

The background features a dark, almost black, field with several large, overlapping, semi-transparent shapes in shades of purple, magenta, and blue. These shapes are curved and layered, creating a sense of depth and movement. Two thin, light-colored lines cross the scene diagonally, adding to the geometric complexity.

AWS re:Invent

DECEMBER 2 - 6, 2024 | LAS VEGAS, NV

DAT426 - NEW

Improve resiliency using Amazon MemoryDB Multi-Region

Itay Maoz

(he/him)

General Manager, In-Memory Databases
AWS

Jean Guyader

(he/him)

Sr. Engineering Manager
AWS



Customers need multi-Region applications

- 34 regions
- 108 availability zones (AZ)
- Every AWS Region is built for high resilience
- Each AZ is physically separated from other AZs



Why build multi-Region applications ?



Increase resilience

- Business continuity
- Regulation/compliance



Global applications

- Low regional local latencies

Why build multi-Region applications ?



Increase resilience

- Business continuity
- Regulation/compliance

Critical applications:

- Financial industry
- Health care
- Large enterprises

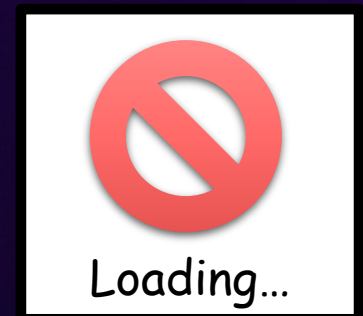
Why build multi-Region applications ?



Global applications

- Low regional local latencies
- Global applications like gaming leaderboards
- Global user data like profiles, history, and preferences

Bad customer experience



Why do you need a multi-Region data layer?



Increase resilience

- Business continuity
- Regulation / compliance

Disaster recovery

- Region failover
- Most applications cannot function without their data
- Even a “standby Region” needs the application’s data



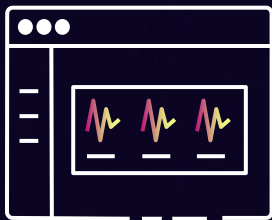
Global applications

- Low regional local latencies

Global data; Local latency

- High latency to fetch data → High application latency
- Local AZ latency: typically < 0.3ms
- Cross-AZ latency: typically < 1ms
- Cross-Region latency: typically < 1s

How to build a multi-Region data layer?



Build your own custom solution

Too hard!



Use a multi-Region database

Which one?

Key considerations for a **multi-Region database**

Consistency & Durability:



