Storage Considerations and Best Practices for SAS® Viya® on AWS

Darryl Osborne

Dilip Rajan

Agenda

AWS and SAS – Technical engagement

SAS 9 on AWS Overview

SAS Viya on AWS

Compute Updates

Storage Overview

SAS Viya Demo with storage recommendations

Resources

SAS - AWS

Challenges – Why SAS on AWS

Traditional on-premises data-analytics infrastructure does not scale well enough and is too expensive.

To run analytics at scale, we need . . .

- Agile and flexible compute options
- Scalable and performant storage
- Easy lift-and-shift for existing workflows
- Modernize monolithic applications using microservices
- The ability to burst workloads to the cloud
- The option to consolidate data from disparate data sources

SAS Grid on AWS – Recap 1

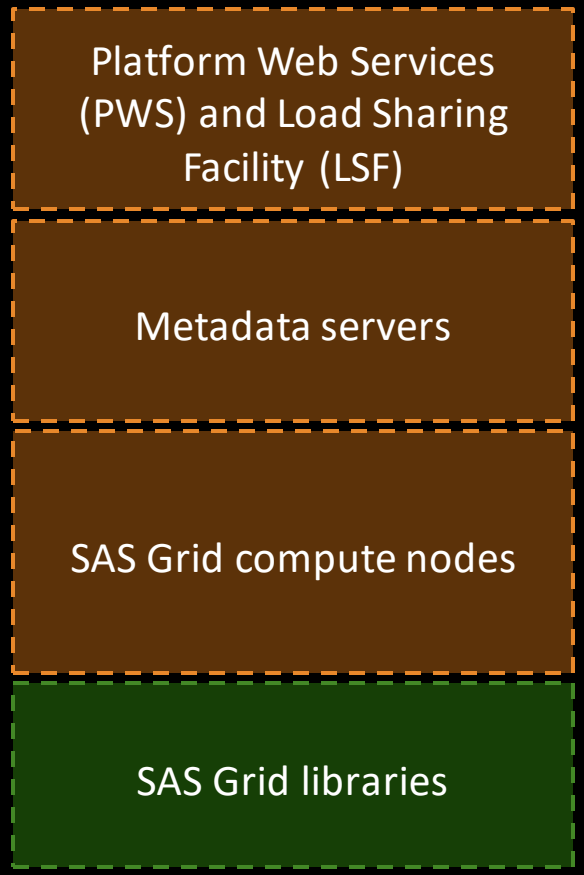

Dispatch jobs submitted and return the status of each job

Contain the metadata repository that stores the metadata definitions of all SAS Grid manager products

