

AWS re:Invent

NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV



STG307-R

Amazon EBS: A tech deep dive

Prarthana Karmakar

Software Development Manager

AWS



AWS storage portfolio



AWS Backup



Amazon EBS snapshots



AWS Transfer Family

Data services

Intelligent-Tiering
Lifecycle management
Compression and dedupe
Replication
Metrics reporting


Core services

Object



Amazon S3 and Amazon S3 Glacier

Block



Amazon EBS

File

FSx

Amazon FSx family



Amazon EFS



AWS DataSync



AWS Storage Gateway

Hybrid/edge storage Data movement services



AWS Snow Family



AWS Outposts



Amazon Elastic Block Store (Amazon EBS) is an easy-to-use, secure, high-performance block storage service designed for use with Amazon EC2. Amazon EBS is optimized for throughput and transaction-intensive workloads at any scale.



Why do customers use Amazon EBS?



High performance

Up to 256,000 IOPS and 4,000 MiB/s throughput with single-digit millisecond latency



Reliability

99.999% availability and between 0.001% and 0.2% AFR



Scalability

Use a single gigabyte or scale up to petabytes of data for dynamic workloads such as stateful containers



Easy to use

Easily increase capacity, performance, switch volume types, and manage backups



Security

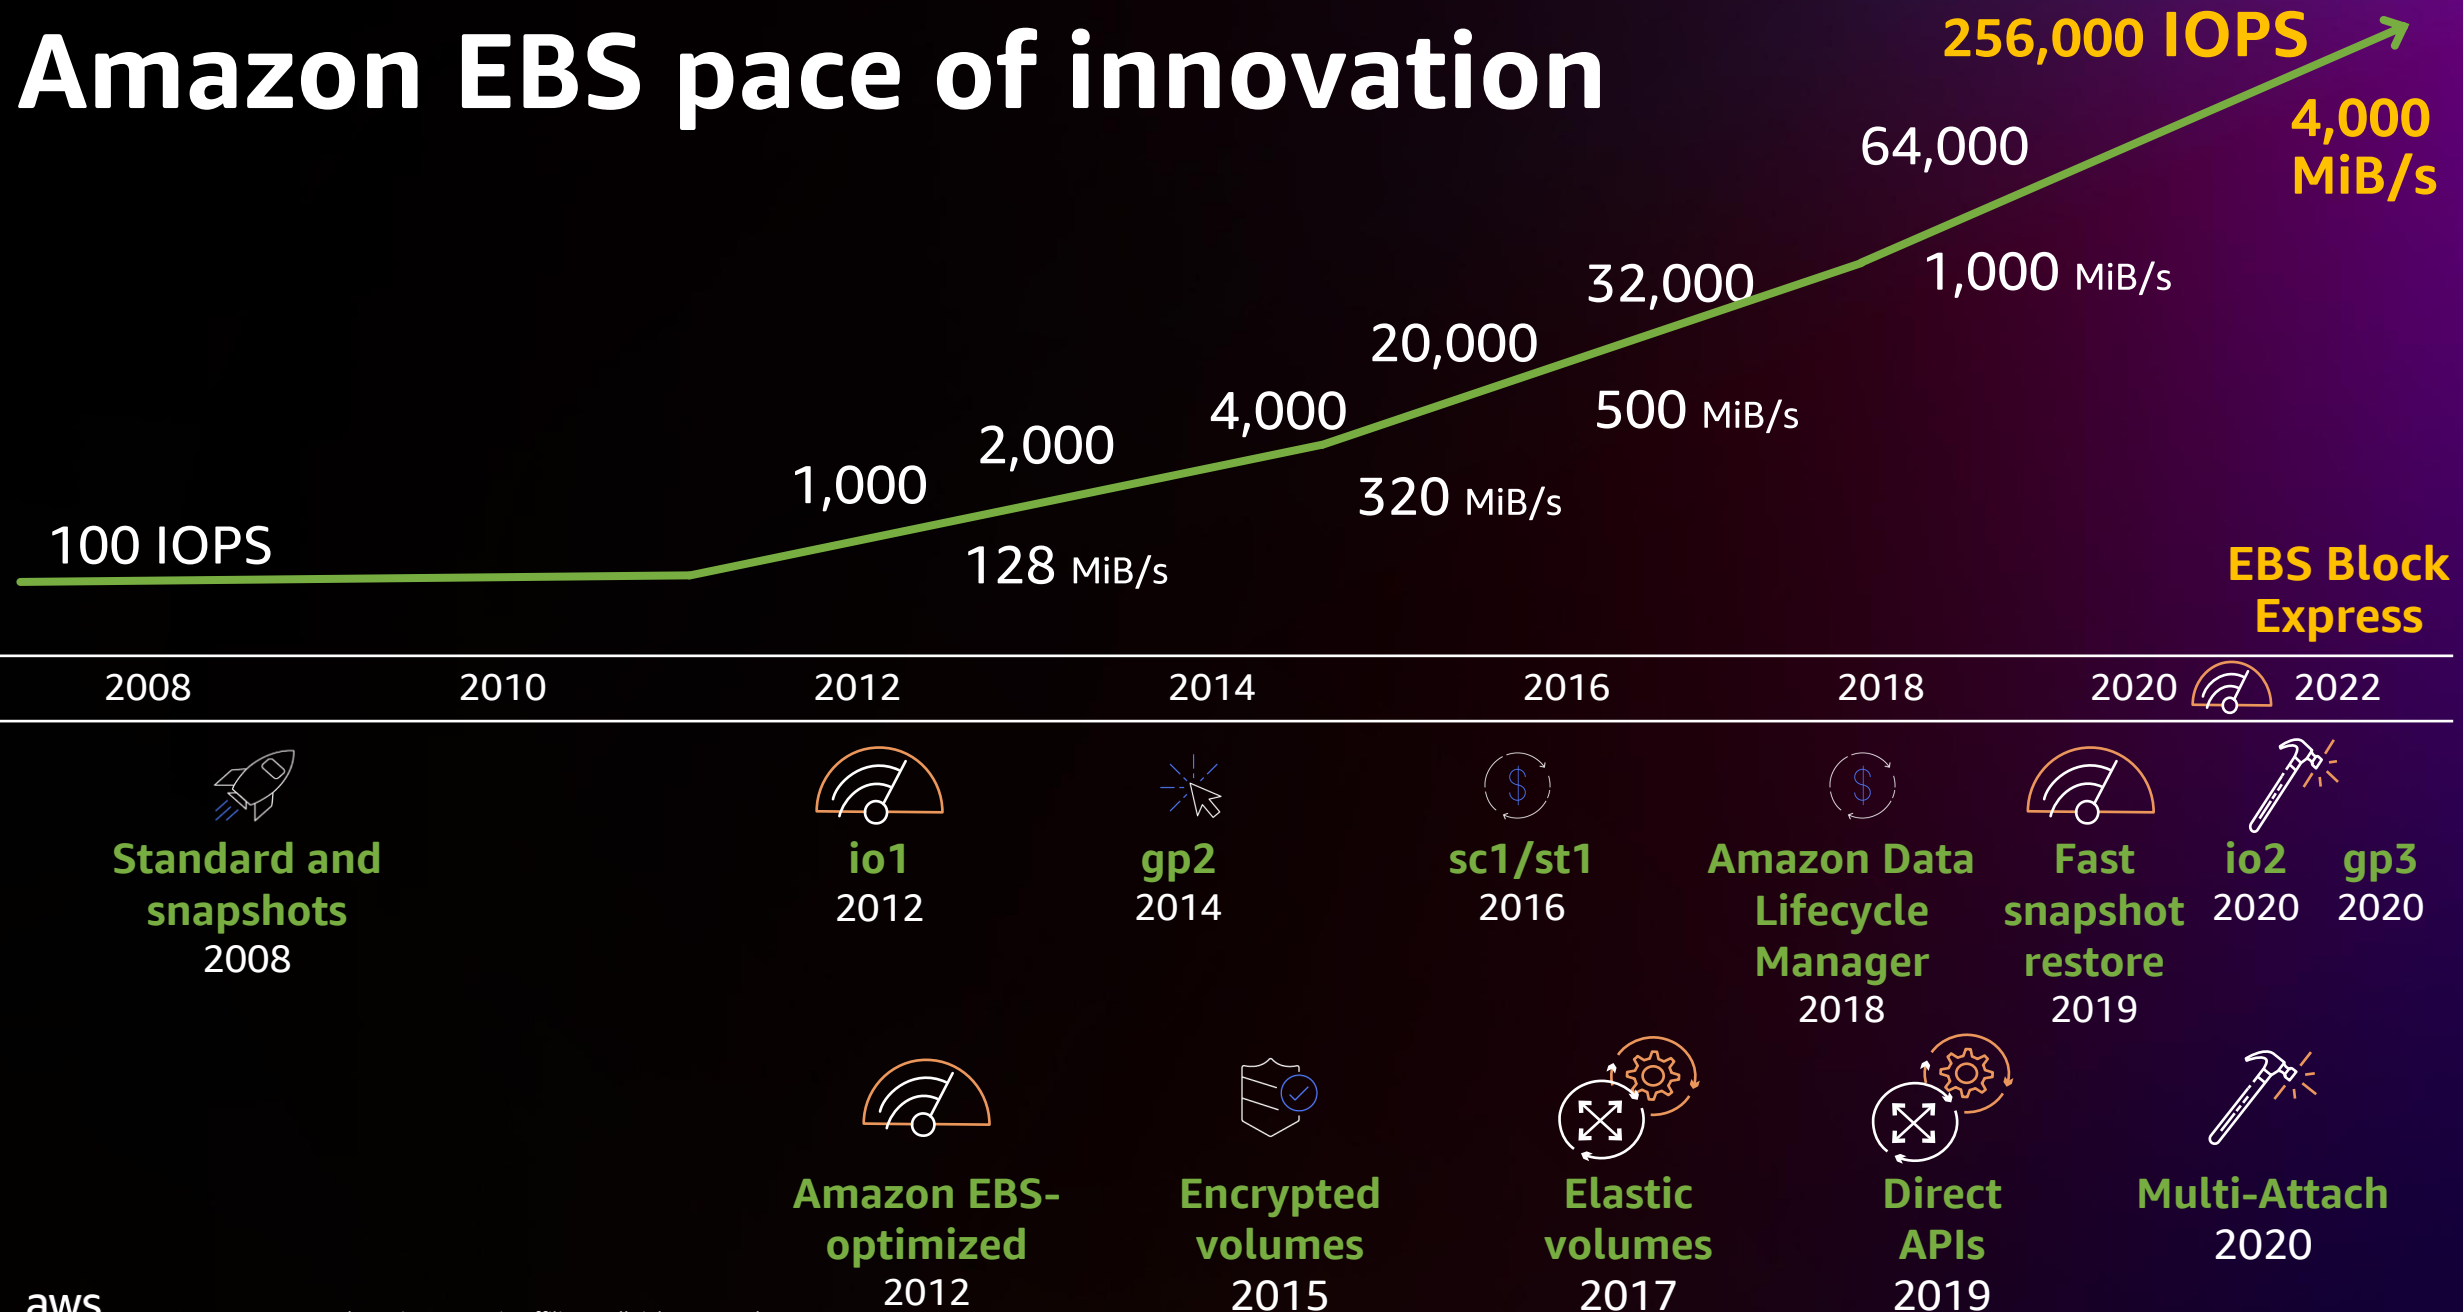
Encrypt all volumes using default encryption keys or custom managed keys



