AWS

SUMMIT

Deep Dive on Amazon Elastic File System

Darryl S. Osborne
AWS - Storage Specialist Solutions Architect
May 18, 2017

amazon web services

Amazon Elastic File System (EFS)

Provides simple, scalable, highly available & durable file storage in the cloud

Petabyte scale file system distributed across an unconstrained number of storage servers in multiple Availability Zones (AZs)

Elastic capacity, automatically growing & shrinking as you add & remove files

Amazon Elastic File System (EFS)

Standard file system interface & semantics

Shared storage

Highly available & highly durable

Consistent low latency

Strong read-after-write consistency

Elastic capacity

Fully managed

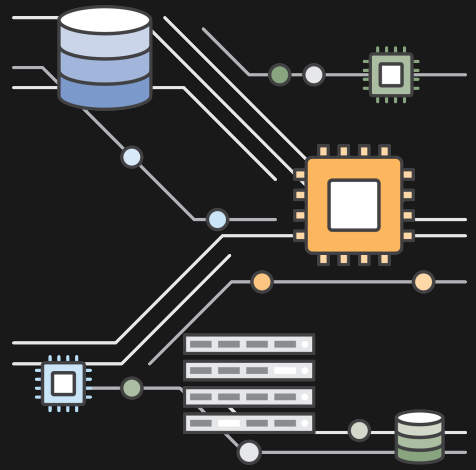

The AWS Storage Portfolio

Amazon EBS
(persistent)

Amazon EC2
Instance Store
(ephemeral)

Block

Amazon EFS

File

Amazon S3

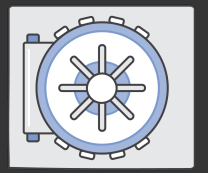

Amazon Glacier

Object

Cloud Data Migration

Snow* data
transport
family

Storage
Gateway

Direct
Connect

3rd Party
Connectors

Transfer
Acceleration

Kinesis Firehose

What customers are using EFS for today

- Web serving
- Content management
- Analytics
- Database backups
- Media and Entertainment workflows
- Container storage
- Workflow management
- Home directories

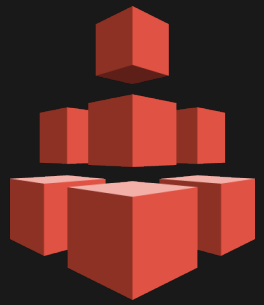

Shared File Solutions in the Cloud… before EFS

3rd Party Software

3rd Party Hardware in AWS Direct Connect locations

Do It Yourself

Do It Yourself – NFS Architecture

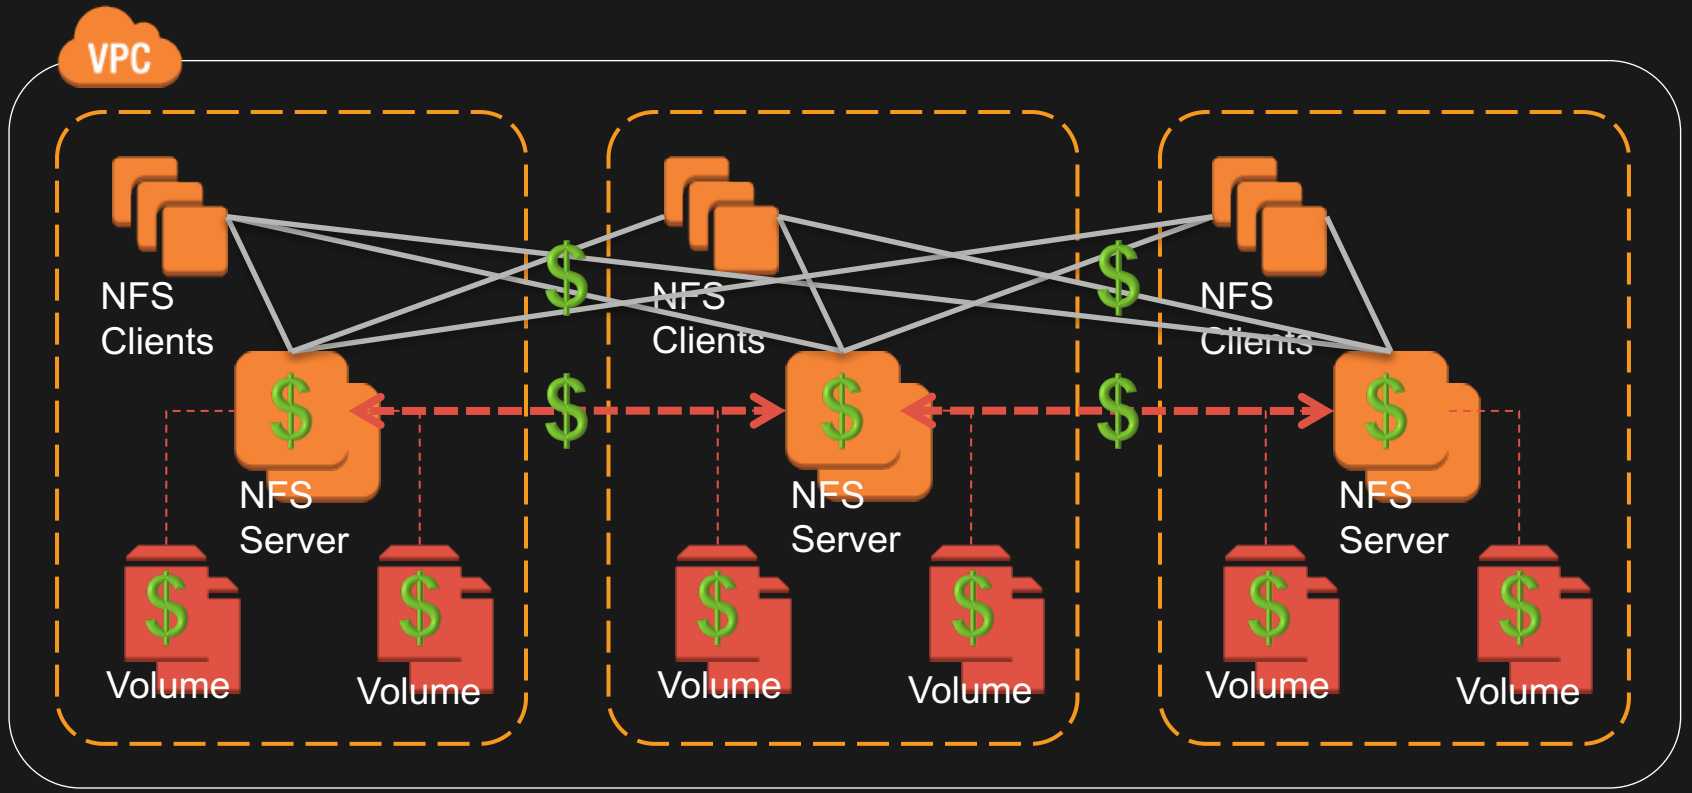

Do It Yourself – NFS Architecture

- ❑ Launch, patch, monitor, & pay for EC2 instances
- ❑ Create, attach, monitor, & pay for provisioned EBS volumes
- ❑ Create, maintain, and monitor auto scaling group
- ❑ Install, patch, monitor, & pay for* file system software
- ❑ Configure, maintain, monitor, & pay for file system data intra/inter-AZ replication
    - IOPS for replication are still IOPS
- ❑ Configure DNS for client HA access to inter-AZ NFS fleet

