



aws **INNOVATE**
FOR EVERY APPLICATION EDITION

25 August, 2022

How to create 100 database clones in minutes with AWS

Naim Mucaj
Senior Solutions Architect
Amazon Web Services

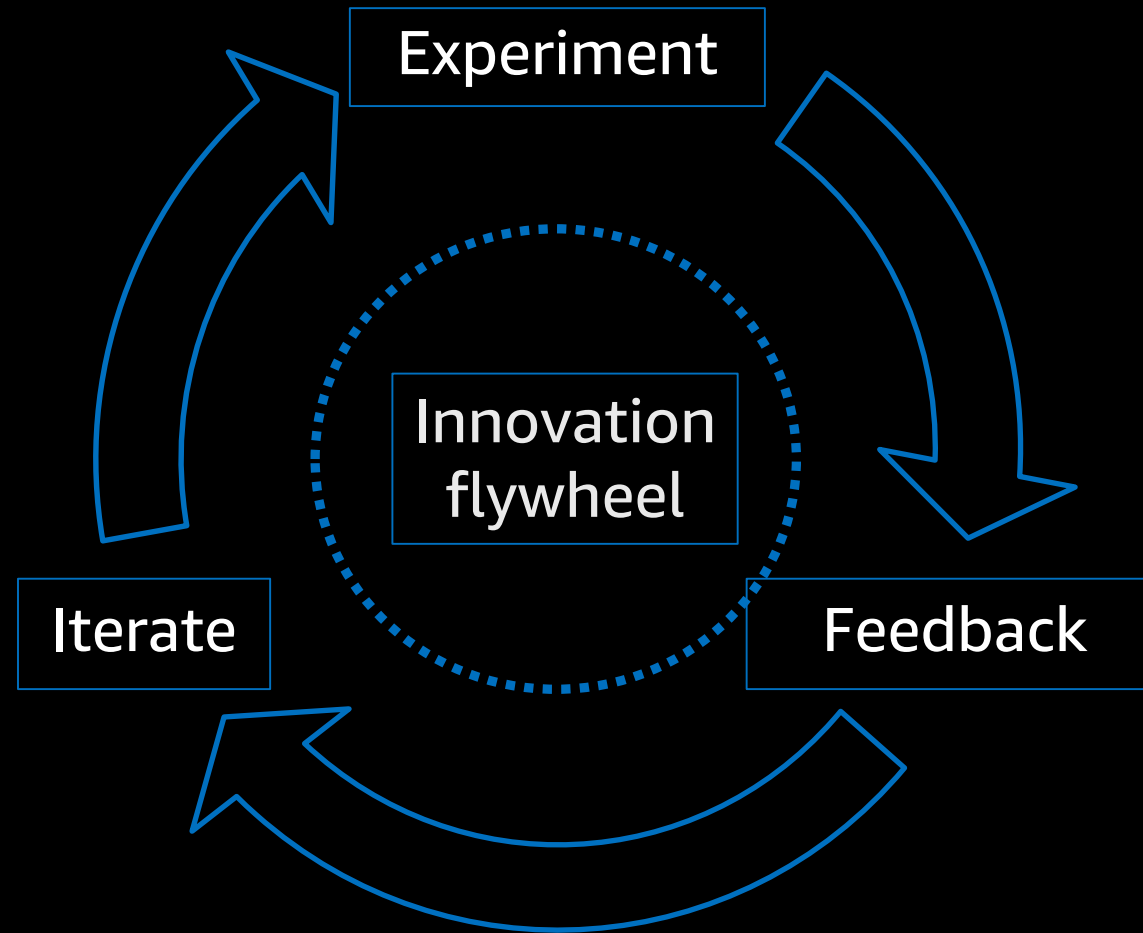


Today's agenda

- Data growth
- Common application development cycles
- Access to data for development and testing
- How Amazon FSx file systems can help accelerate development cycles
- Demo - Creating 100 database clones
- Resources – getting started

How much **data** is out there?

Experiments power the engine of rapid innovation



Common development cycles



```
graph LR; Development --> Test; Test --> Pre-prod; Pre-prod --> Production; Production --> Monitor
```


Development

Test

Pre-prod

Production


Monitor



I want access to more data sets more often!

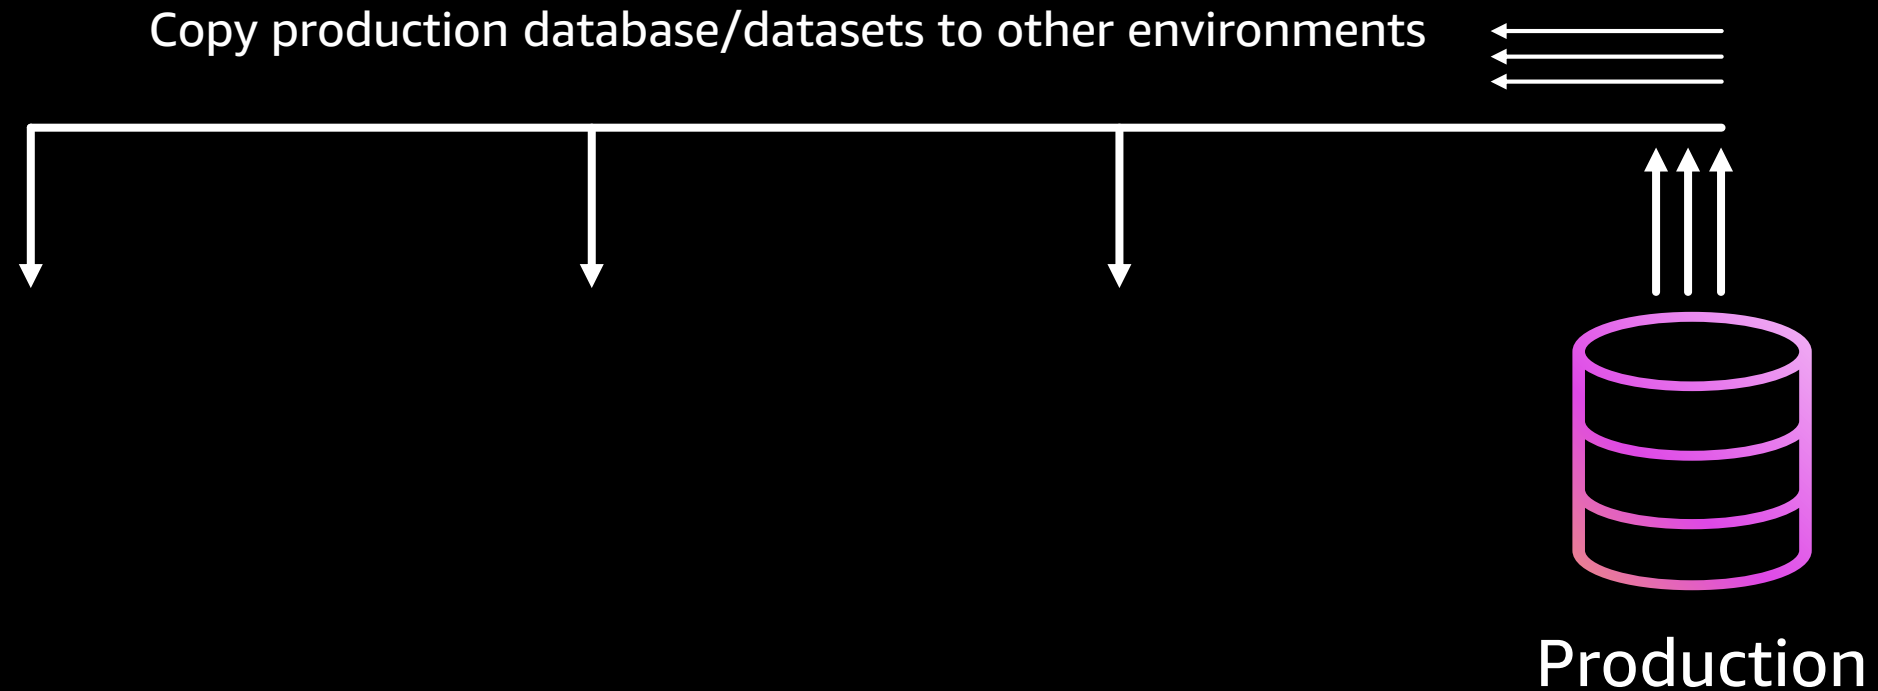
Developers, Database Administrators, Data Analysts, Application Owners