Run a portion of the work allocated to the grid

Store your enterprise data

SAS Grid on AWS – Recap 2

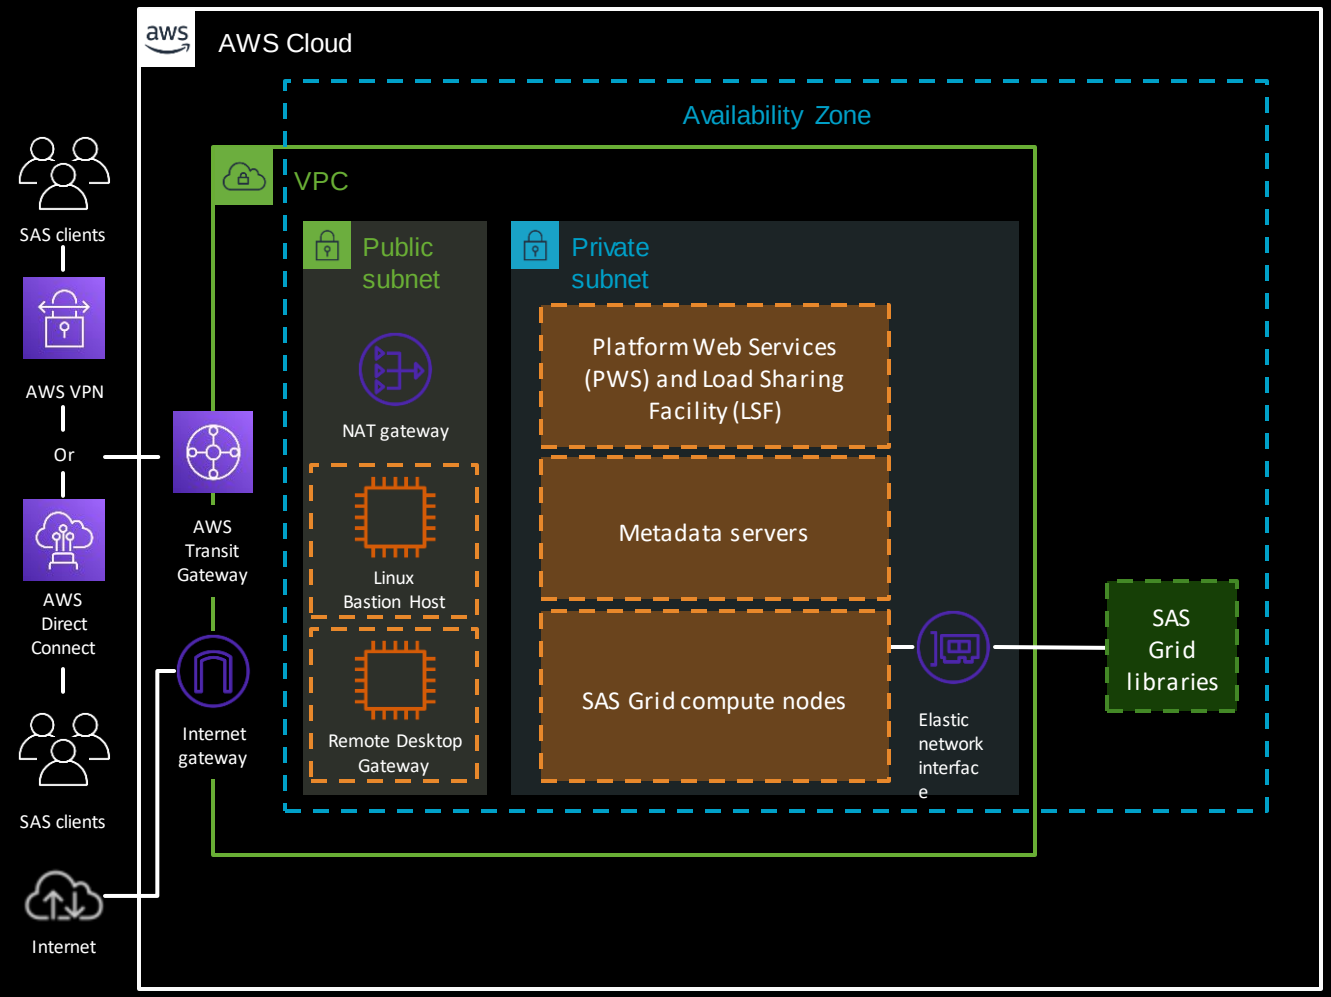

Why Lustre

The world's most popular high-performance file system

Why Lustre on AWS

Increase agility and reduce costs

Why Amazon FSx for Lustre

Reduce storage administration and accelerate innovation

SAS Grid on AWS

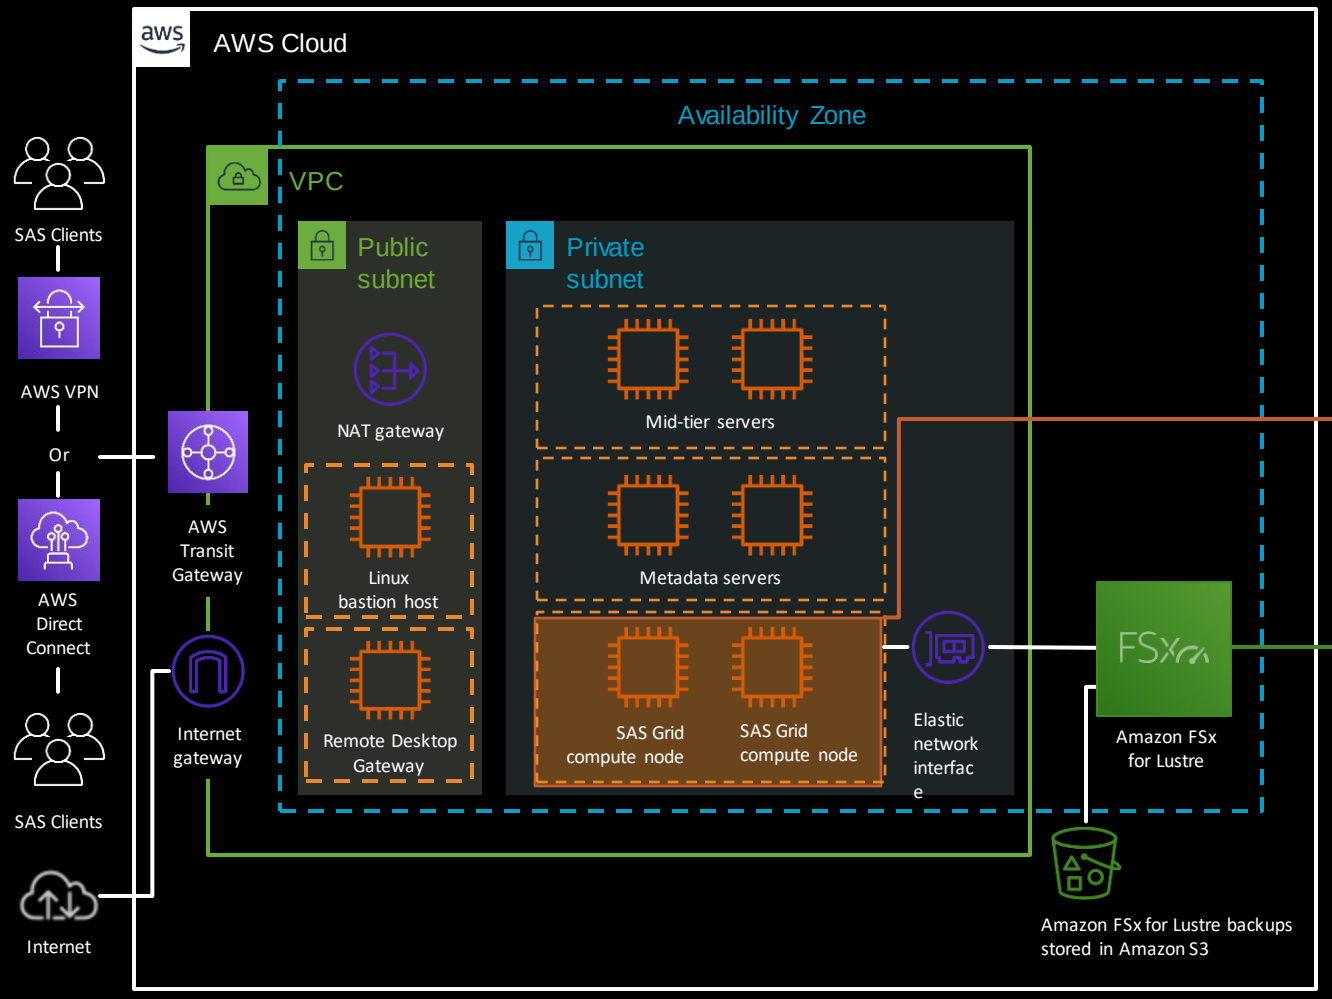

SAS Grid costs and performance are driven by:

SAS Grid compute nodes

Storage
- /SASDATA
- /SASWORK
- /UTILLOC

Instance types vs. compute node requirements

EC2 instance type	Minimum storage (8GB)	Minimum throughput (100 MB/s)
i3en	16 GB	>100 MB/s
i3	15 GB	<100 MB/s
r5n / r5dn	16 GB	>100 MB/s
r5a / r5ad	16 GB	<100 MB/s
r6g / r6gd	16 GB	<100 MB/s
r5	16 GB	<100 MB/s
m5n / m5dn	8 GB	>100 MB/s
m5a / m5ad	8 GB	<100 MB/s
m6g / m6gd	8 GB	<100 MB/s
m5	8 GB	<100 MB/s
c5n	5.25 GB	>100 MB/s
c5a / c5ad	4 GB	<100 MB/s
c6g / c6gd	4 GB	<100 MB/s
c5	4 GB	<100 MB/s

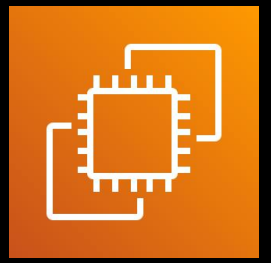

Amazon EC2

Compute node requirements

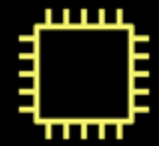

8 GB