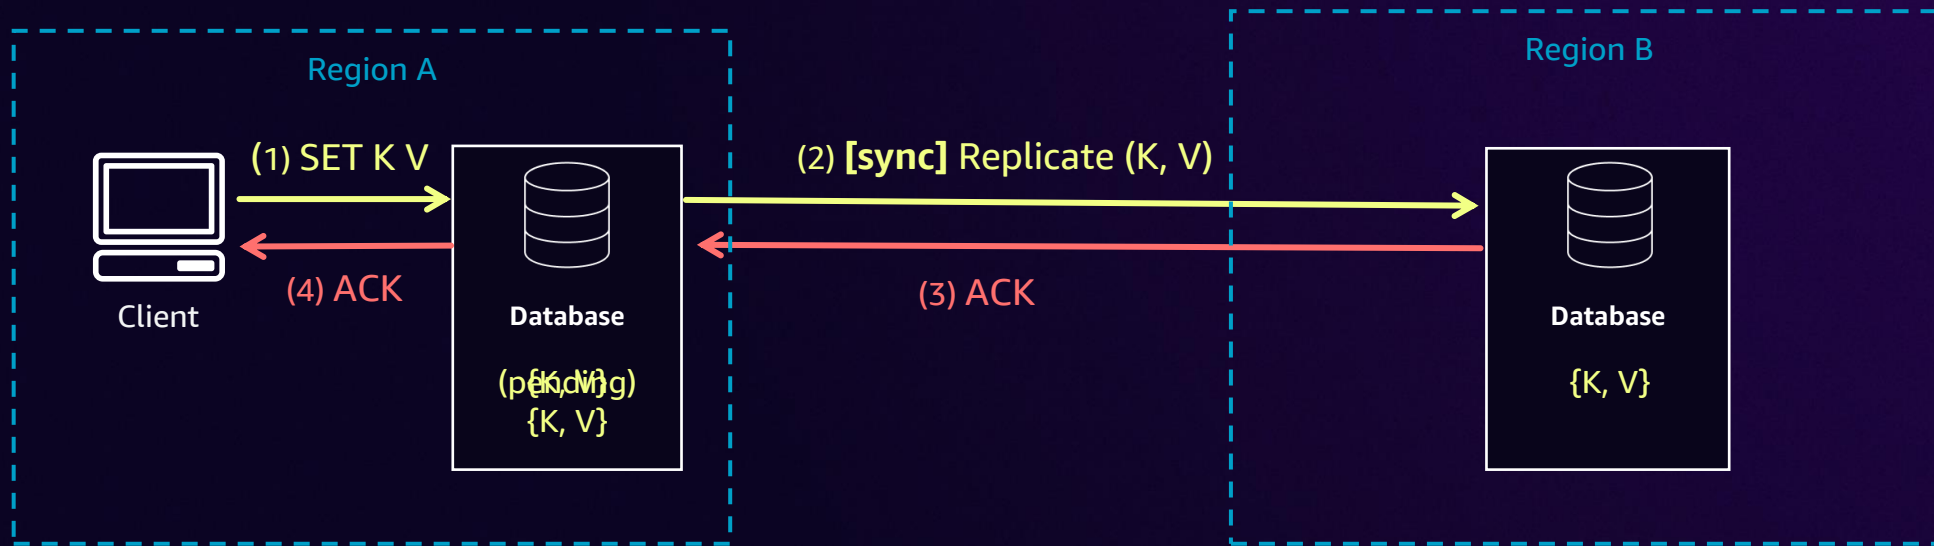
What is your **Recovery Point Objective (RPO)**?

RPO is the maximum data loss you can tolerate if a Region is unavailable

Synchronous or **asynchronous** cross-Region replication?