Do It Yourself

Do It Yourself NFS Architecture

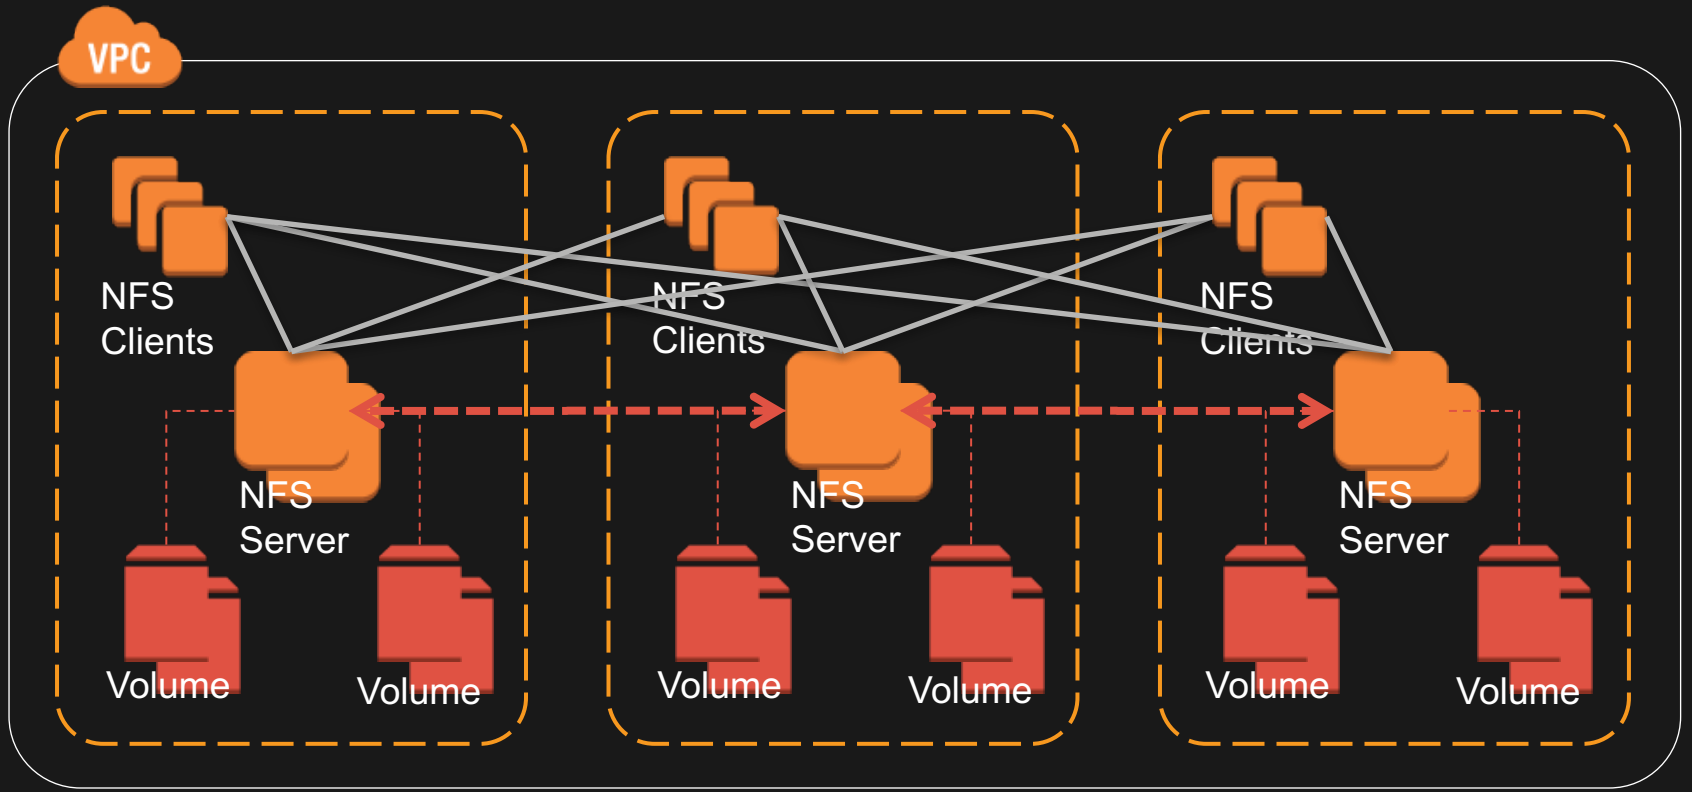

Amazon EFS Architecture

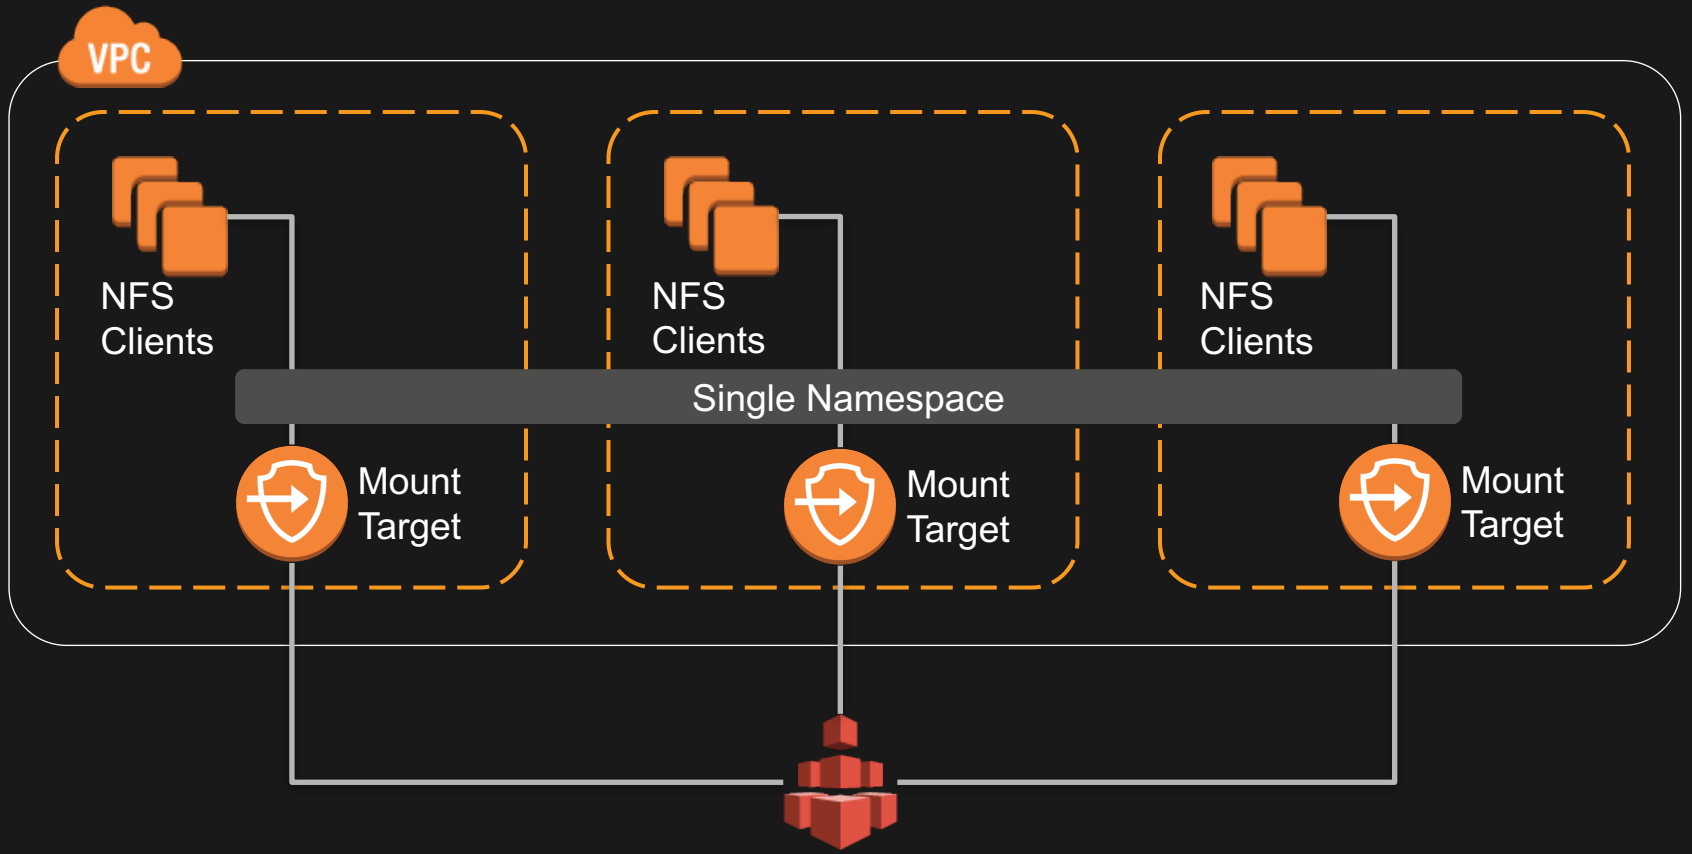

Do It Yourself – Cost

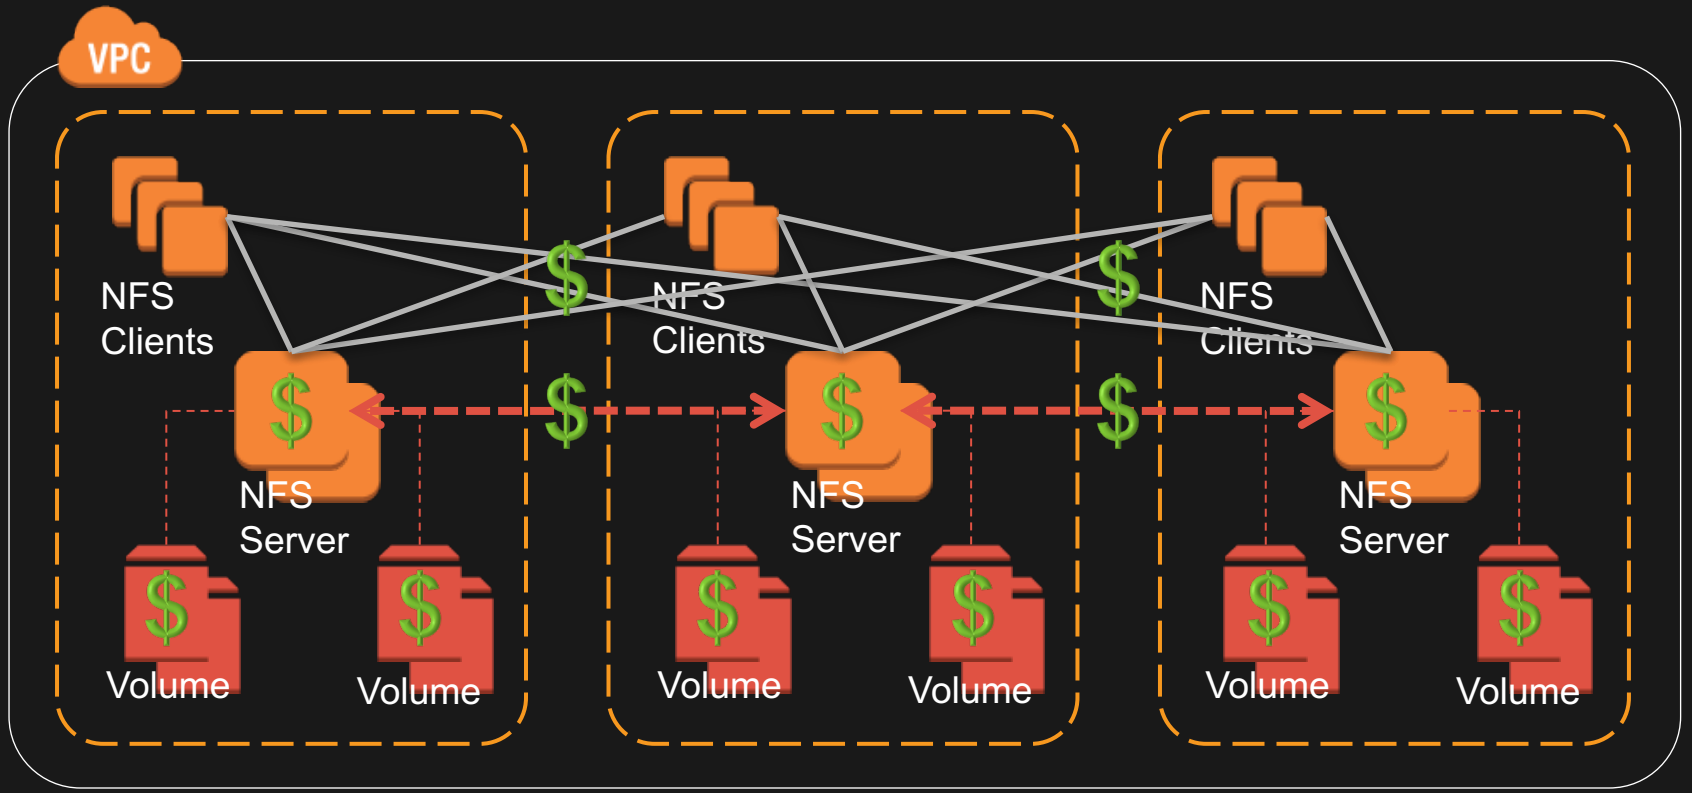

Amazon EFS Architecture

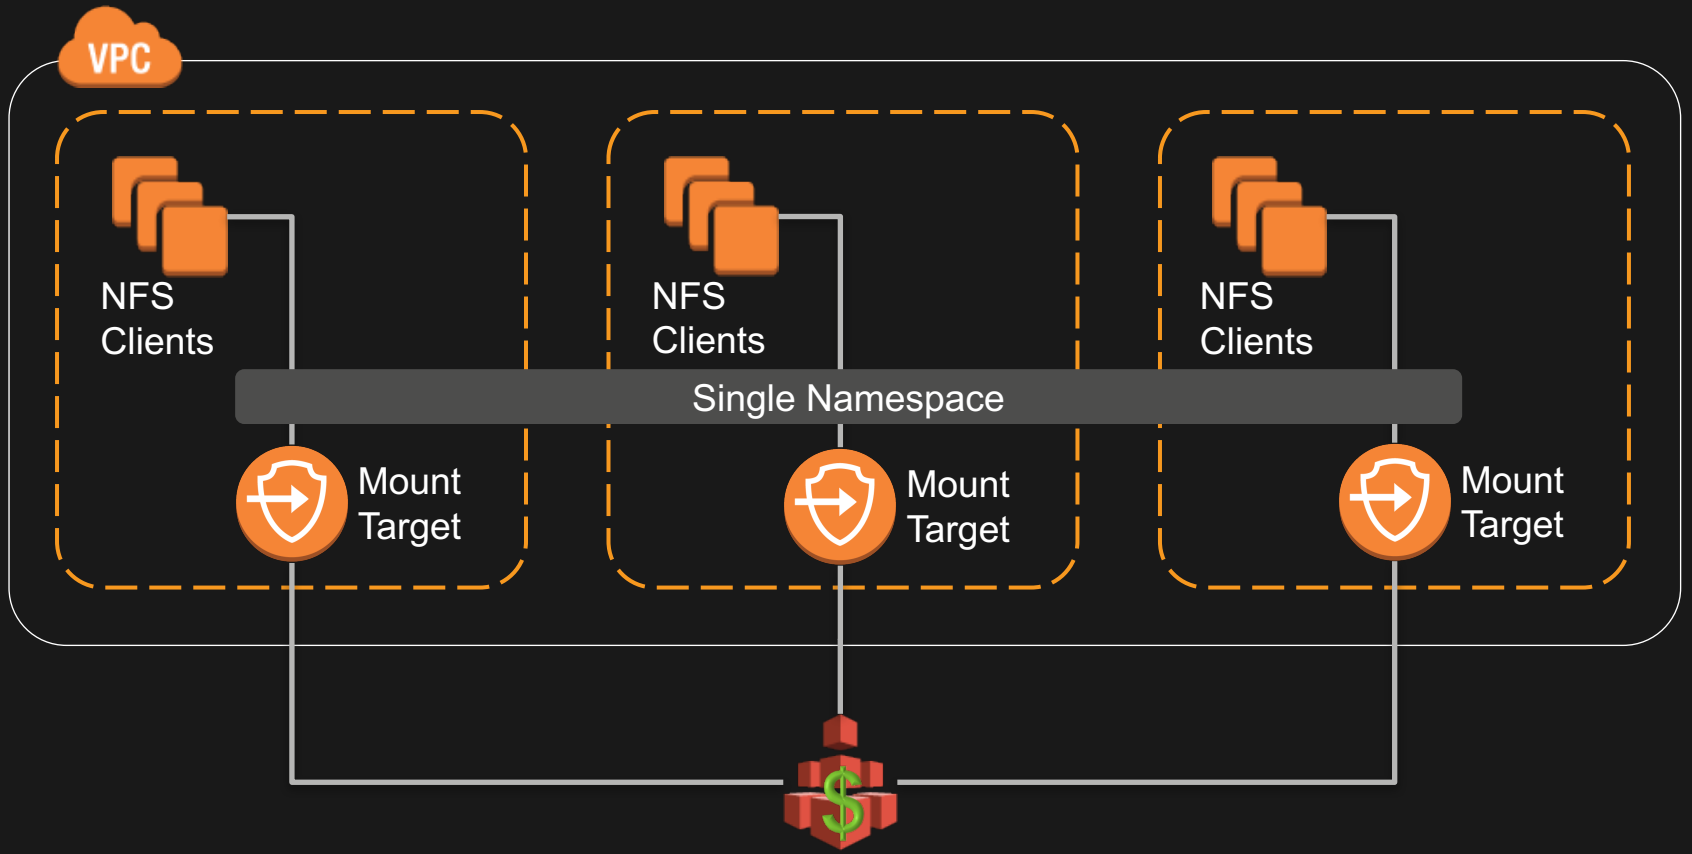

Do you need an *EFS* file system?

If you have an application running on EC2 or use case that requires a file system…

AND

- Requires multi-attach ***OR***
- GBs/s throughput ***OR***
- Multi-AZ availability/durability ***OR***
- Requires automatic scaling (grow/shrink) of storage