minimum memory
per physical core

125 MB/s
minimum throughput
per physical core

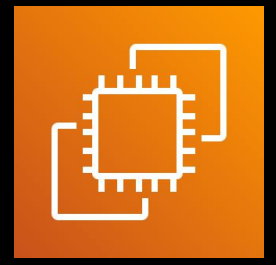

Amazon EC2

Compute node recommendations for SAS Grid

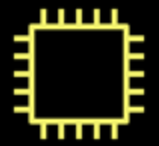

	Physical cores	-Per physical cores specification			Instance store volumes	
	(# of physical cores)	Memory (GiB)	Variable network performance (MB/s)	Consistent network performance (MB/s)	NVMe storage capacity (GiB)	Throughput performance (MB/s)
m5n	1, 2, 4, 8, 16, 24, 32, 48	8	396.18 - 1774.61	258.89	Not available	Not available
m5dn	1, 2, 4, 8, 16, 24, 32, 48	8	396.13 - 1915.05	259.17	75 – 3,600	62.48
r5n	1, 2, 4, 8, 16, 24, 32, 48	16	396.12 - 1789.06	259.05	Not available	Not available
r5dn	1, 2, 4, 8, 16, 24, 32, 48	16	396.14 - 1844.39	258.84	75 – 3,600	62.22
i3en	1, 2, 4, 6, 12, 24, 48	16	528.18 - 1745.51	259.81	1250 – 60,000	154.26

T-Mobile realizes $1.5M in annual savings and doubles the speed of SAS Grid workloads using Amazon FSx for Lustre

The challenge

T-Mobile managed their own storage for SAS Grid, which proved to be unscalable and cost prohibitive