Refreshing development environments

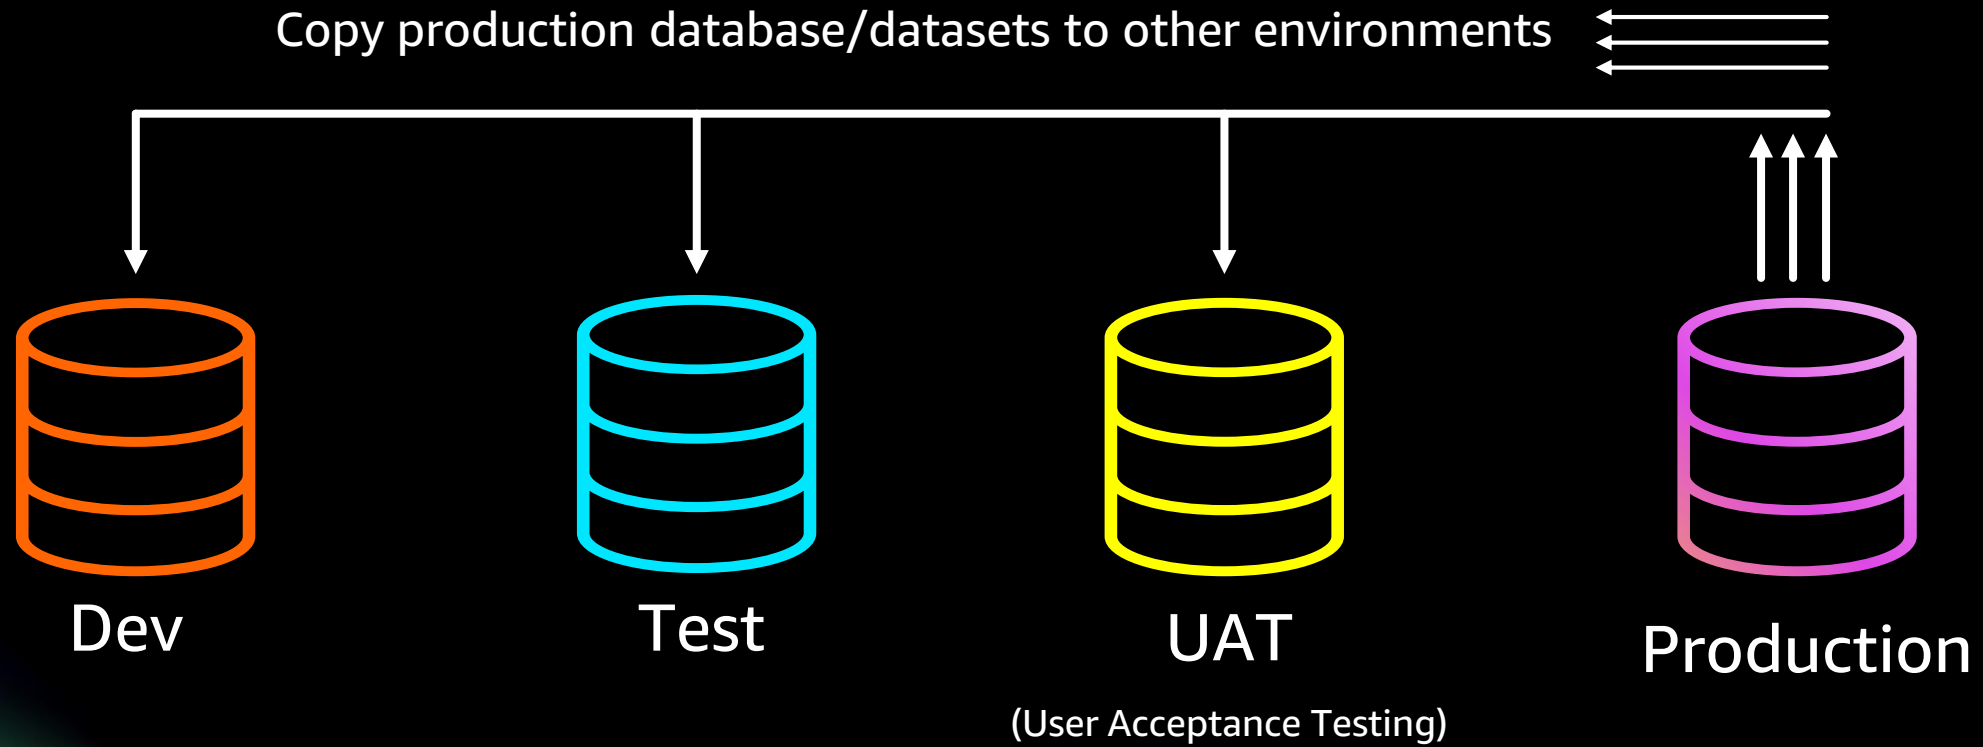


Production

Refreshing development environments



Refreshing development environments



Refreshing development environments



Dev



Test



UAT

(User Acceptance Testing)