Cost-effective


Pay as low as \$0.025/GB-month for HDD and \$0.08/GB-month for SSD

Amazon EBS pace of innovation




Amazon EBS volume types

SSD




gp3
General
Purpose SSD



io2
Block
Express

HDD

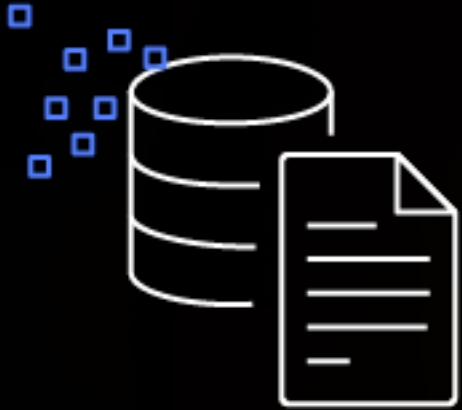


st1
Throughput
Optimized HDD



sc1
Cold
HDD

Understand your mission



Databases

PostgreSQL, MySQL,
Cassandra, MongoDB



Data and analytics

Kafka, Splunk, Hadoop,
data warehousing



Media

Transcoding, encoding,
render farms




File

CIFS, NFS
archive

Select the right volumes

- Multiple Amazon EBS volumes can be attached to an instance
- Mix and match volume types for optimum performance
- Use Elastic Volumes to change your volume characteristics