The solution

T-Mobile deployed Amazon FSx for Lustre, a fully-managed high-performance file system, to migrate and scale their SAS Grid infrastructure

The benefits

Reduced TCO by 83% and reduced storage costs by 67%, resulting in $1.5M in annual cost savings

Cut the run time of SAS Grid analytics workloads by half

> "Amazon FSx for Lustre helped us double the speed of our SAS Grid workloads, reduce our Total Cost of Ownership by 83% and completely eliminate our operational burden"

Dinesh Korde, Sr. Manager Software Development, T-Mobile

Compute Overview

Running your data analytics at scale

Microservices

Containers and serverless compute

Serverless

AWS Lambda

Containers

AWS Fargate

Red Hat OpenShift

Amazon EKS

Amazon ECS

Compute

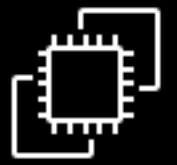

Amazon EC2

Performant, scalable, persistent
shared storage

Comparison of operational responsibility

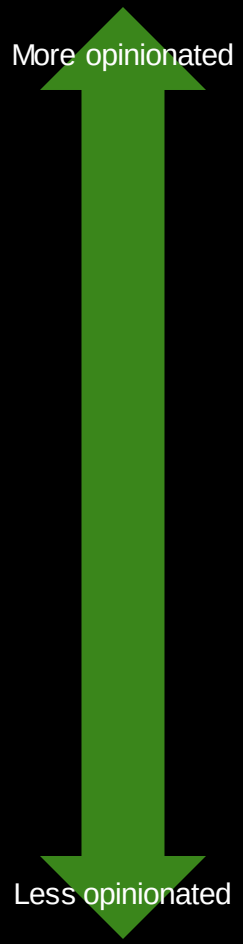

	AWS manages	Customer manages
AWS Lambda Serverless functions	• Data source integrations • Physical hardware, software, networking, and facilities • Provisioning	• Application code
AWS Fargate Serverless containers	• Container orchestration, provisioning • Cluster scaling • Physical hardware, host OS/kernel, networking, and facilities	• Application code • Data source integrations • Security config and updates, network config, management tasks
ECS/EKS Container-management as a service	• Container orchestration control plane • Physical hardware software, networking, and facilities	• Application code • Data source integrations • Work clusters • Security config and updates, network config, firewall, management tasks
EC2 Infrastructure-as-a-Service	• Physical hardware software, networking, and facilities	• Application code • Data source integrations • Scaling • Security config and updates, network config, management tasks • Provisioning, managing scaling and patching of servers

Compute node recommendations (new compute families)

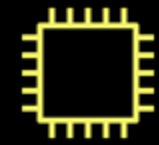

	Physical cores	-Per physical cores specification			Instance store volumes	
	(# of physical cores)	Memory (GiB)	Variable network performance (MB/s)	Consistent network performance (MB/s)	NVMe storage capacity (GiB)	Throughput performance (MB/s)
M6i/m6in 3rd Gen, Xeon	1, 2, 4, 8, 16, 24, 32, 48	8	396.18 - 1774.61	258.89	Not available	Not available
M6id/m6idn 3rd Gen, Ice lake	1, 2, 4, 8, 16, 24, 32, 48	8	396.13 - 1915.05	259.17	118 – 7,200	62.48
R6i/r6id 3rd Gen, Xeon	1, 2, 4, 8, 16, 24, 32, 48	16	396.12 - 1789.06	259.05	Not available	Not available
R6id/r6idn 3rd Gen, Ice lake	1, 2, 4, 8, 16, 24, 32, 48	16	396.14 - 1844.39	258.84	118 – 7,200	62.22
I4i 3rd Gen, Ice Lake	1, 2, 4, 6, 12, 24, 48	16	528.18 - 1745.51	259.81	468 – 30,000	154.26

Storage Overview

We’re razor-focused on providing the best file solutions to serve your usage patterns

Amazon FSx

File systems purpose-built for NAS or scale-out workloads, cloudified

1. Run NAS workloads in the cloud
2. Leverage the cloud’s virtually unlimited compute scale

Amazon EFS

Serverless, fully elastic file storage designed for builders

3. Easily share data in cloud-native apps and workflows

Amazon File Cache

High-speed cache for on-premises or AWS datasets

4. Easily burst your workloads to cloud

Amazon FSx

Amazon FSx for Lustre

Amazon FSx for Windows File Server

Amazon FSx for NetApp ONTAP

Amazon FSx for OpenZFS

Amazon FSx for Lustre

Fully managed Lustre file system for high performance workloads

Massively scalable performance

Seamless access to your data repositories

Simple and fully managed

Native file system interface

Cost-optimized for compute-intensive workloads

Secure and compliant

Customers continue to increase the size of their workloads on AWS across industry verticals and application areas

Industries and example use cases

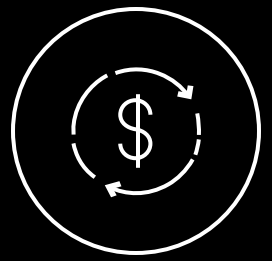

Financial services:
Modeling and analytics

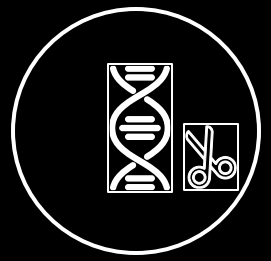

Life Sciences:
Genome analysis

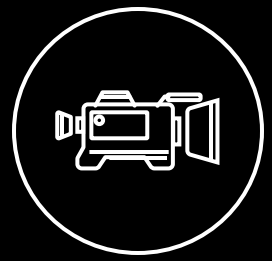

Media and Entertainment:
Rendering and transcoding

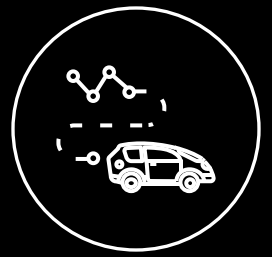

Automotive:
ECU simulations and object detection

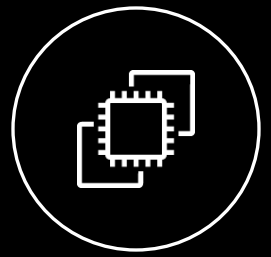

Semiconductor:
Electronic design automation

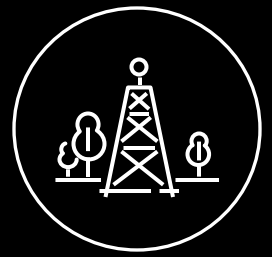

Oil and gas:
Seismic data processing

Application areas

Big data analytics

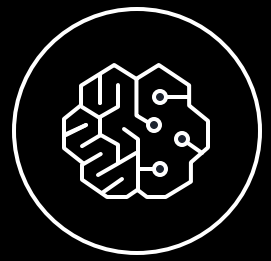

Machine learning

High-performance computing

For every $1 spent on high performance computing, businesses see $463 in incremental revenues and $44 in incremental profit[1]

FSx for Lustre deployment options

High and scalable performance

In all options, we support encryption at-rest and in-transit*

Multiple FSx for Lustre throughput options and deployment types allow customers to optimize storage cost and performance

Storage type	Baseline throughput	Price per GB-month (in IAD)	
		Persistent storage	Scratch Storage
HDD	12 MB/s/TiB	$0.025 $0.041 (with SSD cache)	-
	40 MB/s/TiB	$0.083 $0.099 (with SSD cache)	-
SSD	125 MB/s/TiB	$0.145	-
	250 MB/s/TiB	$0.210	-
	500 MB/s/TiB	$0.340	-
	200 MB/s/TIB		$0.140

- Scratch file systems are ideal for temporary storage and shorter-term processing of data.
- Data is not replicated and does not persist if a file server fails.