Where is EFS available today?

- **US West** (Oregon)
- **US East** (N. Virginia)
- **US East** (Ohio)
- **EU** (Ireland)
- **Asia Pacific** (Sydney)

More coming soon!

Hands-on: Create an EFS File System (Console)

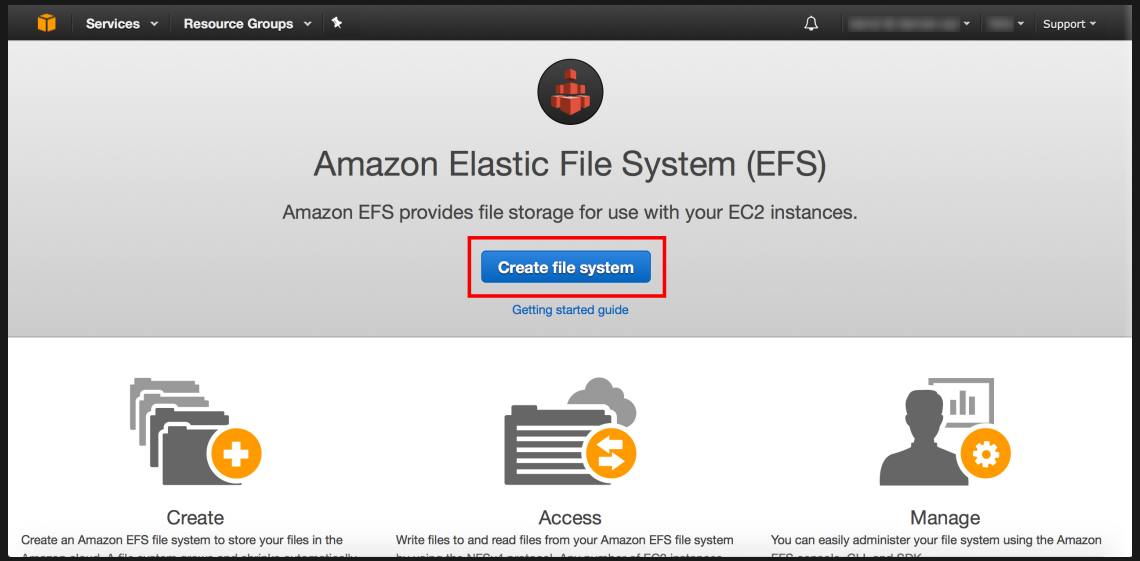

Resources for Amazon EFS

File System

- Regional construct
- Ten per account per region (soft)
- Default throughput limit 3 GB/s (soft)
- Metered size updates approx. every hour
- Accessible from EC2
  - VPC, EC2-Classic via ClassicLink
- Accessible from on-premises
  - AWS Direct Connect