Workloads in container orchestrators



Simplified persistent storage management

Provision capacity that you
need now **with ability**
to grow on demand

Seamless integration of Amazon EBS with CSI driver

Amazon EBS CSI driver



EKS > Clusters > test-cluster

test-cluster

[Refresh](#) [Delete cluster](#)

[New versions are available for 2 add-ons.](#)

Cluster info [Info](#)

Kubernetes version Info 1.23	Status ✔ Active	Provider EKS
-------------------------------------------------	--------------------	-----------------

Overview | Resources | Compute | Networking | **Add-ons** | Authentication | Logging | Update history | Tags

Add-ons (3) [Info](#) [Edit](#) [Remove](#) [Add new](#)

[1](#)

aws-ebs-csi-driver <input type="radio"/>	kube-proxy <input type="radio"/>	vpc-cni <input type="radio"/>
Type Storage	Type Networking	Type Networking
Status ⌚ Creating	Status ✔ Active	Status ✔ Active
Version v1.11.4-eksbuild.1	Version v1.23.7-eksbuild.1	Version v1.10.4-eksbuild.1

Workloads on Kubernetes clusters



```
~/go/s/g/to/k8s.io main kubectl apply -k "github.com/kubernetes-sigs/aws-ebs-csi-driver/deploy/kubernetes/overlays/stable/?ref=release-1.12"
serviceaccount/ebs-csi-controller-sa created
serviceaccount/ebs-csi-node-sa created
clusterrole.rbac.authorization.k8s.io/ebs-csi-node-role created
clusterrole.rbac.authorization.k8s.io/ebs-external-attacher-role created
clusterrole.rbac.authorization.k8s.io/ebs-external-provisioner-role created
clusterrole.rbac.authorization.k8s.io/ebs-external-resizer-role created
clusterrole.rbac.authorization.k8s.io/ebs-external-snapshotter-role created
clusterrolebinding.rbac.authorization.k8s.io/ebs-csi-attacher-binding created
clusterrolebinding.rbac.authorization.k8s.io/ebs-csi-node-getter-binding created
clusterrolebinding.rbac.authorization.k8s.io/ebs-csi-provisioner-binding created
clusterrolebinding.rbac.authorization.k8s.io/ebs-csi-resizer-binding created
clusterrolebinding.rbac.authorization.k8s.io/ebs-csi-snapshotter-binding created
deployment.apps/ebs-csi-controller created
poddisruptionbudget.policy/ebs-csi-controller created
daemonset.apps/ebs-csi-node created
csidriver.storage.k8s.io/ebs.csi.aws.com created
```

```
~/go/s/g/to/aws-ebs-csi-driver master helm upgrade --install aws-ebs-csi-driver --namespace kube-system ./charts/aws-ebs-csi-driver --values ./charts/aws-ebs-csi-driver/values.yaml
Release "aws-ebs-csi-driver" does not exist. Installing it now.
```

Overcome on-premises challenges



Scalable

Virtually **unlimited** scalability

Provision capacity that you need now with **ability to grow on demand**



Simplify management

Simplified provisioning of resources

Migrate data between storage tiers **seamlessly** to address changing workload requirements

Built-in security and recovery features that can be configured with a few clicks



Optimize operations

Eliminate investments in infrastructure management and complex data center operations

Eliminate the infrastructure lifecycle tax

No more disruptions and complex migration projects from EOL infrastructure

Seamless integration into AWS services



Security




Built-in cloud security

Robust encryption



Protect data at-rest, data in-transit, and all volume backups with Amazon EBS encryption

Encryption by default



A single opt-in setting enables you to easily achieve compliance goals by encrypting all new volumes as they are created

Simple and easy protection



Secure your data with ease and with zero impact on performance; encryption is supported by all volume types and it includes built-in key management

SECURE FOR MISSION-CRITICAL WORKLOADS

Amazon EBS volume encryption



EBS encryption:
data volumes

Volumes > Create Volume

Create Volume

Volume Type ⓘ

Size (GiB) (Min: 1 GiB, Max: 16384 GiB) ⓘ

IOPS 300 / 3000 (Baseline of 3 IOPS per GiB with a minimum of 100 IOPS, burstable to 3000 IOPS) ⓘ