Production

The process burden

App/DB/Development/Test teams:


*Copy, sync, dump, transfer, custom scripts, waiting on others,
want to maintain more versions*

Infrastructure teams:


*Network saturation, performance issues, backup window conflicts,
storage capacity, server/compute limitations*

The process can take hours, days, weekends to refresh


Legacy data infrastructure requires undifferentiated heavy lifting



Hardware and software installation




Performance and availability




Capacity planning and cluster scaling




Security and compliance management



Expensive and proprietary, with high amounts of lock-in



Data when and where you need it



Amazon FSx file systems

Many common use cases for file storage



Database
applications



Software build
environments

Amazon FSx for NetApp ONTAP

THE ONLY COMPLETE, FULLY MANAGED NETAPP FILE SYSTEM IN THE CLOUD

Sept.
2021

FSx
N

NetApp
ONTAP



Leverage ONTAP's data management capabilities and multi-protocol support



Optimize cost with automatic tiering and built-in data compression and deduplication



Leverage ONTAP snapshots and cloning to run efficient data environments



**When you
need it**

Amazon FSx for OpenZFS

HIGH-PERFORMANCE STORAGE POWERED BY AWS GRAVITON AND OPENZFS

re:Invent
2021

Powered by
Graviton

FSx
ZFS

OpenZFS



Migrate ZFS and other Linux-based file servers



Up to **1 million IOPS** at latencies of just a few hundred microseconds



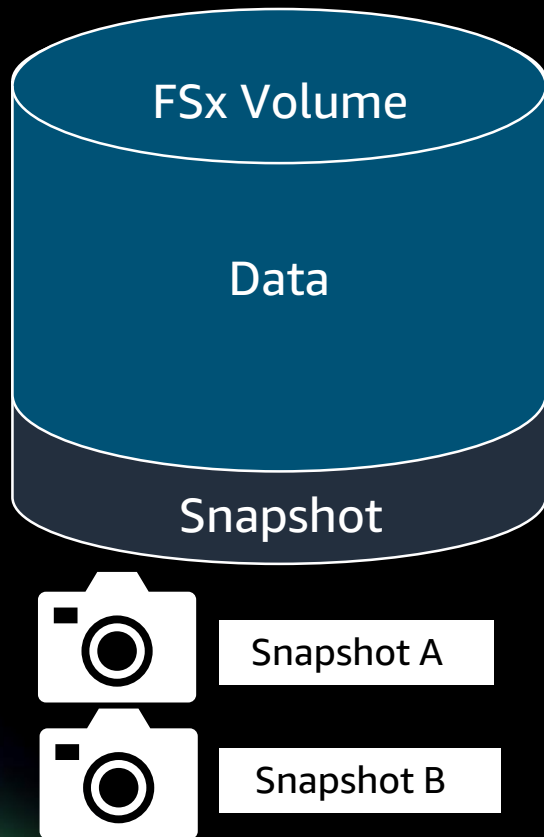
Leverage ZFS snapshots and cloning to run data intensive workloads



