

Preserve existing security (ACLs) and DFS Namespaces, and continue to use existing AD user identities

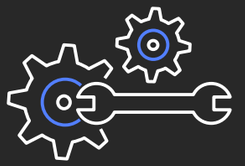

SMB share migration tooling

Migration transparent to your applications and users
i.e., no need to modify your existing applications or user workflows

Amazon FSx in Action

Summary and Q&A

Related sessions

STG 201
AWS leadership session: Storage state of the union

STG 202
What's new in AWS file storage

STG 238
File storage for business-critical applications

STG 211
How to use AWS storage for on-premises file applications

STG 362
Migrating Windows file servers to Amazon FSx

STG 326
Deploy fully managed Amazon FSx for Windows File Server instances in minutes

STG 324
Map and drive performance with Amazon FSx for Windows File Server

Thank you!

Edward Naim
edwanaim@amazon.com

Prashanth Bungale
bungale@amazon.com

Darryl Osborne
darrylo@amazon.com

AWS re:Invent

aws

Please complete the session survey in the mobile app.