- File systems with SSD storage can burst up to 1.3 GB/s per TiB

Sample pricing for - US East (N. Virginia)

[1] *Prices are subject to change without notice. Pricing varies by AWS Region. For current pricing information, see the Amazon FSx for Lustre Pricing page on the AWS website.*

What is Amazon FSx for NetApp ONTAP?

Fully-featured NetApp ONTAP

With the simplicity, agility, and scalability of an AWS service

Amazon FSx for ONTAP: Benefits

Launch and run fully managed file storage built on NetApp ONTAP

Familiar and fully managed

- Fully Managed
- Use AWS and NetApp tools
- Data replication (SnapMirror),snapshot (SnapVault), caching (FlexCache), and cloning (FlexClone) capabilities

Fast and cost effective

- Multiple GB/s throughput, 100K+ IOPS, sub-ms latencies
- Automatic tiering to low cost, reduce costs 90%
- Deduplication, compression, compaction, thin provisioning

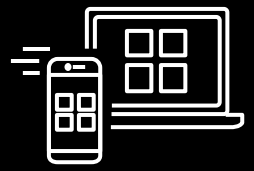

Accessible and Integrated

- Multi-protocol (NFS, SMB, iSCSI)
- Accessible from Linux, Windows, MacOS
- EC2, EKS, Workspaces, Appstream 2.0, VMware Cloud
- Concurrent, multi-protocol access

Secure and Compliant

- Encrypted at-rest and in-transit
- Integration with software for Anti-Virus and Auditing
- Active Directory for identity-based authentication
- ISO, PCI-DSS, SOC compliant and HIPPA eligible

Amazon FSx for NetApp ONTAP: Automatic performance and cost optimization

Intelligent policy-based data movement between tiers

Primary Tier

SSD, Multi-AZ

Up to 192 TB

Optimized for performance

~20%

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

Capacity Pool Tier

Elastic, Multi-AZ

Unlimited capacity (PB+ file systems)

Cost-optimized for less accessed files

~80%

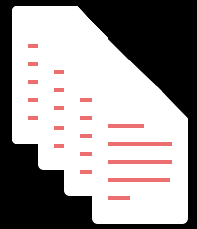

Automated Tiering Policies

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

Amazon FSx for ONTAP: Getting Started

Backup, DR Migration

SnapVault
SnapMirror

On-prem ONTAP

Hybrid: Cloud bursting

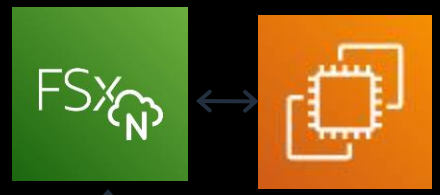

In-AWS compute
(e.g., EC2)

FlexCache
origin (source)
volume

On-prem ONTAP

Hybrid: On-prem caching

Source volume

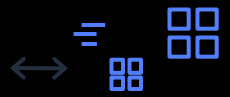

On-prem FlexCache or
Global File Cache

On-prem
compute

Cross-region

Replication (in-cloud DR)

SnapVault
SnapMirror

Region A

Region B

Low-latency caching

FlexCache

Region A

Region B

Amazon FSx for NetApp ONTAP: Pricing Dimensions

Provisioned pricing dimensions

- SSD storage ($0.250 per GB-month)
- Throughput capacity ($1.200 per MBps-month)
- [Optional] SSD IOPS $0.0340 per IOPS-month)

Elastic pricing dimensions

- Capacity pool storage ($0.0438 per GB-month stored)
- Capacity pool requests $0.0004 per 1,000 read requests
- Capacity pool requests $.005 per 1,000 write requests
- Backup storage ($0.05 per GB-month stored)

Introducing Amazon FSx for OpenZFS

FSx for ZFS

Shared file storage that delivers **high speeds at a low cost**, accessible through NFS

Built on the **AWS Graviton** family of processors and the popular open-source **OpenZFS** file system

Deliver results faster with high-performance storage

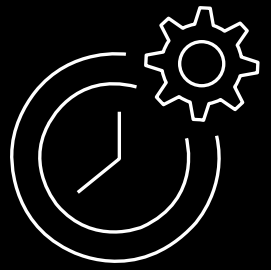

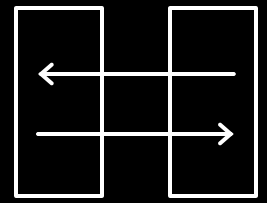

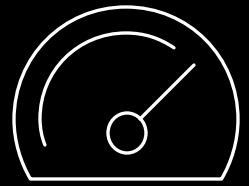

	Latency	IOPS	Throughput
What is it?	Average time to return the first byte of data	Number of operations (I/O) per second	Amount of data transferred per second
What can **FSx for OpenZFS** deliver?	Latencies of a **few hundred microseconds**	Up to **1 million IOPS**	Up to **12.8 GB/s**
When is it important?	Sequential transaction-heavy applications like code repositories (Git) and content management systems (Wordpress)		Scalable data-intensive workloads like data and analytics, as well as ML

Do more with less and adapt faster to changing business needs

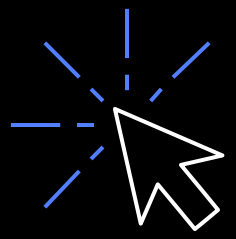

Storage and performance scaling in minutes*

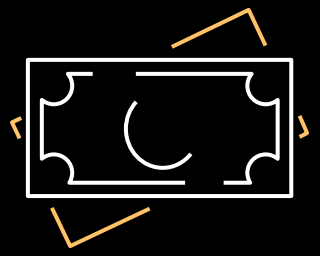

Zstandard compression (reduce storage usage by up to ~50%)

SSD storage

$0.09/GB-month

($0.045/GB-month w/ compression*)

Throughput

$0.26/ MBps-month

*Pricing assumes average compression savings of ~50% and is an effective price.
*Pricing is for Single-AZ

SAS Viya on AWS

SAS Viya on AWS with Amazon Elastic File System (IaC)