When you
need it

Data protection: Volume snapshots

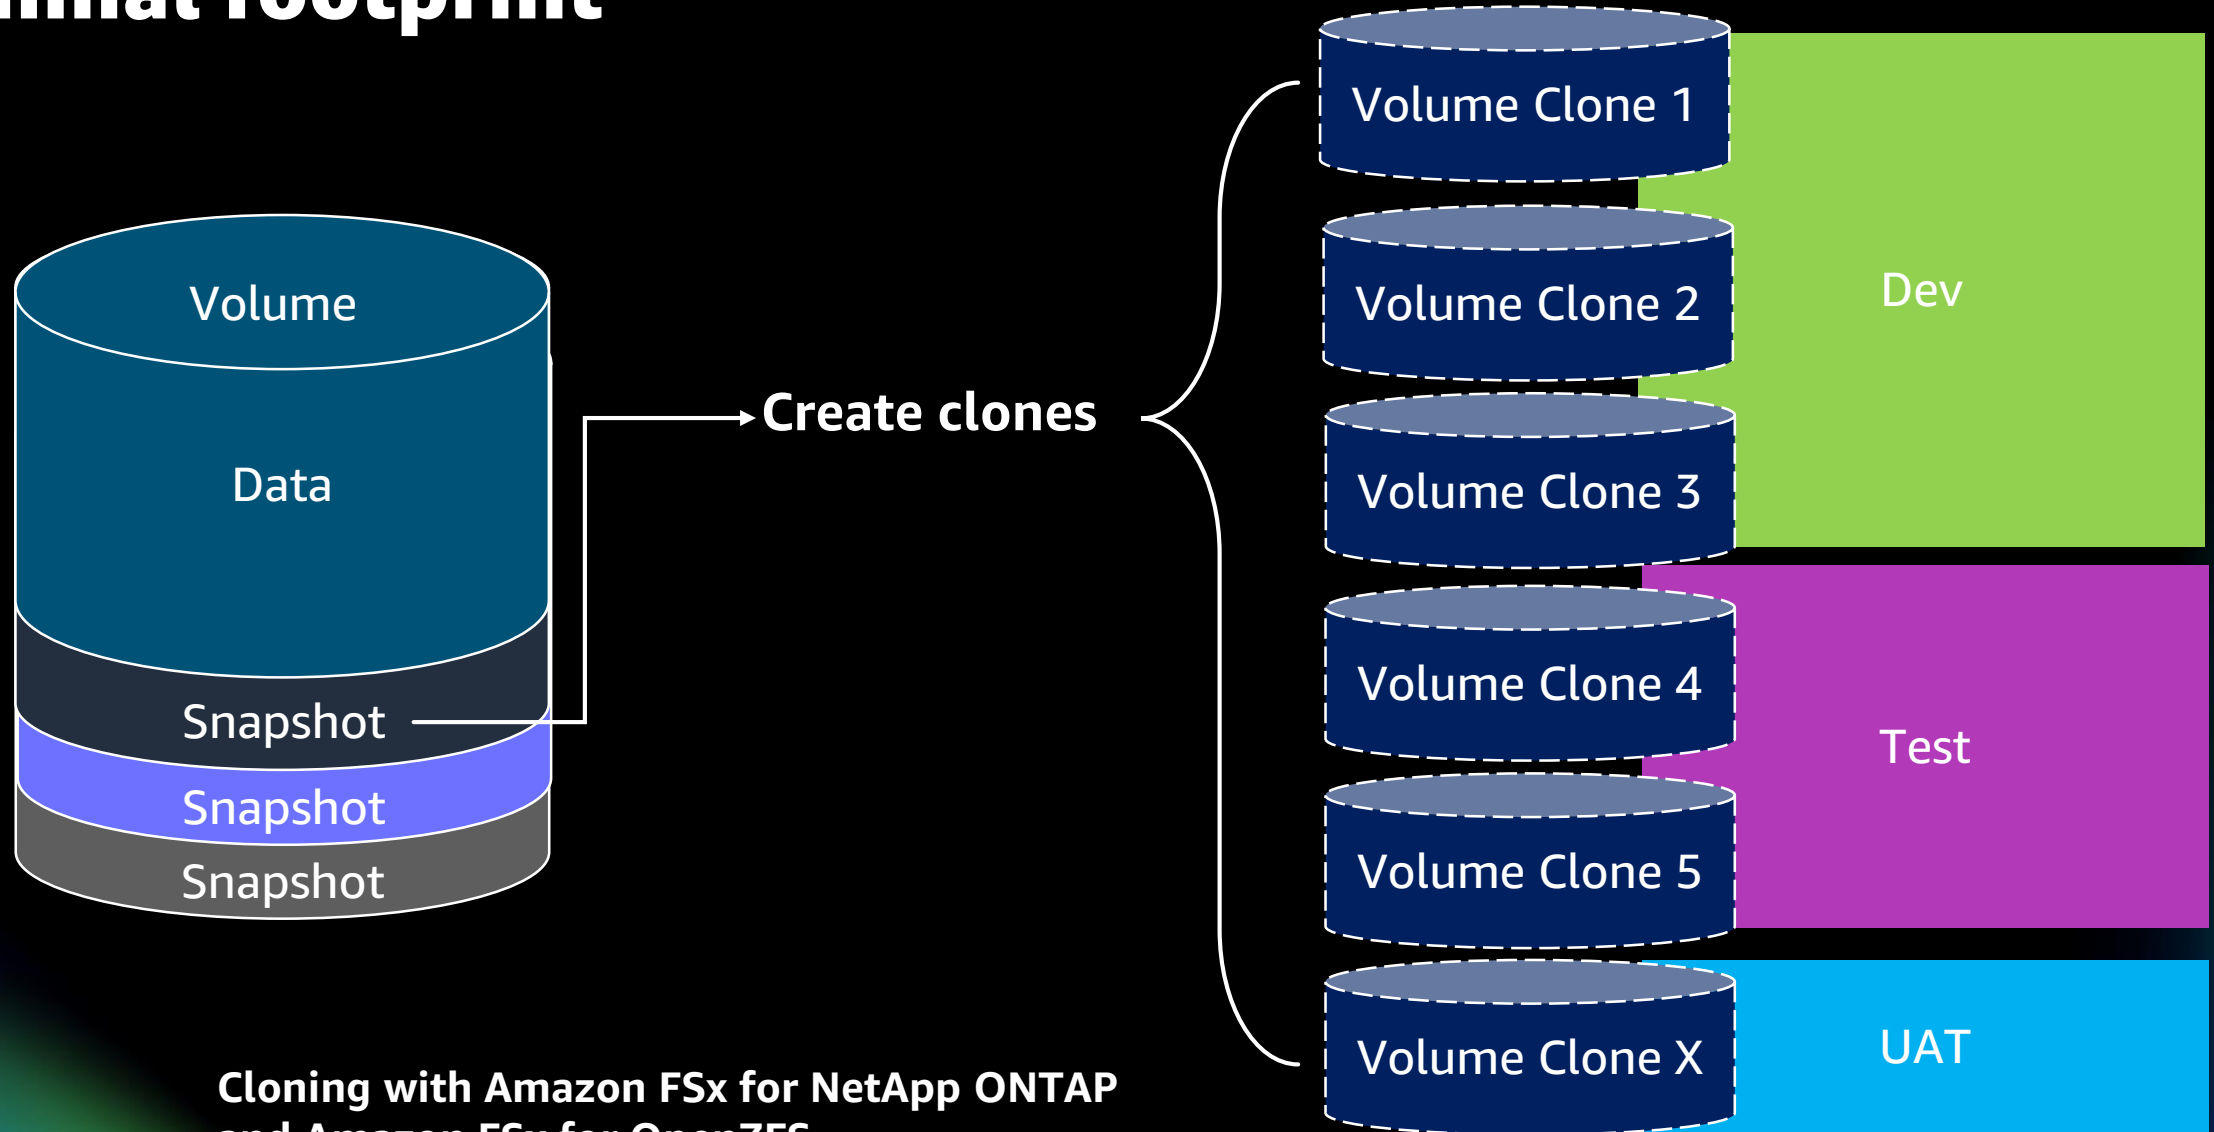
Amazon FSx file system



Snapshots

- Fast to create
- Shorter-term recovery
- Internal, stored on file system volume
- Built-in, automatic or user scheduled snapshots
- Individual file or entire volume restore
- Crash-consistent & incremental

Volume cloning with Amazon FSx – Efficient datasets, minimal footprint



Cloning with Amazon FSx for NetApp ONTAP
and Amazon FSx for OpenZFS

Volume cloning benefits – Access more data more often

- Create multiple space efficient, minimal footprint copies of your data sets
- Writable, point-in-time copies of a volume
- Share data blocks with parent volume
- Consume no storage, until changes are made
- Fast – created almost instantly
- You can then split a clone into it's own independent volume

Multiple datasets (without cloning)



Multiple datasets (with clones)



Pearson reduced RTO/RPO from hours to minutes

“FSx for NetApp ONTAP improved the agility of our mission-critical ERP application by enabling us to **refresh our data 10x faster**. We also **increased financial and HR job performance by 3x** and **reduced our application recovery time and recovery point (RTO/RPO) objectives from hours to minutes**. As a fully managed service, FSx for NetApp ONTAP removed the operational overhead of our physical IT infrastructure freeing us to focus more on our application data and drive more value for our customers.”