Resources for Amazon EFS

File System cont…

- Scenarios for on-prem via Direct Connect

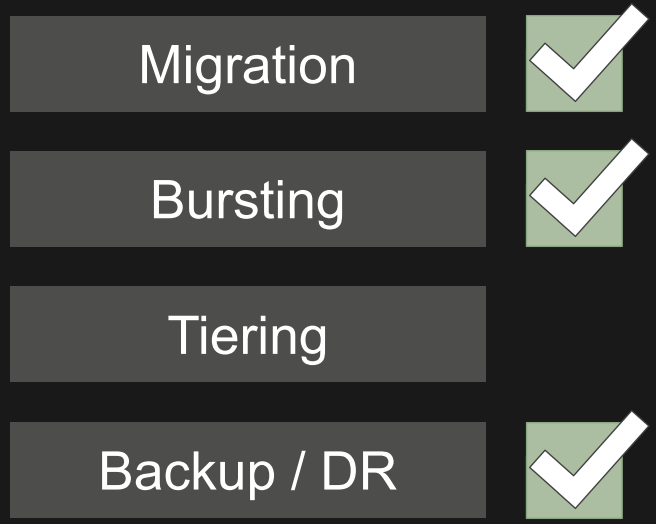

Resources for Amazon EFS

Mount Targets

- One or more per file system
- Create in a VPC Subnet
- One per Availability Zone
- Must be in the same VPC

Resources for Amazon EFS

Security Groups

- Standard VPC Security Group
- Same VPC as subnet
- Up to five per mount target
- Allow inbound TCP port 2049 from NFS clients

Mount EFS

NFSv4.0

NFSv4.1

Linux Kernel 4+

Hands-on: Mount an EFS File System (Console)