Key considerations for a multi-Region database

Synchronous Cross-Region Replication

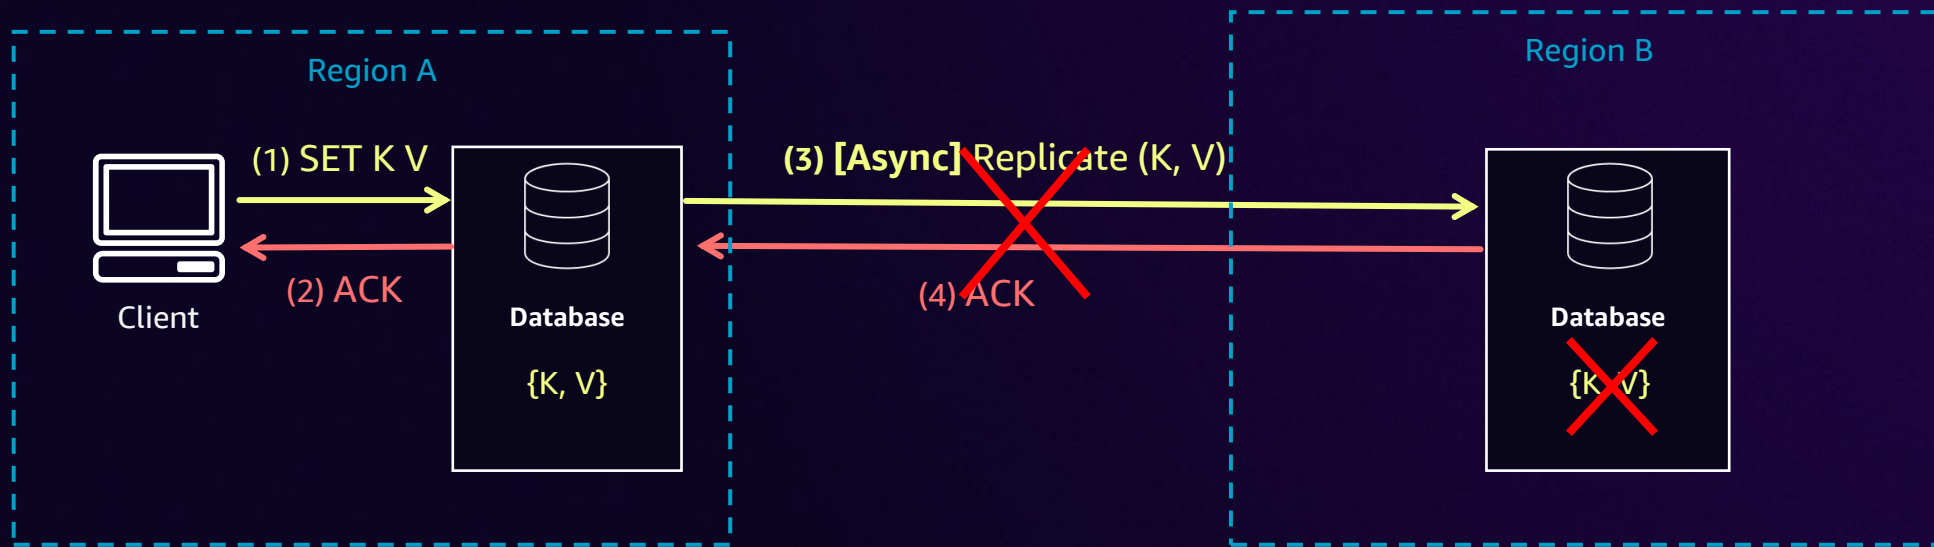


The client receives an ACK only after the data is durable in **both** Regions

Strong consistency

Key considerations for a multi-Region database

Asynchronous Cross-Region Replication



The client receives an ACK after the data is durable in **one** Region

Replication failure → database will retry

Eventual consistency

Key considerations for a **multi-Region database**

Synchronous Replication

+ **RPO = ~ 0 is possible**, the client gets an ack only after the data is in **ALL** Regions

- **High write latency**, up to **hundreds** of milliseconds

Strong consistency

Asynchronous Replication

- **RPO > 0**, the client gets an ack after the data is durable in only **ONE** Region

+ **Low write latency**, a **few** milliseconds

Usually the replication lag is small (seconds) → **RPO** can be seconds

Eventual consistency

Key considerations for a **multi-Region database**

High Availability?

What is your **Recovery Time Objective (RTO)**?

RTO is the maximum time to restore operations following a Region failure

Active/active or **active/passive** cross-Region replication?