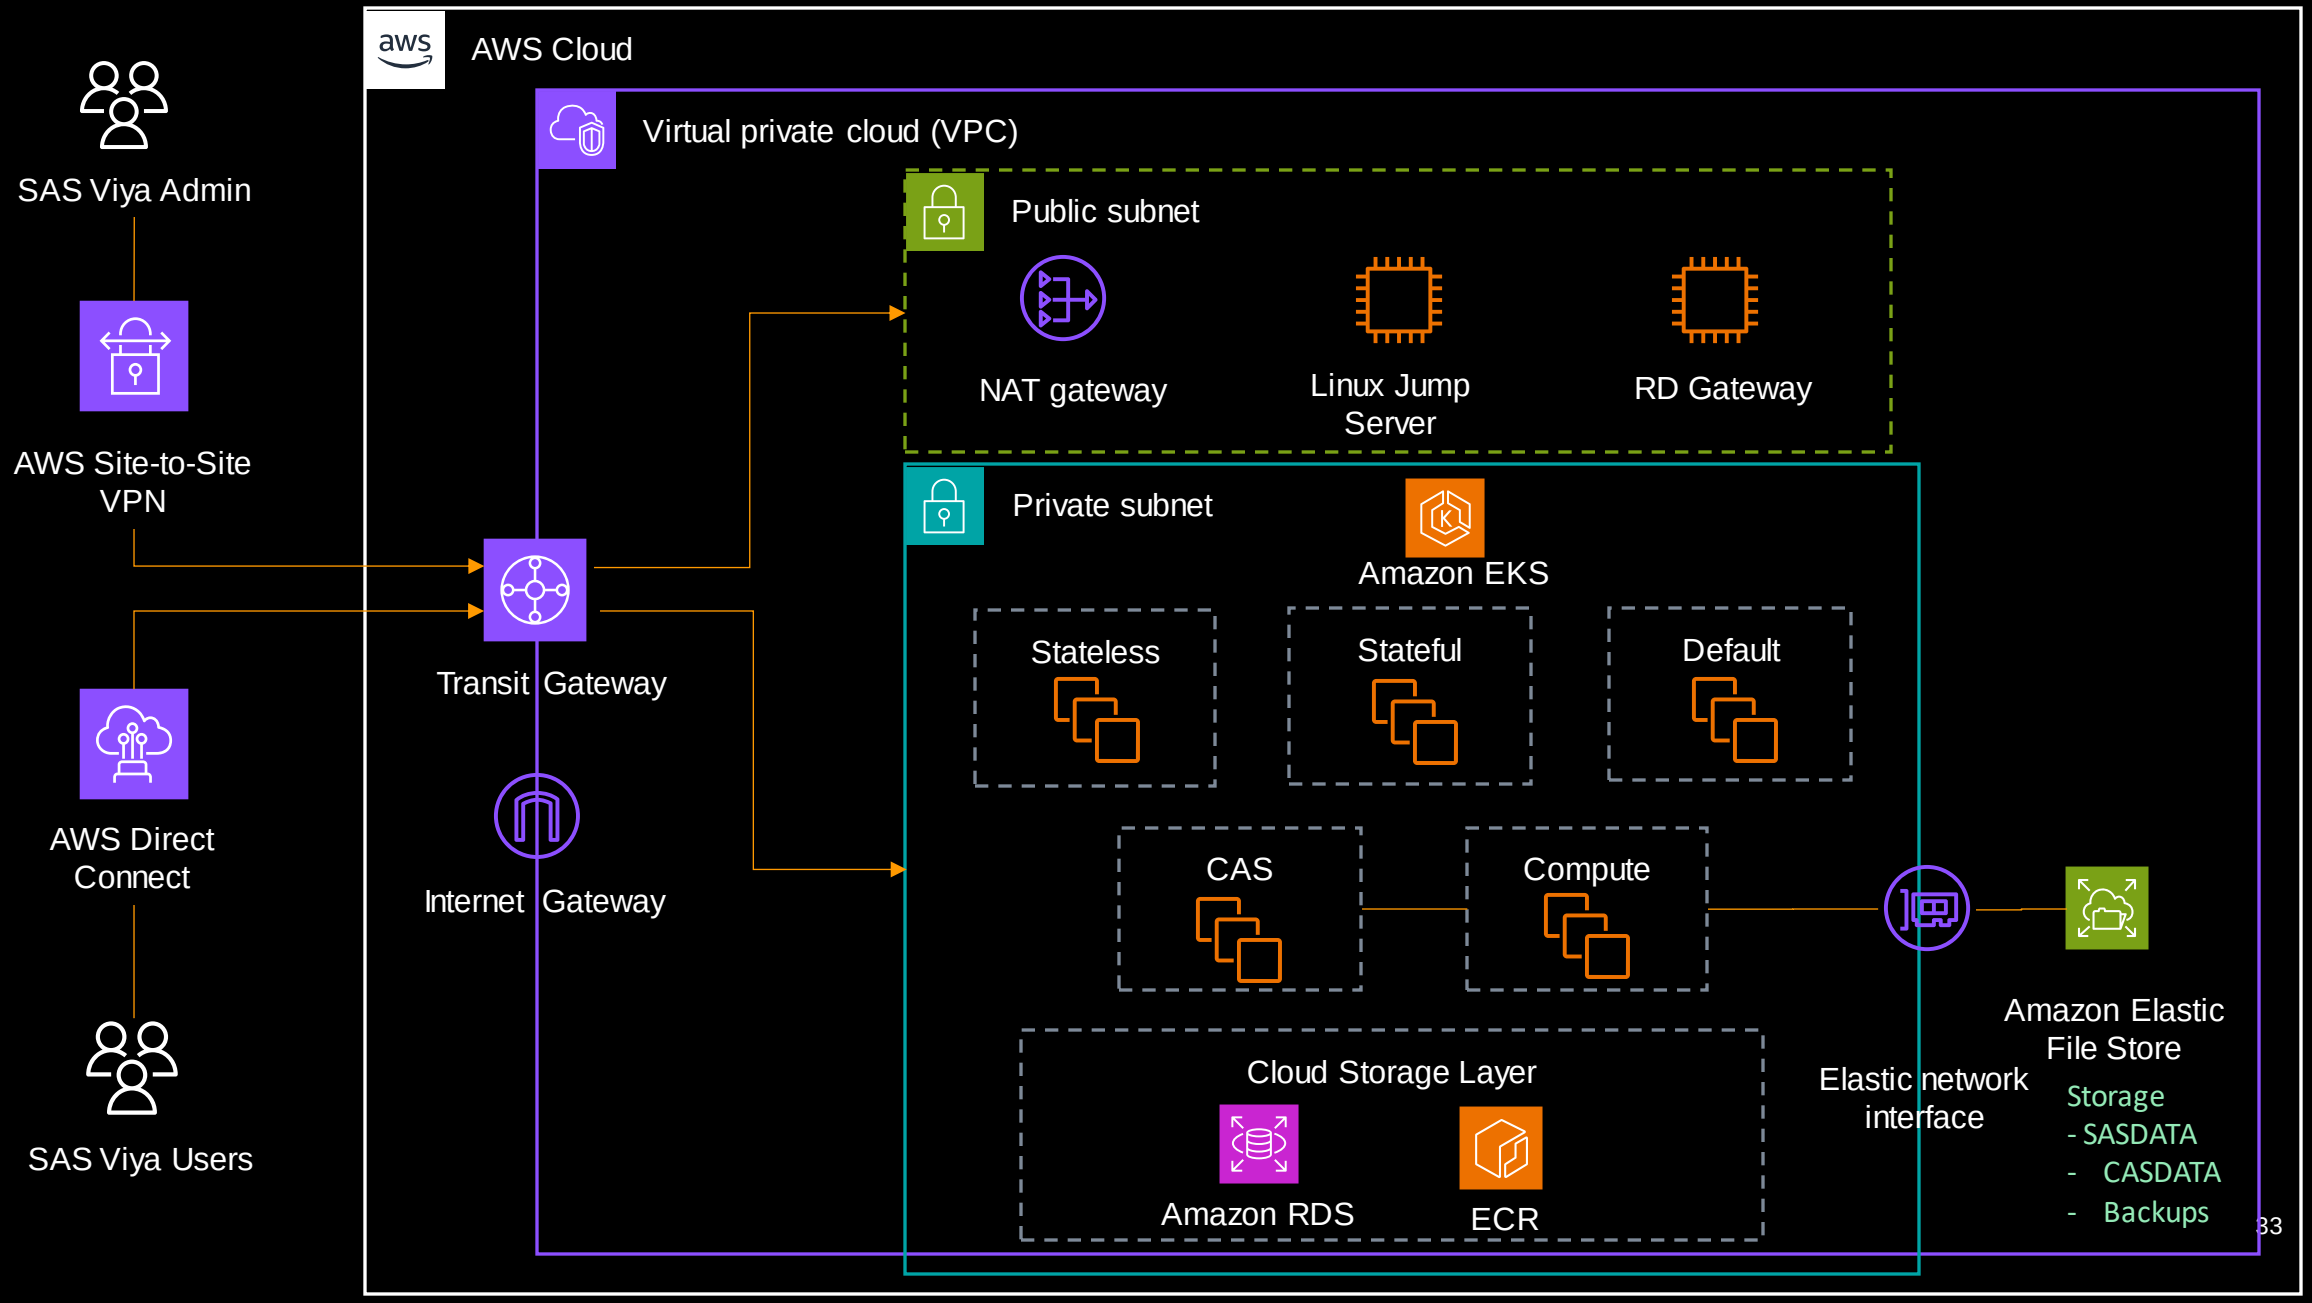

SAS Viya on AWS with FSx for NetApp (IaC)