Norm Nadeau

VP of Site Reliability Engineering, Pearson

[Learn more](#)



Pearson

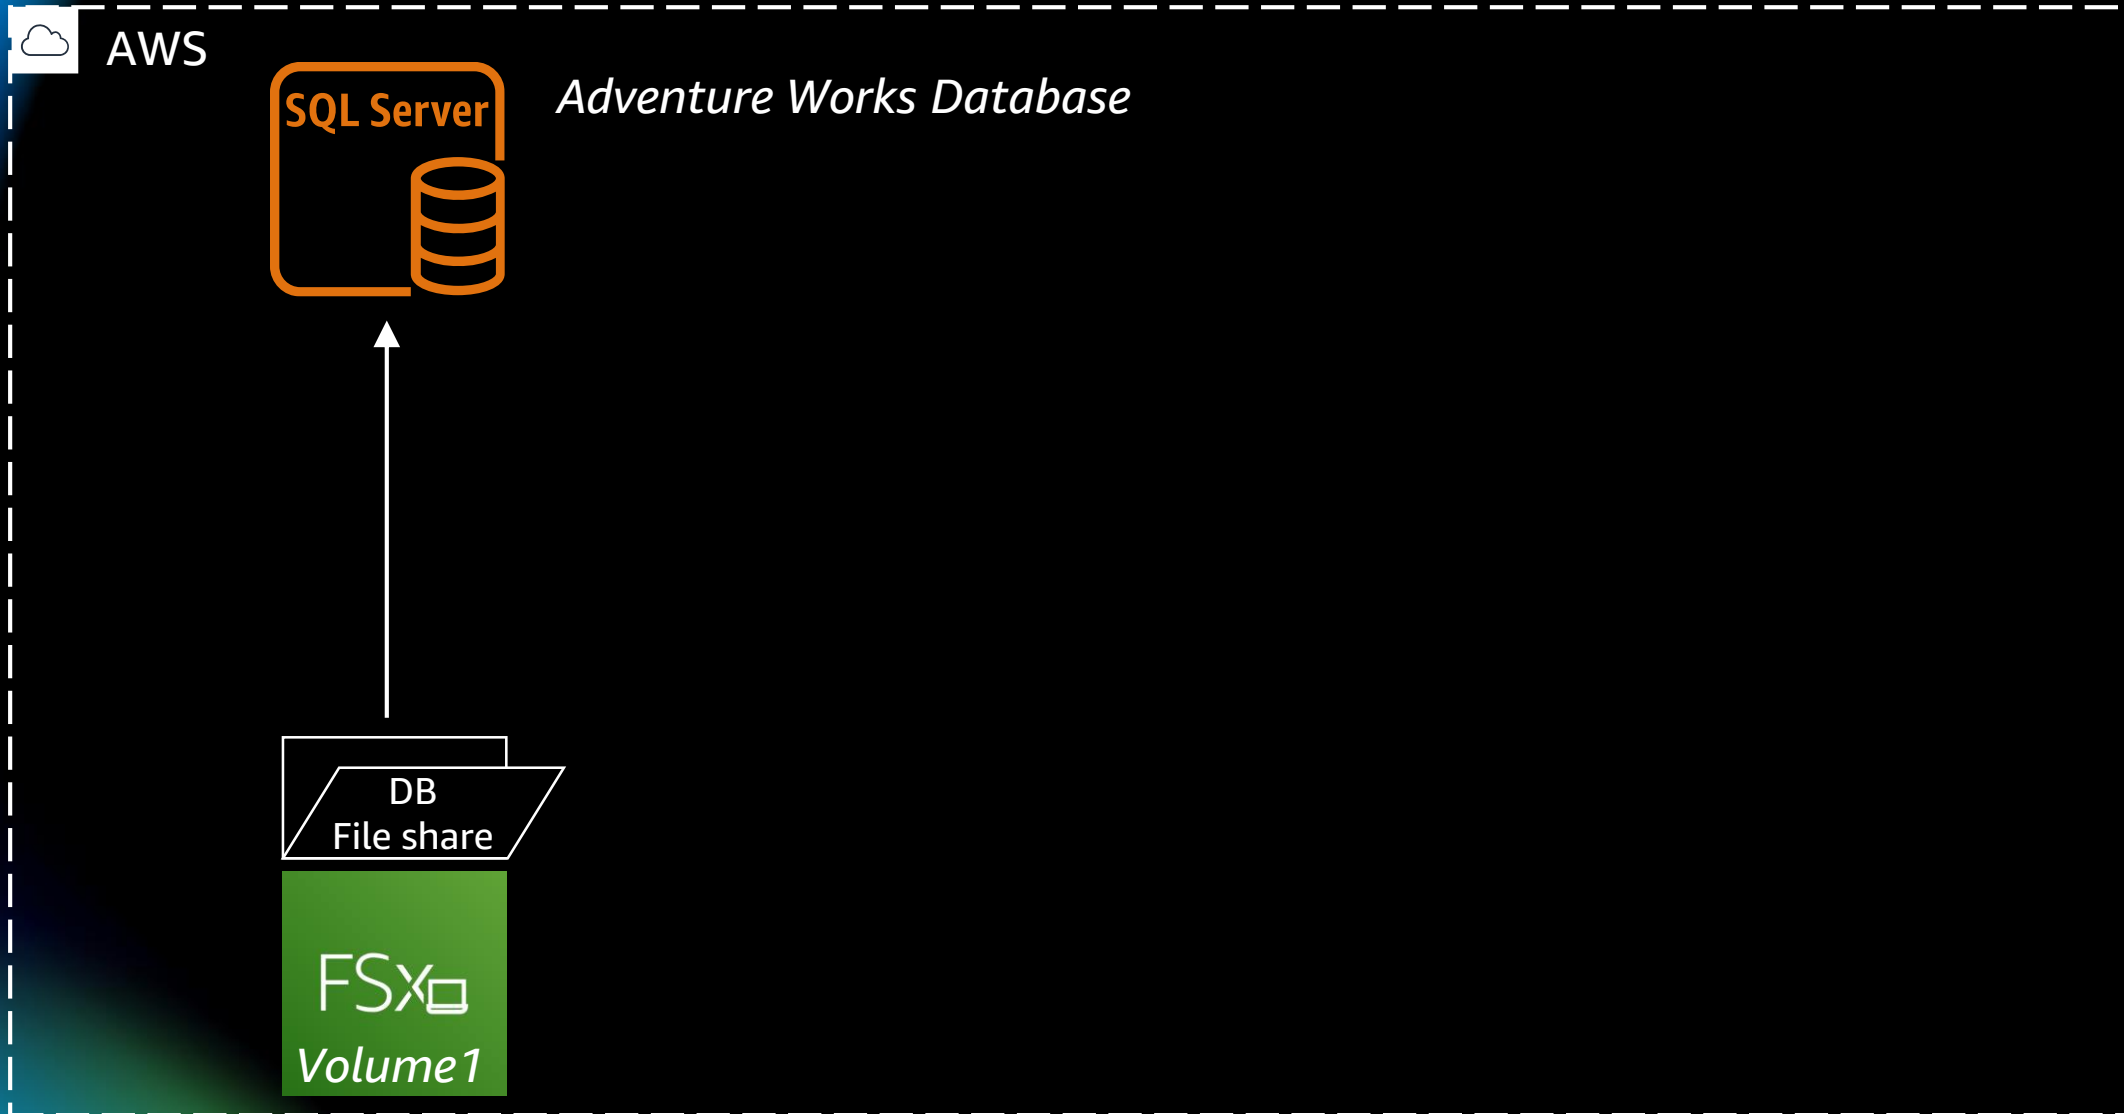
DEMO



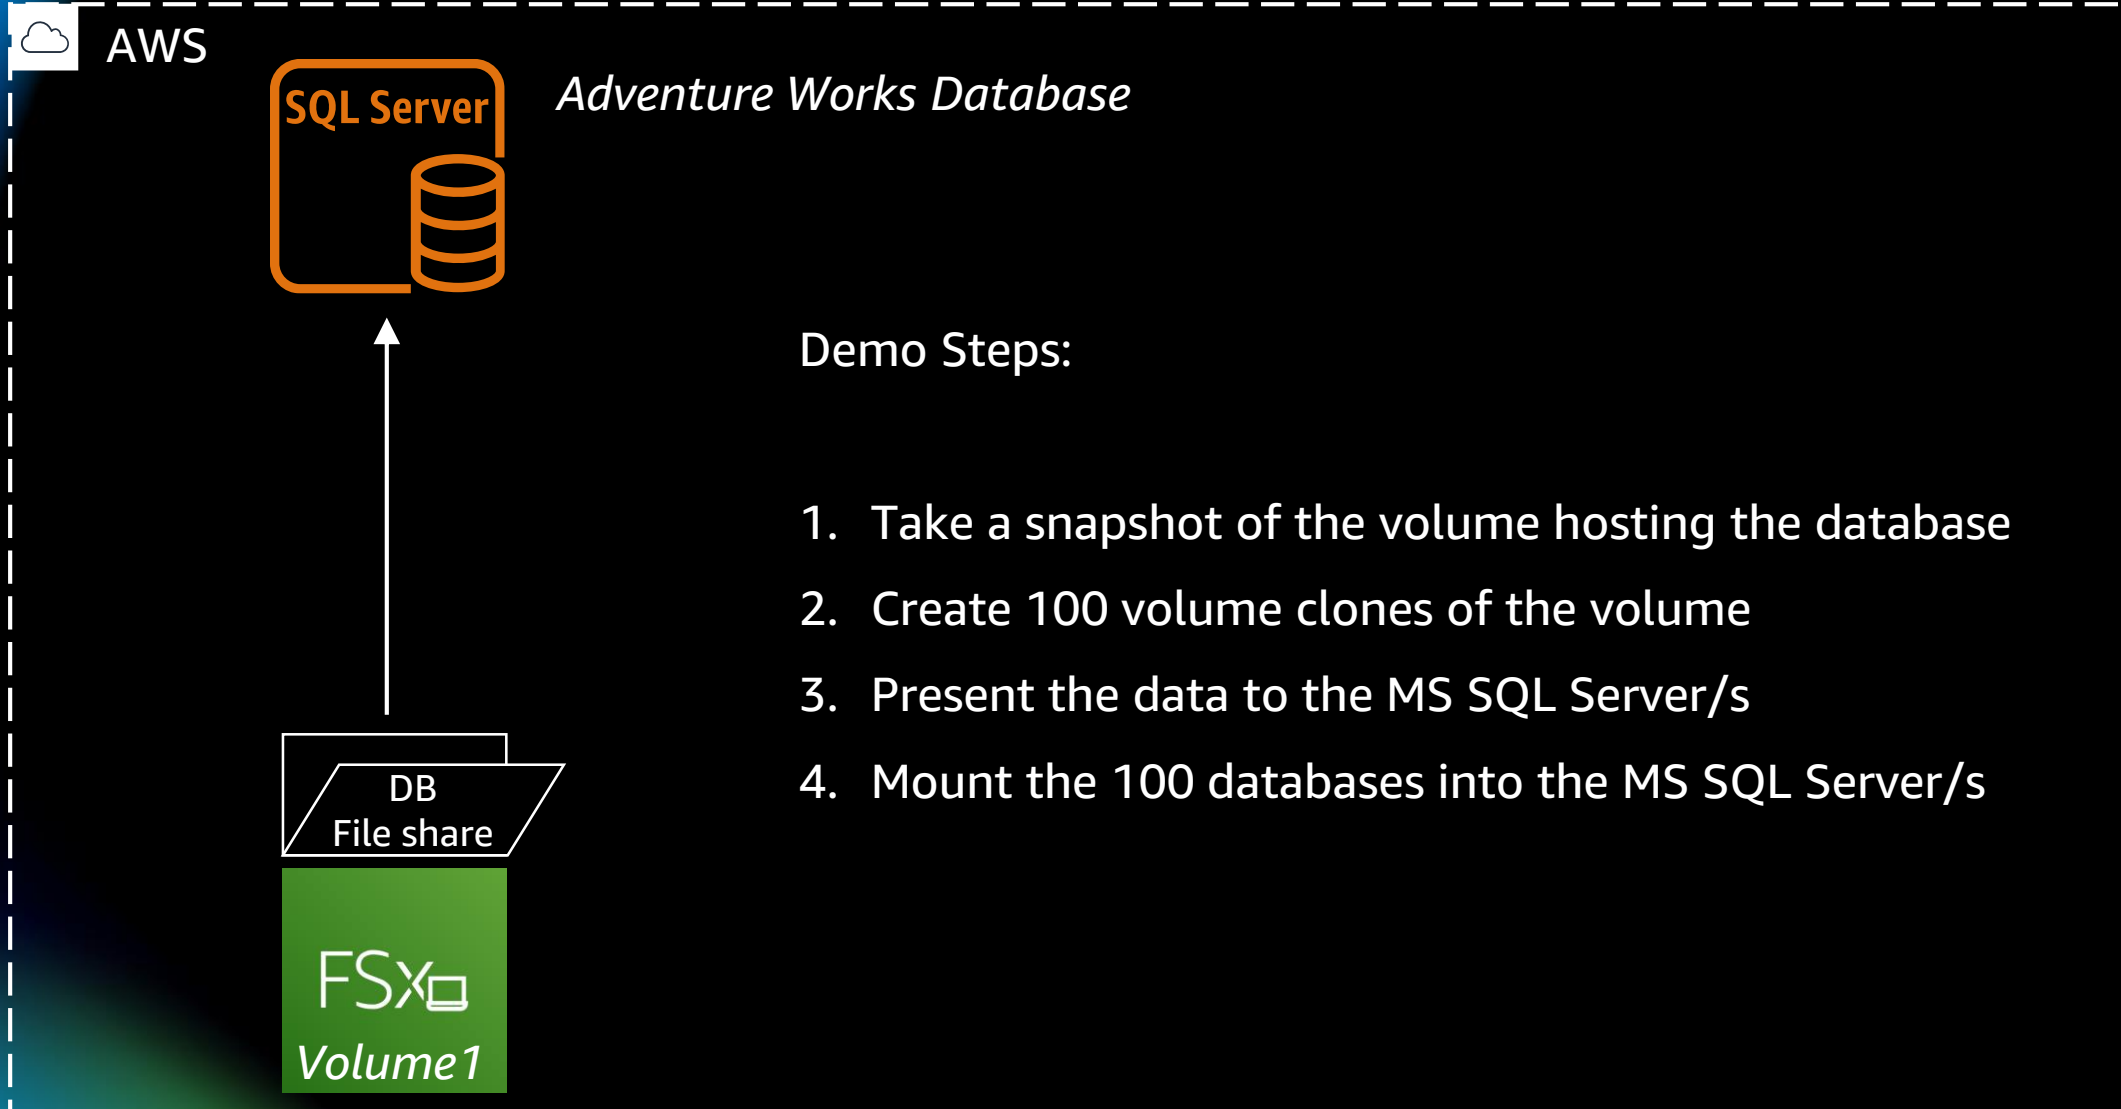
Demo scenarios – Sophia, Joe and MS SQL

- **Sophia** manages IT compute and storage infrastructure with multiple APP/DB environments.
- **Joe** is a Database Administrator and has a need to **create multiple environments** from a Production MS SQL Database data set.
- **MS SQL** database is currently located on an Amazon FSx for NetApp ONTAP volume. Multiple Amazon EC2 Instances can be spun up to **burst develop / test** the cloned databases leveraging **volume cloning**.