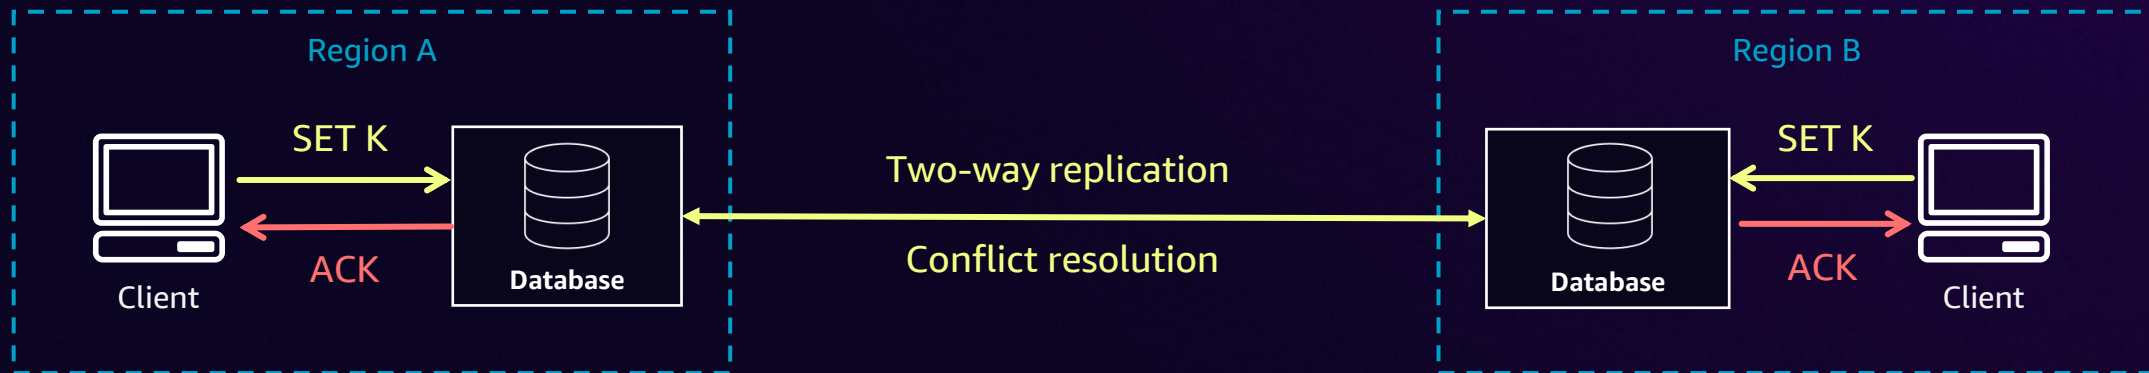


Key considerations for a multi-Region database

Active/Passive Cross-Region Replication, Single-Writer

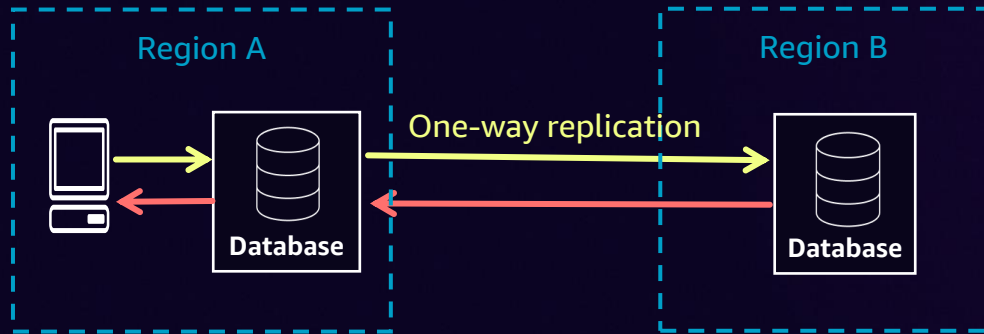


Active/Active Cross-Region Replication, Multi-Writer



Key considerations for a multi-Region database

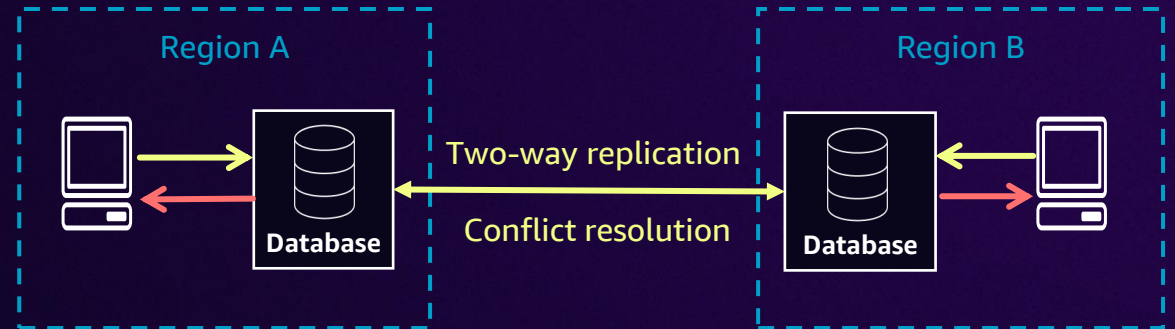
Active/Passive Replication, Single-Writer



Simpler:

- + Easier to build, fewer constraints and limitations
- + No write conflicts

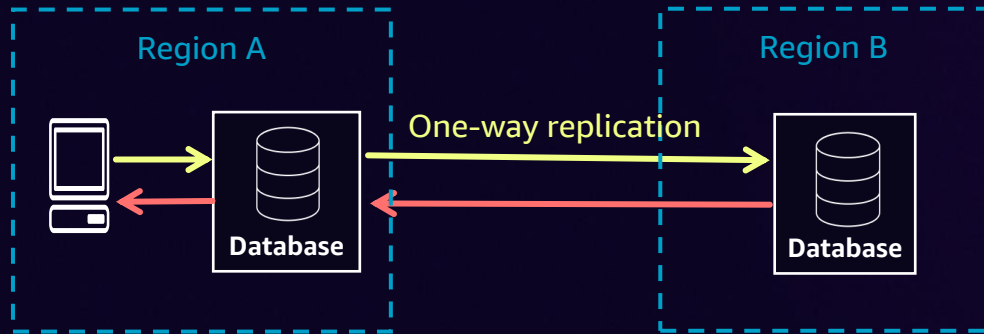
Active/Active Replication, Multi-writer



- Harder to build, more constraints and limitations
- Potential write conflicts

Key considerations for a multi-Region database

Active/Passive Replication, Single-Writer

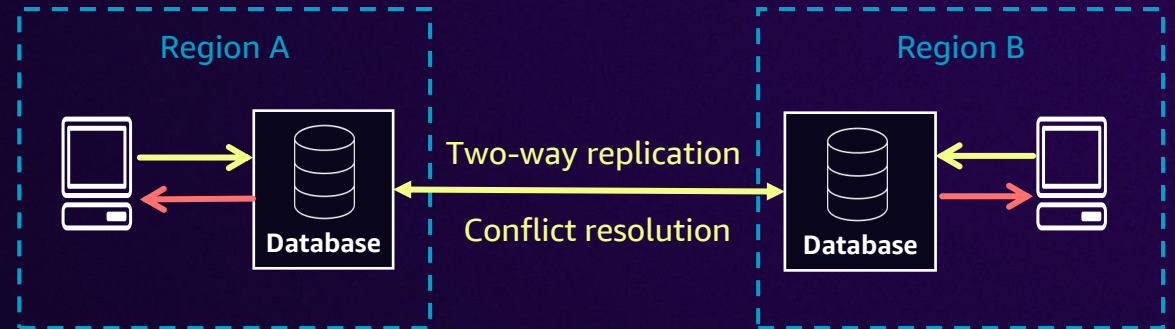


Simpler:

- + Easier to build, fewer constraints and limitations
- + No write conflicts

- Complex to manage outages – requires application side failure detection, complex failover, and disaster recovery
- RTO > 0, minutes or even hours

Active/Active Replication, Multi-writer



- Harder to build, more constraints and limitations
- Potential write conflicts

Better Availability:

- + Easy to manage Region failure – no need application side database failover
- + RTO ≈ 0 is possible

Key considerations for a **multi-Region database**

- Number of regions supported in a single database
 - For **disaster recovery**, two Regions are sufficient for many use cases
 - **Global applications** with global data usually require the data in all the application's Regions



Other considerations for a multi-Region database



Enterprise security, compliance



The data layer API and simplicity of the solution



Scalability



Performance

Most complete set of purpose-built databases

RELATIONAL



Amazon RDS



Amazon Aurora

KEY-VALUE



Amazon DynamoDB

DOCUMENT



Amazon DocumentDB

CACHING



Amazon ElastiCache

GRAPH



Amazon Neptune

TIME-SERIES



Amazon Timestream

LEDGER



Amazon QLDB

WIDE COLUMN



Amazon Keyspaces

MEMORY



Amazon MemoryDB

Overview of Amazon MemoryDB



Amazon MemoryDB

VALKEY- AND REDIS OSS- COMPATIBLE, IN-MEMORY DATABASE WITH MULTI-AZ DURABILITY



Valkey & Redis OSS compatibility

Intuitive open source APIs and flexible data structures



Ultra-fast performance

Microsecond read and single-digit millisecond write latencies with millions of RPS