Launch EC2 instance from EC2 Console

Connect to the instance

Make a directory

Mount EFS file system

Query disk file system & file system table

- df; df -hT; df -h -t nfsv4; mount -t nfsv4

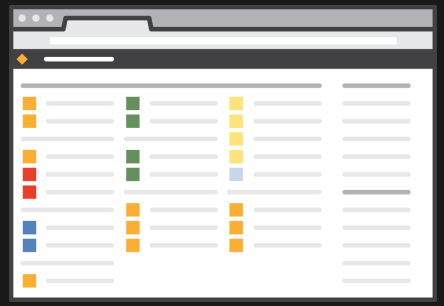

Amazon EFS Architecture

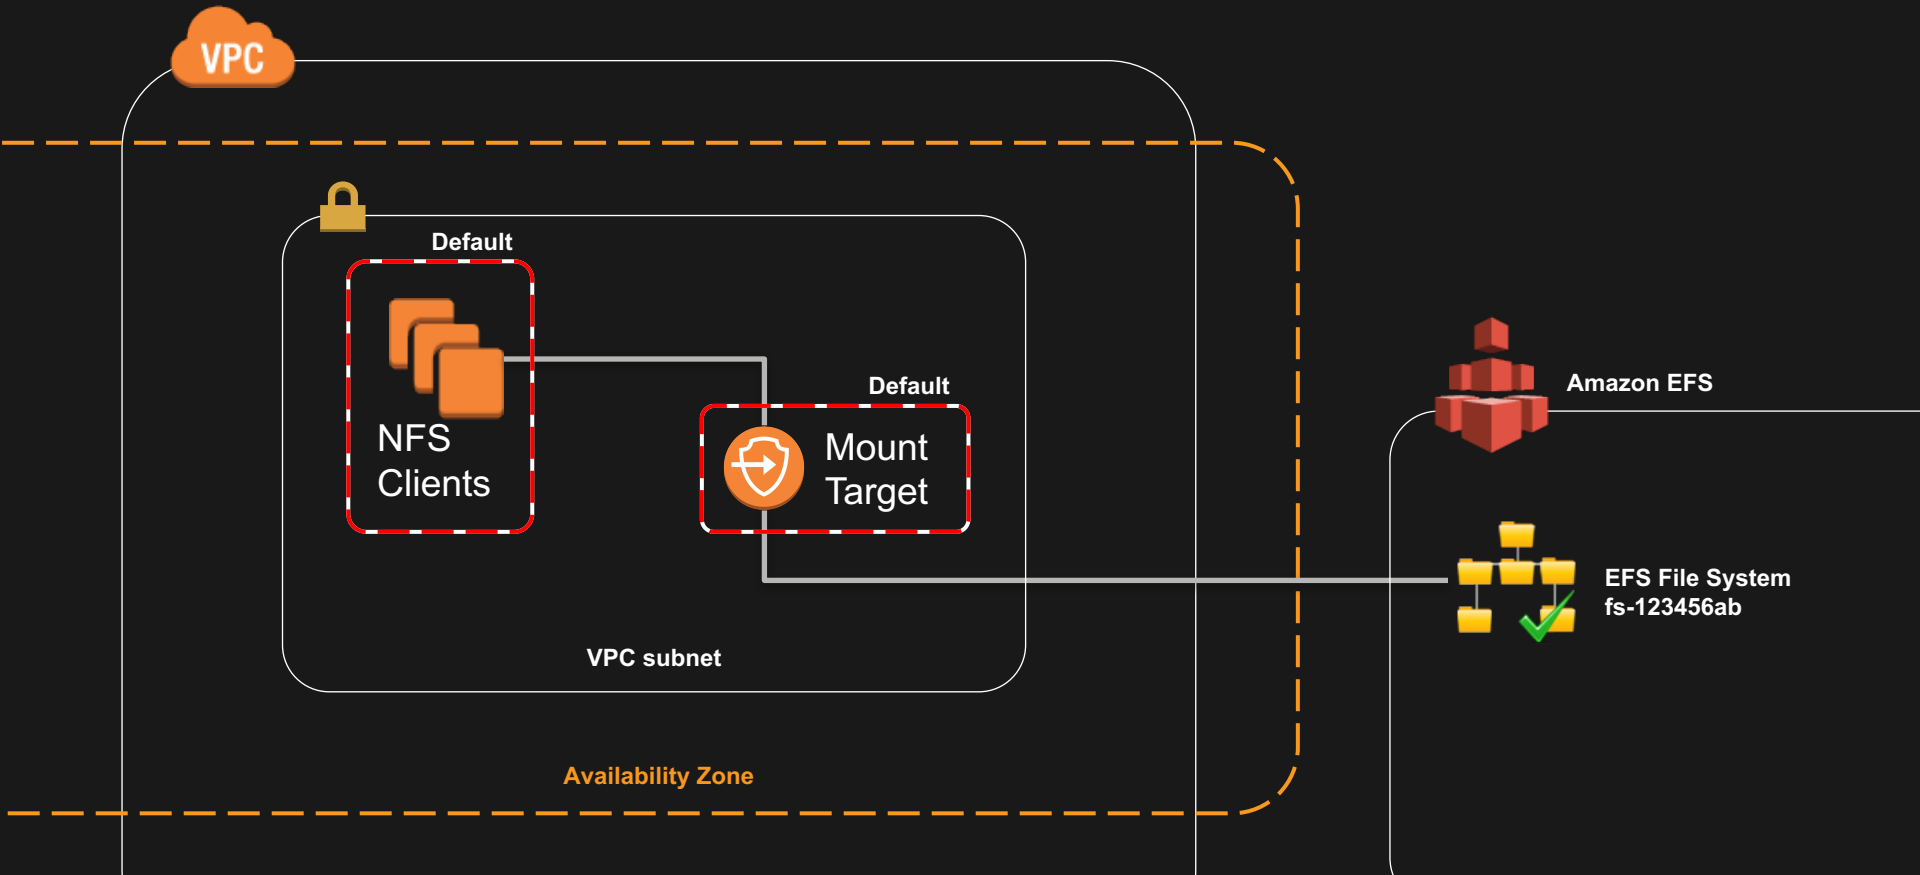

Amazon EFS Architecture