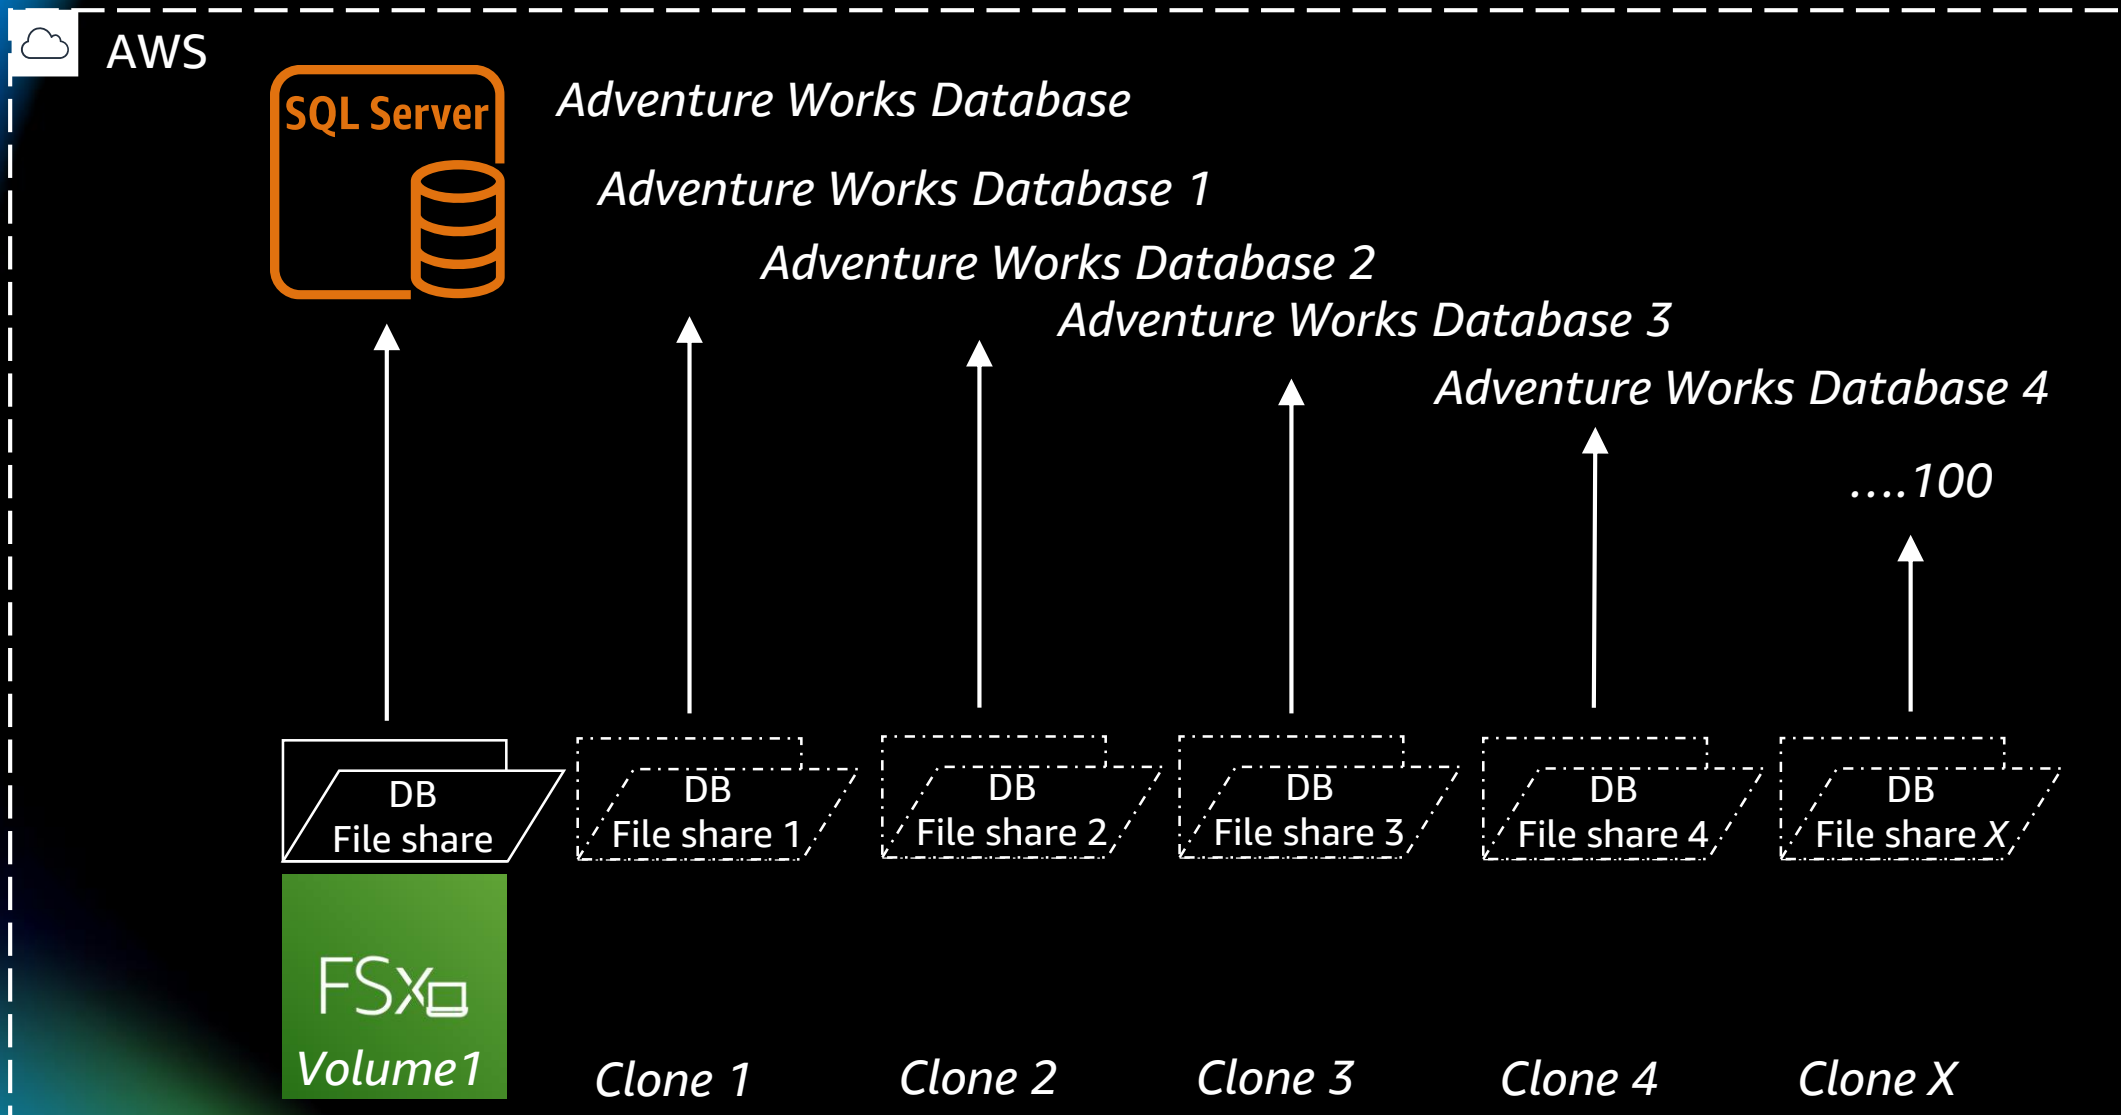
Demo environment



Demo environment



Demo recap



Summary

- Challenges of maintaining and refreshing datasets environments.
- You can reduce operational complexity, improve agility, and reduce costs by using the in-built features to clone your datasets with Amazon FSx for NetApp ONTAP & Amazon FSx for OpenZFS file systems.
- Accelerate **test & development refresh cycles**.
- Use snapshots and clones to address **point-in-time checkpoints (Snapshots)**, and for presentation of **efficient read/write datasets for dev/test/UAT (clones)**.

Other resources

Learn more about Amazon FSx file systems:

<https://aws.amazon.com/fsx/>

<https://aws.amazon.com/fsx/netapp-ontap/>

<https://aws.amazon.com/fsx/openzfs/>

Checkout "Pearson" Case Study Video:

<https://www.youtube.com/watch?v=4FZpoxN3sJQ>

Learn more about Amazon Storage Services:

<https://aws.amazon.com/products/storage/>

Visit the AWS resource hub

Start building upon a scalable, reliable, and globally available infrastructure so that you can focus on innovation and bringing new applications to market. Dive deeper with these resources today.

- Accelerate innovation with AWS
- Get more performance for your applications at lower costs with AWS
- Global-scale solutions
- How startups succeed with AWS



<https://tinyurl.com/for-every-app-hub-aws>

Visit resource hub

AWS Training and Certification



Self – Paced Digital Training on AWS

Explore learning plans and 500+ digital courses from our new learning center, AWS Skill Builder, to help you achieve your goals on your schedule.

bit.ly/3lzVj0g



AWS Certification

Validate technical skills and cloud expertise to grow your career and business.

go.aws/3PwN3ff

Thank you for attending AWS Innovate – For Every Application Edition

We hope you found it interesting! A kind reminder to **complete the survey**.
Let us know what you thought of today's event and how we can improve the event
experience for you in the future.



aws-apj-marketing@amazon.com



twitter.com/AWSCloud



facebook.com/AmazonWebServices



youtube.com/user/AmazonWebServices



slideshare.net/AmazonWebServices



twitch.tv/aws

Thank you!