Durability and high availability

Multi-AZ transaction log for durability, replicas for high availability with 99.99% SLA



Fully managed

AWS-managed hardware and software setup, configuration, monitoring, and snapshots



High scalability

Up to 500 nodes and 105 TB of in-memory storage per cluster, with 1 replica per shard



Security

Amazon VPC, encryption at rest and in transit, Access Control Lists (ACLs), IAM auth

Valkey: A community replacement for Redis

Open source, high-performance, key-value data store



Open source

- ▶ Built by **existing** Redis OSS contributors
- ▶ **Drop-in** replacement for Redis OSS 7.2
- ▶ Stewarded by **Linux Foundation**
- ▶ **Permissively** licensed (BSD)



Momentum

- ▶ **40+** organizations
- ▶ **150+** code contributors
- ▶ **1000s** of contributions
- ▶ **1M+** container pulls

Valkey: GA on ElastiCache and MemoryDB

Lower prices for Valkey engine

ElastiCache

33% lower data storage and ECPUs for serverless
20% lower instance price for node-based

MemoryDB

80% lower price for data written, with a new 10 TB/month free tier
30% lower instance price for node-based



Zero downtime upgrade from Redis OSS with a few clicks



Switch easily from Redis OSS reserved nodes to Valkey

Amazon MemoryDB

VALKEY- AND REDIS OSS- COMPATIBLE, IN-MEMORY DATABASE WITH MULTI-AZ DURABILITY



Valkey & Redis OSS compatibility

Intuitive open source APIs and flexible data structures



Ultra-fast performance

Microsecond read and single-digit millisecond write latencies with millions of RPS



Durability and high availability

Multi-AZ transaction log for durability, replicas for high availability with 99.99% SLA



Fully managed

AWS-managed hardware and software setup, configuration, monitoring, and snapshots