Availability Zone* ⓘ

Throughput (MB/s) Not applicable ⓘ

Snapshot ID ⓘ

Encryption Encrypt this volume ⓘ

Master Key ⓘ

KMS Key Description Default master key that protects my EBS volumes when no other key is defined

KMS Key Account This account: ██████████

KMS Key ID bae63fd5-6121-411b-84d1-fd93f8d8b6a4

KMS Key ARN arn:aws:kms:us-east-1:██████████:key/bae63fd5-6121-411b-84d1-fd93f8d8b6a4

Key	Value
Name	Encrypted Volume
Application Name	App01

Add Tag 48 remaining (Up to 50 tags maximum)

[Cancel](#) [Create Volume](#)

EBS is designed for...

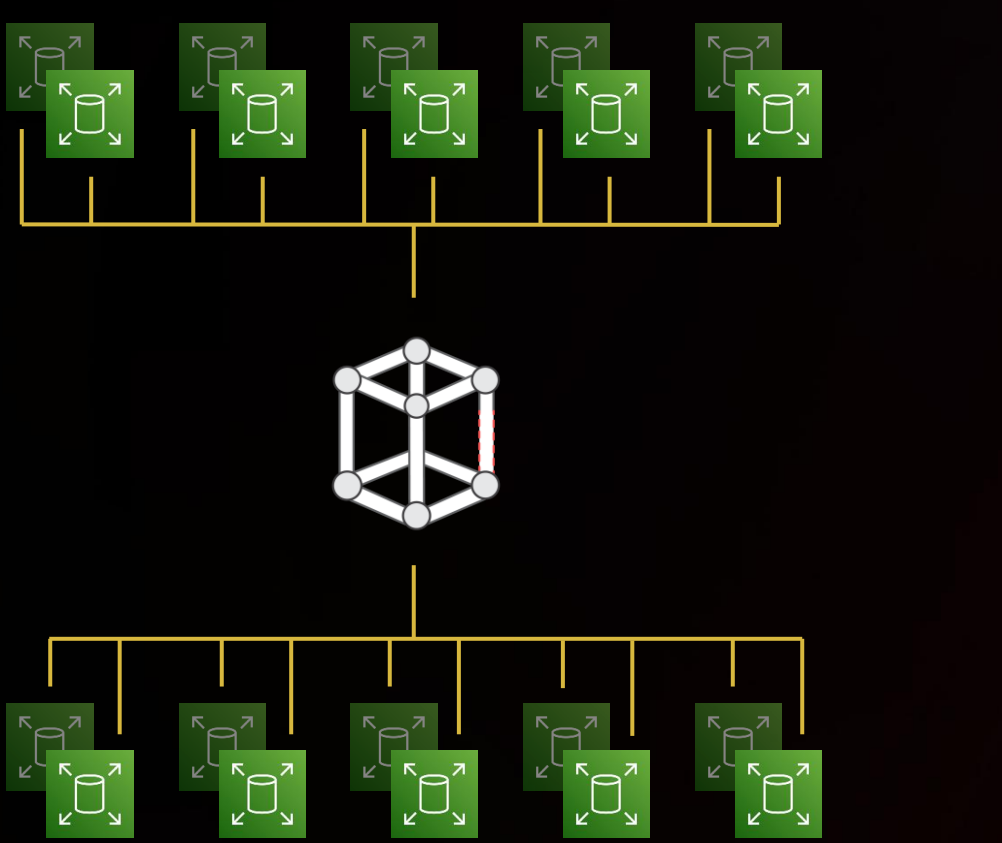


99.999% service availability



Up to 99.999% durability

Configuration manager



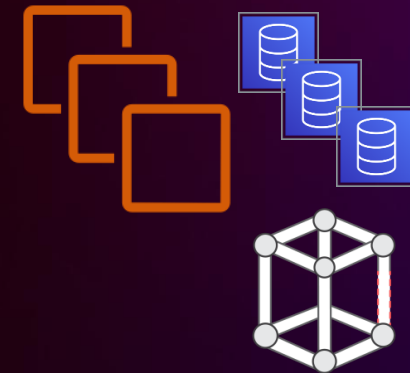

- Authority map for host and volume
- Wholly contained in Availability Zone
- Standalone service
- Paxos-based replication