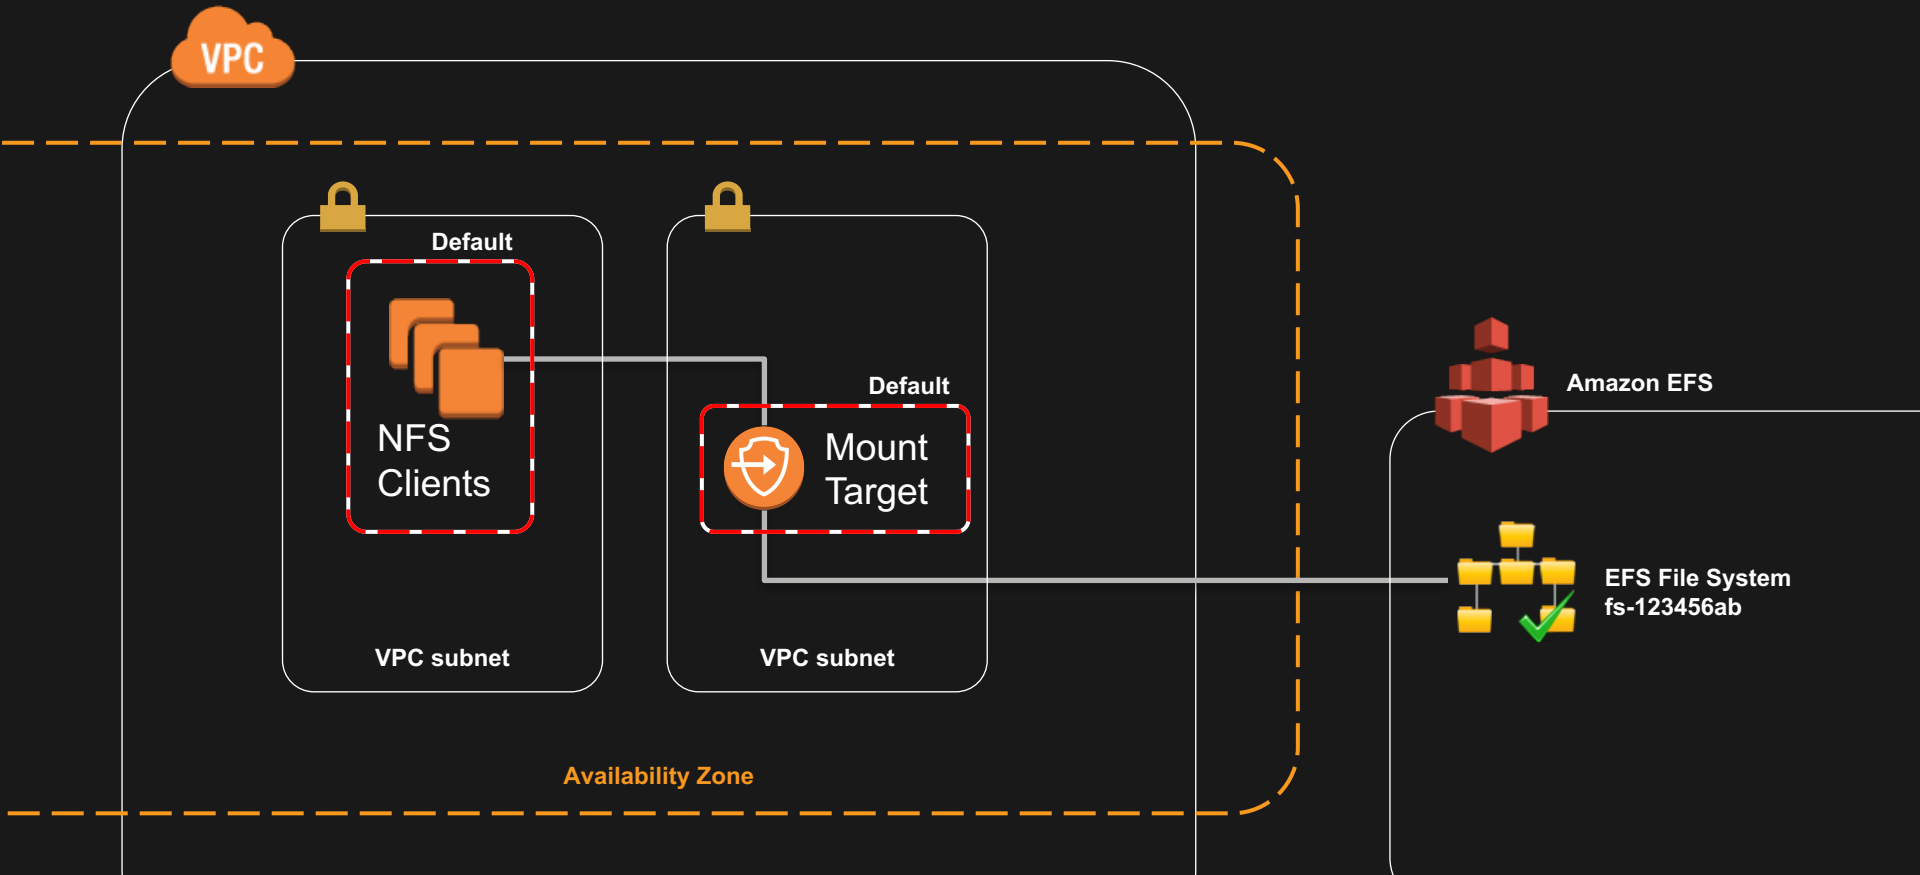

Amazon EFS Architecture

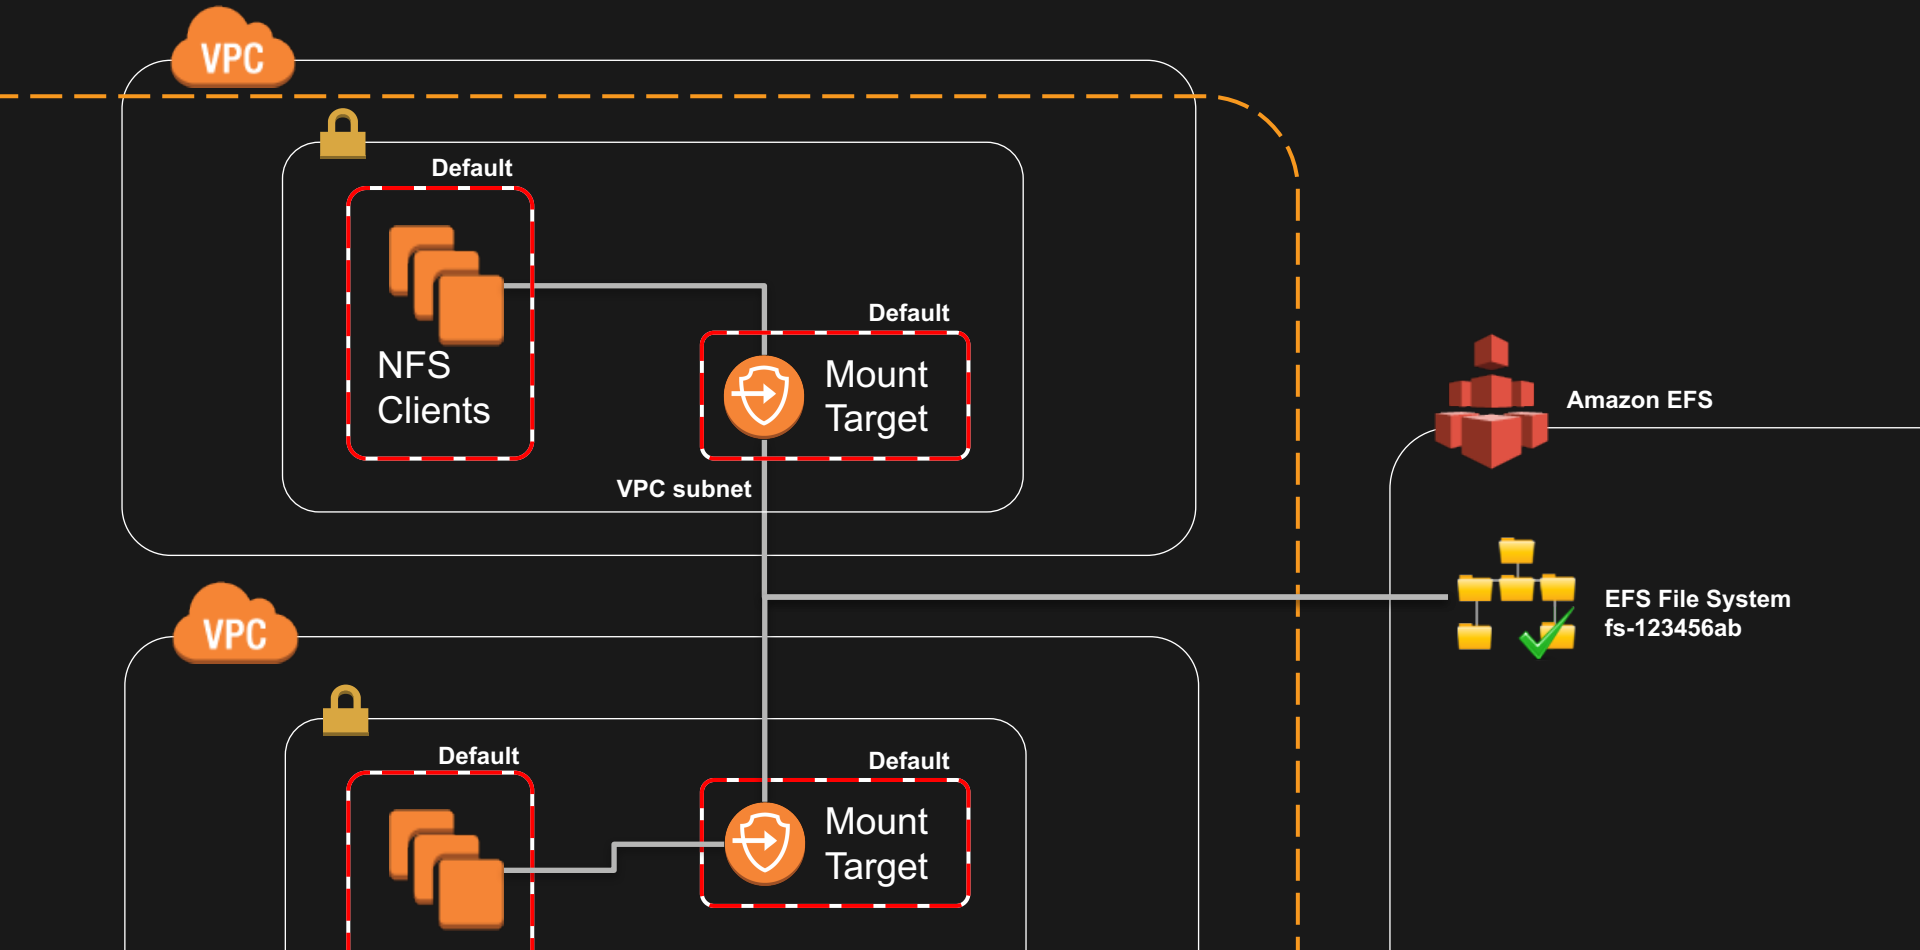

Security

Control network traffic using VPC security groups and network ACLs

Control file and directory access by using POSIX permissions

Control administrative access (API access) to file systems by using AWS Identity and Access Management (IAM)