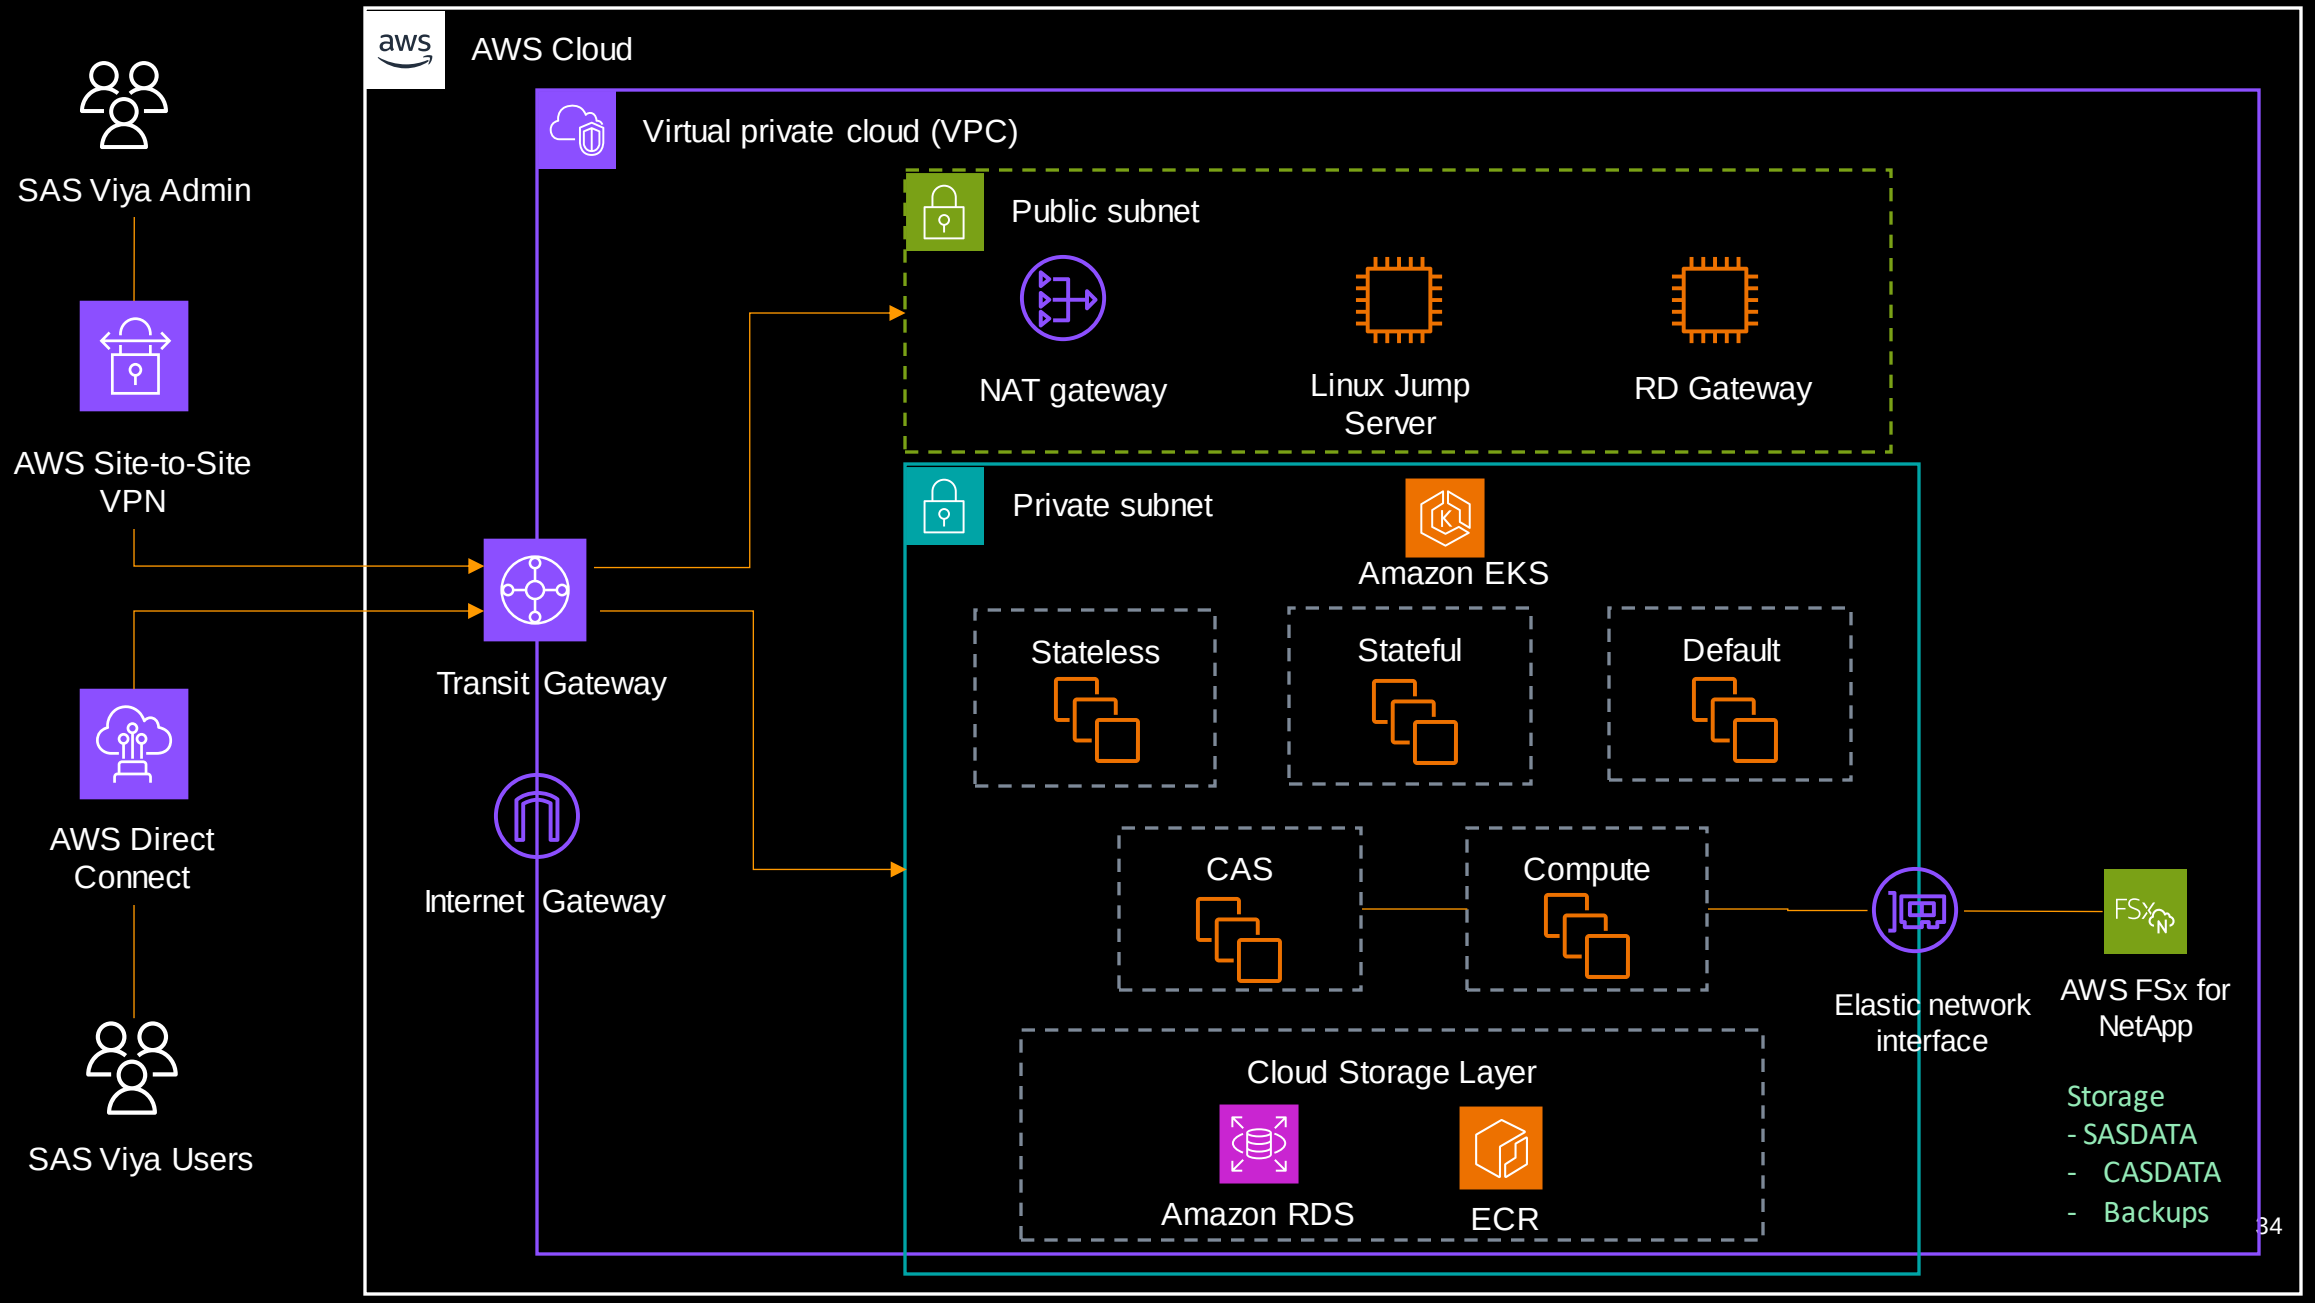

SAS Viya on AWS with FSx for Lustre (Manual deployment)