Impact reduction: AZ independence

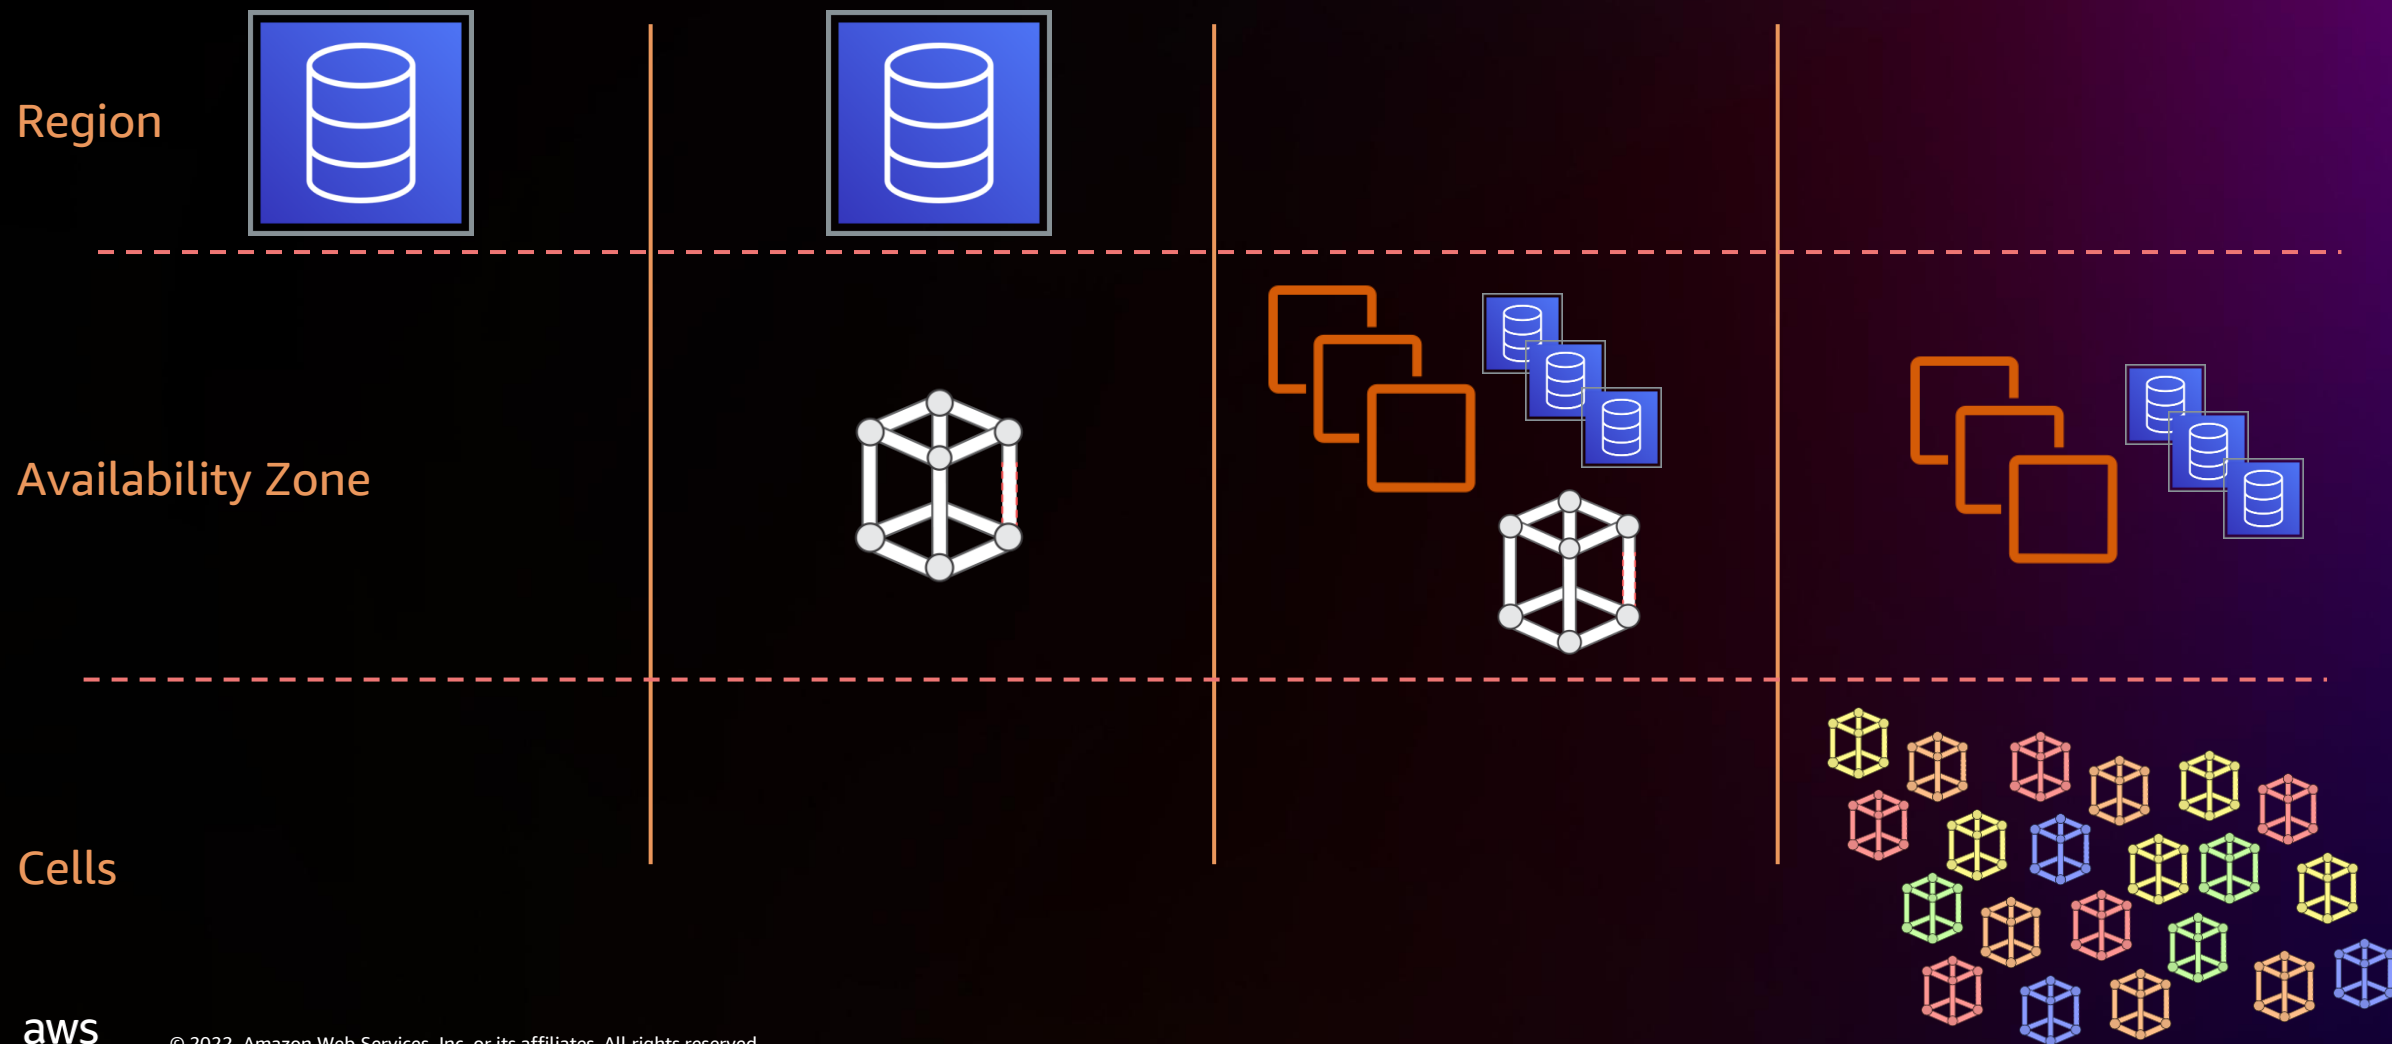
Region




Availability Zone



Impact reduction: Physalia



Physalia: Millions of tiny databases



- Small number of clients per database
- Simple to manage
- Cell can reside near clients
- Reduced area of impact
- Each cell is a 7-node Paxos cluster