action-level and *resource-level* permissions

Amazon EFS is designed for wide spectrum of performance needs

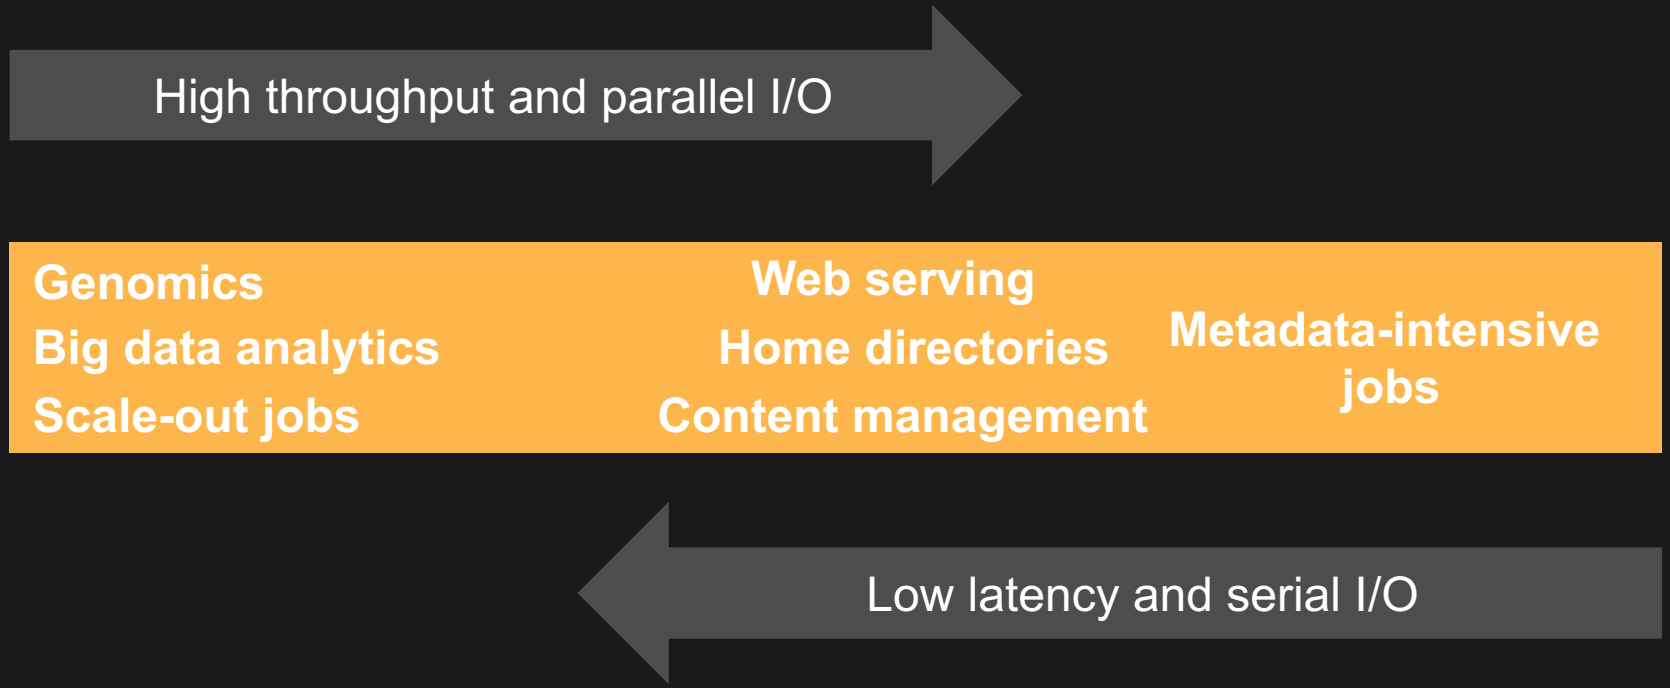

Performance modes for different workloads

Mode	What's it for	Advantages	Tradeoffs	When to use
General purpose (default)	Latency-sensitive applications and general-purpose workloads	Lowest latencies for file operations	Limit of 7K ops/sec	Best choice for most workloads
Max I/O	Large-scale and data-heavy applications	Virtually unlimited ability to scale out throughput / IOPS	Slightly higher latencies	Consider for large scale-out workloads

Amazon EFS - distributed data storage design

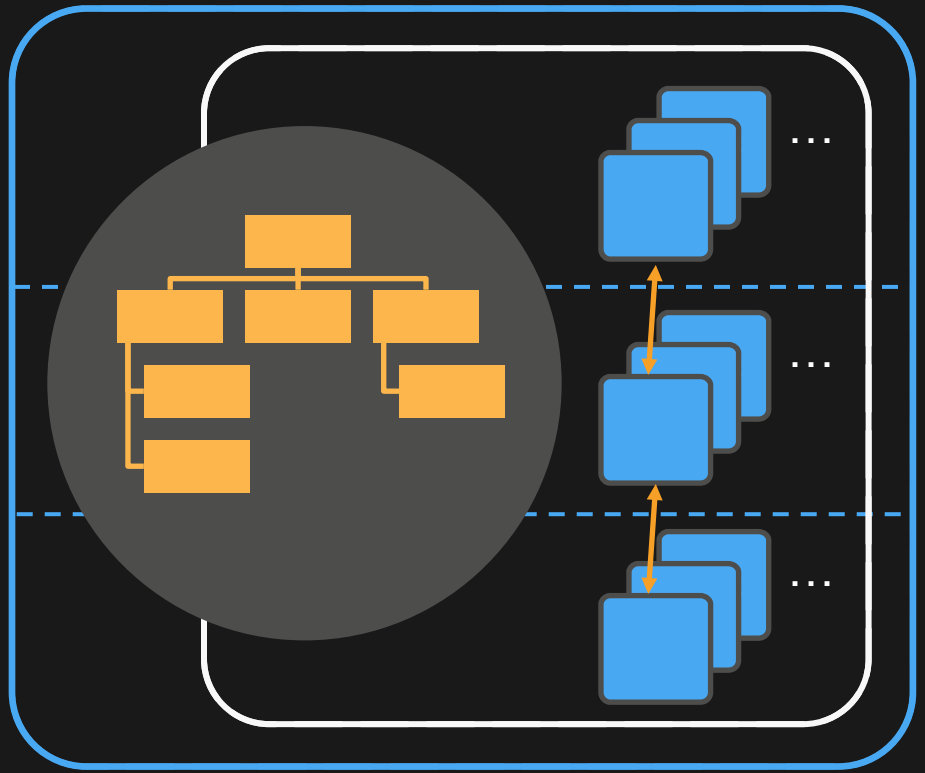

File systems distributed across unconstrained number of servers

- Avoids bottlenecks/constraints of traditional file servers
- Enables high levels of aggregate IOPS/throughput

Data also distributed across Availability Zones (durability, availability)

How to think about EFS perf relative to EBS

		Amazon EFS	Amazon EBS PIOPS
Performance	Per-operation latency	Low, consistent	Lowest, consistent
	Throughput scale	Multiple GBs per second	Single GB per second
Characteristics	Data availability / durability	Stored redundantly across multiple AZs	Stored redundantly in a single AZ
	Access	1 to 1000s of EC2 instances, from multiple AZs, concurrently	Single EC2 instance in a single AZ
	Use cases	Big Data and analytics, media processing workflows, content management, web serving, home directories	Boot volumes, transactional and NoSQL databases, data warehousing & ETL