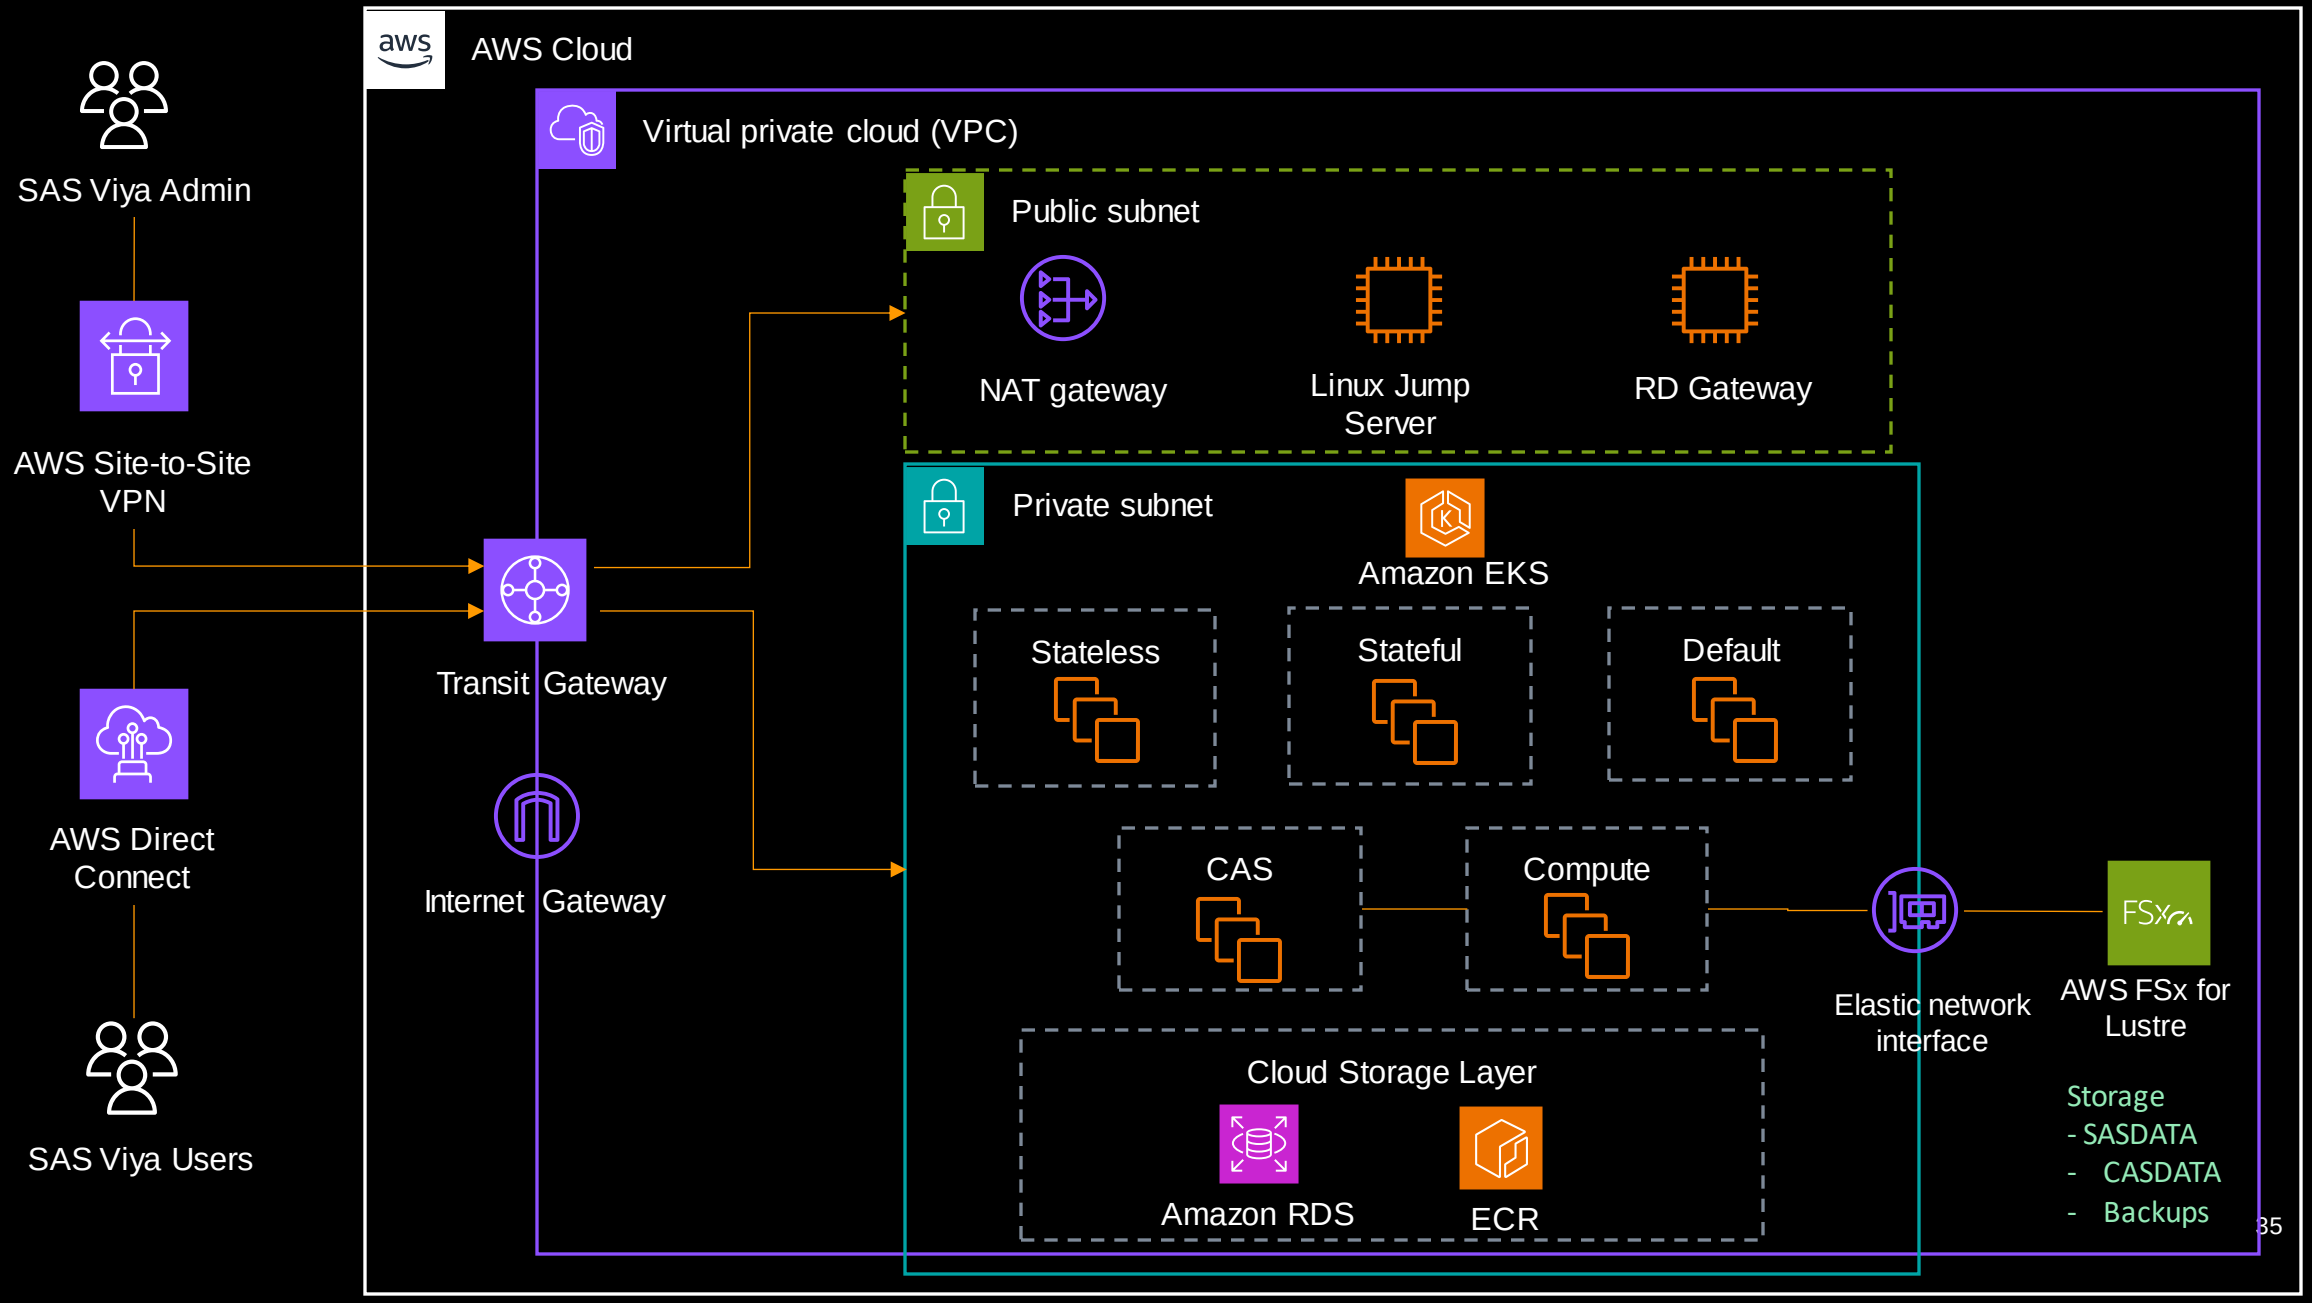

SAS Viya with FSx Ontap Demo

Resources

Infrastructure as Code: https://github.com/sassoftware/viya4-iac-aws

Coming soon – Blog and prescriptive guidance

Thank you

Dilip Rajan – rajand@amazon.com

Darryl Osborne – darrylo@amazon.com