SRD storage fabric



- Native multipathing enabled
- Congestion control for large-scale cloud deployments
- High throughput; minimizes latency jitter
- Fast and transparent recovery from network failures

Scalable Reliable Datagram (SRD) is a high-performance, low-latency network transport protocol

Nitro card for Amazon EBS



NVMe controller

Standard drivers broadly available


Amazon EBS data plane

Encryption support

NVM to remote storage protocol


Amazon EBS-optimized by default

New generation architecture



- Disaggregated infrastructure
- Disaggregated data plane and data services
- Nitro and SRD
- More consistent I/O latency

Instance family: X2iedn



- Amazon EBS-optimized by default
- 32:1 ratio of memory to vCPU on all sizes
- Up to 260,000 IOPS
- Up to 100 Gbps network bandwidth
- Up to 80 Gbps of bandwidth to Amazon EBS

Disaggregated data plane

- Reducing code in the fast path improves performance
Think critically about what absolutely needs to be there!
- With smaller code, it is easier to reason about security and correctness
- Can move faster with defined APIs

Protecting volumes with Amazon EBS snapshots

What are EBS snapshots?



Snapshot use cases

Point-in-time copy of EBS volume

Incremental copies
Only changed blocks are copied

Stored in Amazon S3
11 9s of durability

Crash consistent

Contains **all data necessary** to recover a volume

Back up data on EBS volumes

Meet **RPO and RTO** SLA objectives

Copy volumes within or across Availability Zones for **redundancy**

Copy volumes to another Region for **disaster recovery**

Redeploy production data for test/dev

Amazon EBS snapshots features



Track

Tag snapshots at creation

Track and monitor
snapshot resources

Enforce IAM rules based on tagging

Cost allocation tags provide visibility
into snapshot storage costs




Secure

IAM policy-driven privileges for
snapshot management

Encrypted snapshots

Snapshots of encrypted volumes
are encrypted at creation

Unencrypted snapshots can be
encrypted during copy process



Share

Allow other accounts to copy a
snapshot or create a volume
from a snapshot

Enable test/dev activities

Enable redeployment of
custom AMIs

Amazon EBS snapshots archive tier



- Up to 75% lower storage costs compared to standard tier
- Designed for long-term retention of snapshots
- Snapshots need to be retrieved from archive before use

Amazon EBS snapshots archive tier provides low storage costs for long-term retention of rarely accessed EBS snapshots

Recycle Bin for EBS snapshots



Quick recovery of deleted snapshots

Automatically retain deleted snapshots for a **retention period** you specify

Use **retention rules** to specify retention periods for **all or some tagged snapshots** in your account

Recover deleted snapshots before retention time expiry


Restrict access to Recycle Bin rules and resources using **administrator IAM role**

Previously...



With IAM policies, you can prevent unencrypted volumes from launching

Amazon EBS encryption by default



Without change to workflows, newly launched volumes, plus snapshots are encrypted

Amazon EBS key takeaways

Continue to **lower** your costs



Innovative performance and reliability



Flexible tools to scale and manage



Simple to use and manage **securely**



If you remember nothing else...



Monitor access

Keep tabs on which snapshots you are sharing and why



Account-level encryption

Encryption is a checkbox – literally! Use it


Amazon EBS at re:Invent

Date	Time	ID	Session Title
Tuesday, Nov. 29	4:15 PM	STG404	Reduce time to recovery with Amazon EBS direct APIs & flexible snapshot proxy
Wednesday, Nov. 30	8:30 AM	STG314	Amazon EBS snapshots: Build protection and cost-optimize
Wednesday, Nov. 30	9:15 AM	STG220	Best practices for moving on-premises applications to Amazon EBS
Wednesday, Nov. 30	12:15 PM	STG406	Automated data protection & backup compliance with Amazon EBS snapshots
Wednesday, Nov. 30	1:45 PM	STG219	Build resilient architectures with Amazon EBS
Wednesday, Nov. 30	5:30 PM	STG306	Best practices for EBS volumes & performance monitoring using CloudWatch




Continue your AWS storage learning

Build a learning plan



Set your AWS storage learning plans via **AWS Skill Builder**

Increase your knowledge



Use our **ramp-up guides** to build your storage knowledge

Earn AWS storage badges



Demonstrate your knowledge by achieving **digital badges**

aws.training/storage



Thank you!



Please complete the session survey in the **mobile app**