I/O Size Implication

I/O size impacts throughput of serialized operations

Parallelize

Take advantage of EFS's distributed architecture

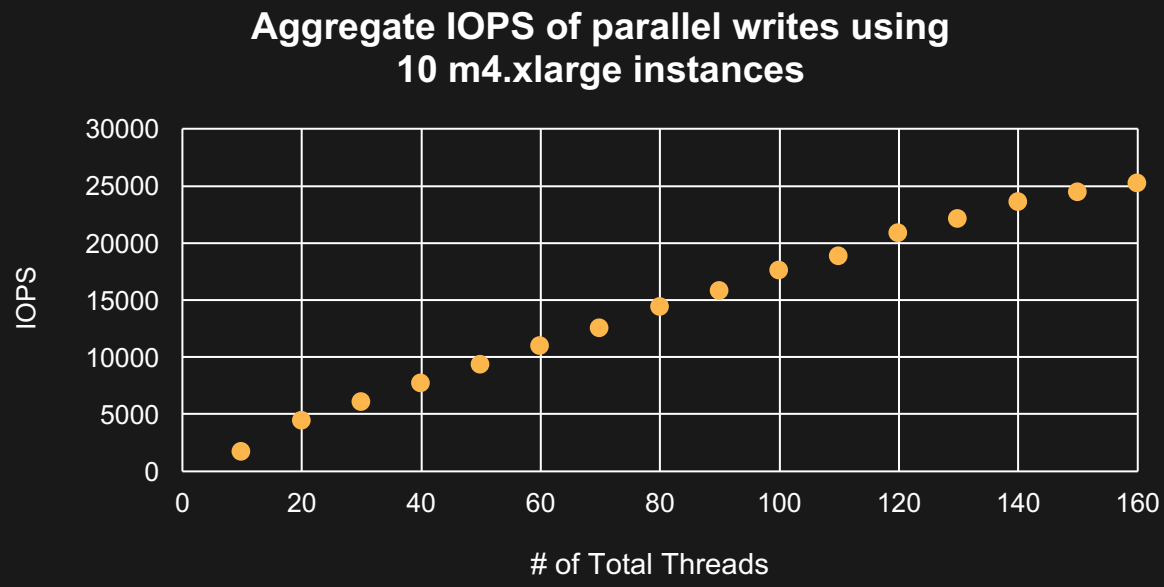

Parallelize via multiple threads and/or multiple instances

Previous Scalability Test

Large files – 50 instances

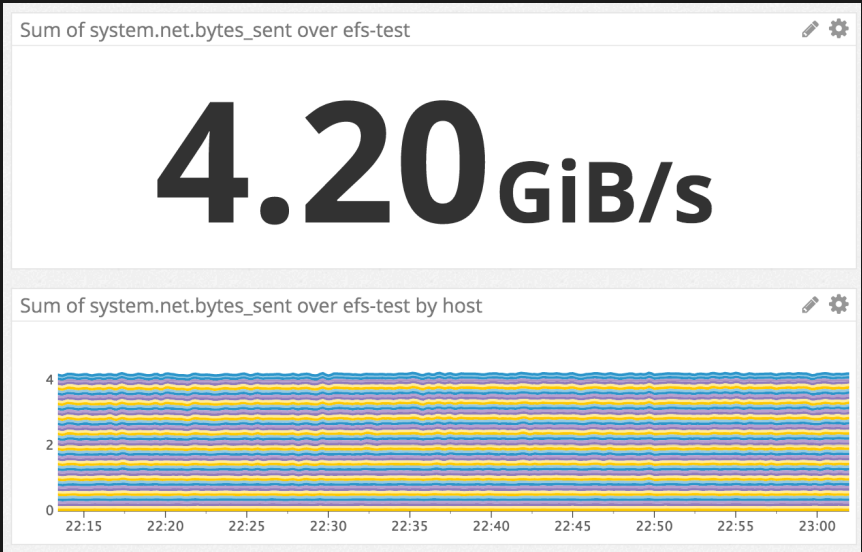

Small files – 300 instances

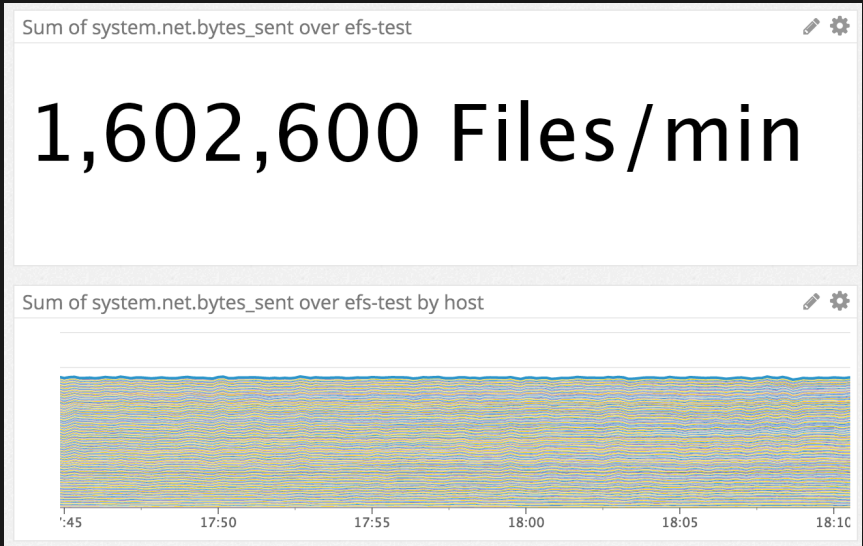

EFS CloudWatch Metrics

- DataReadIOBytes
- DataWriteIOBytes
- MetaDataIOBytes
- TotalIOBytes
- BurstCreditBalance
- PermittedThroughput
- ClientConnections
- PercentIOLimit

Amazon CloudWatch

EFS Economics

No minimum commitments or up-front fees

No need to provision storage in advance

No other fees, charges, or billing dimensions

Price: $0.30/GB-Month (US Regions)
$0.33/GB-Month (EU Ireland)
$0.36/GB-Month (AP Sydney)

EFS TCO example

Let’s say you need to store ~500 GB and require high availability and durability

Using a shared file layer on top of EBS, you might provision 600 GB (with ~85% utilization) and fully replicate the data to a second Availability Zone for availability/durability

Example comparative cost:

Storage (2x 600 GB EBS gp2 volumes):	**$120 per month**
Compute (2x m4.xlarge instances):	**$350 per month**
Inter-AZ data transfer costs (est.):	**$129 per month**
Total	**$599 per month**

EFS cost is (500GB * $0.30/GB-month) = **$150 per month**, with no additional charges

Key recommendations

- Test your application!
- Use General Purpose mode for lowest latency, Max-I/O for scale-out
- Use Linux kernel version 4.0 or newer, mount via NFSv4.1
- To optimize, look for opportunities to:
    - Aggregate I/O
    - Perform async operations
    - Parallelize
    - Cache
- Don’t forget to check your burst credit earn/spend rate when testing – ensure sufficient amount of storage

Reference

AWS Loft EFS Hands-on Walk-through - https://bit.ly/awsloft2017

AWS 10-minute Tutorials - https://aws.amazon.com/getting-started/tutorials/

Amazon EFS Web page - https://aws.amazon.com/efs/

YouTube AWS Channel - https://www.youtube.com/user/AmazonWebServices

Reference Architecture - https://aws.amazon.com/architecture/

QuickStarts - https://aws.amazon.com/architecture/

*qwik*LABS - https://aws.qwiklabs.com/

AWS

SUMMIT

Danke!

Darryl S. Osborne
darrylo@amazon.com

amazon web services