High scalability

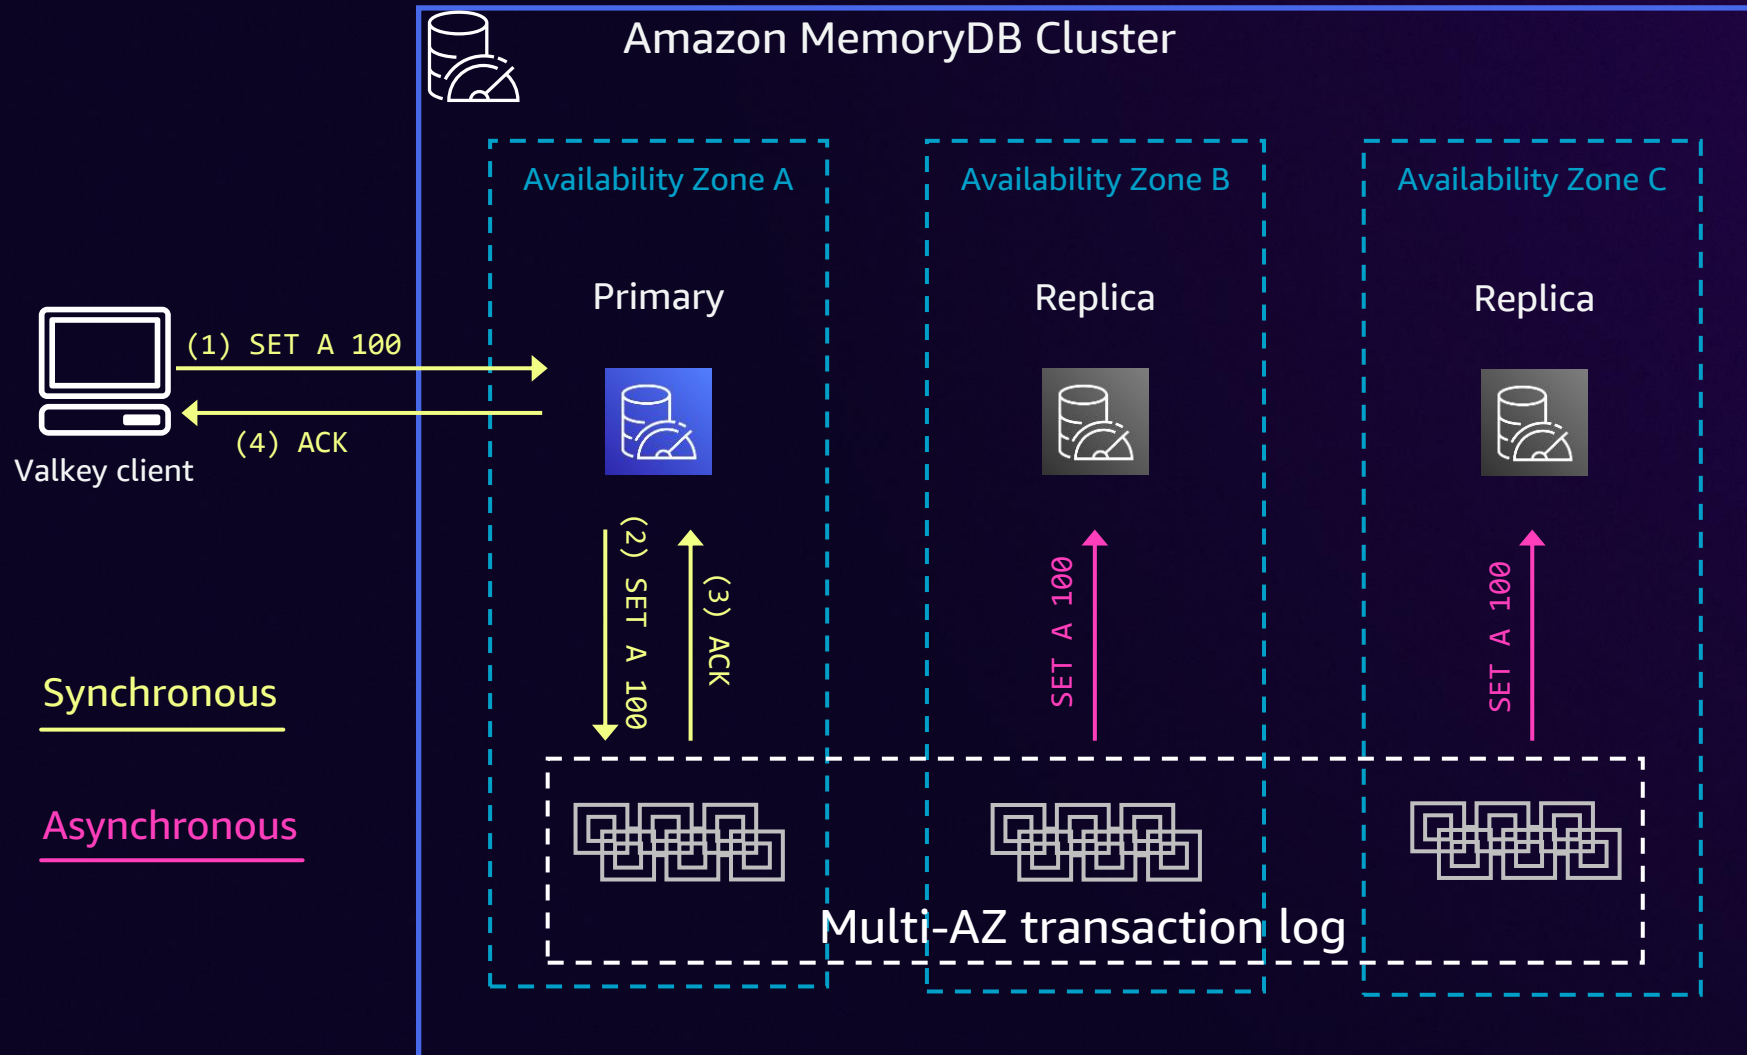
Up to 500 nodes and 105 TB of in-memory storage per cluster, with 1 replica per shard



Security

Amazon VPC, encryption at rest and in transit, Access Control Lists (ACLs), IAM auth

MemoryDB durability



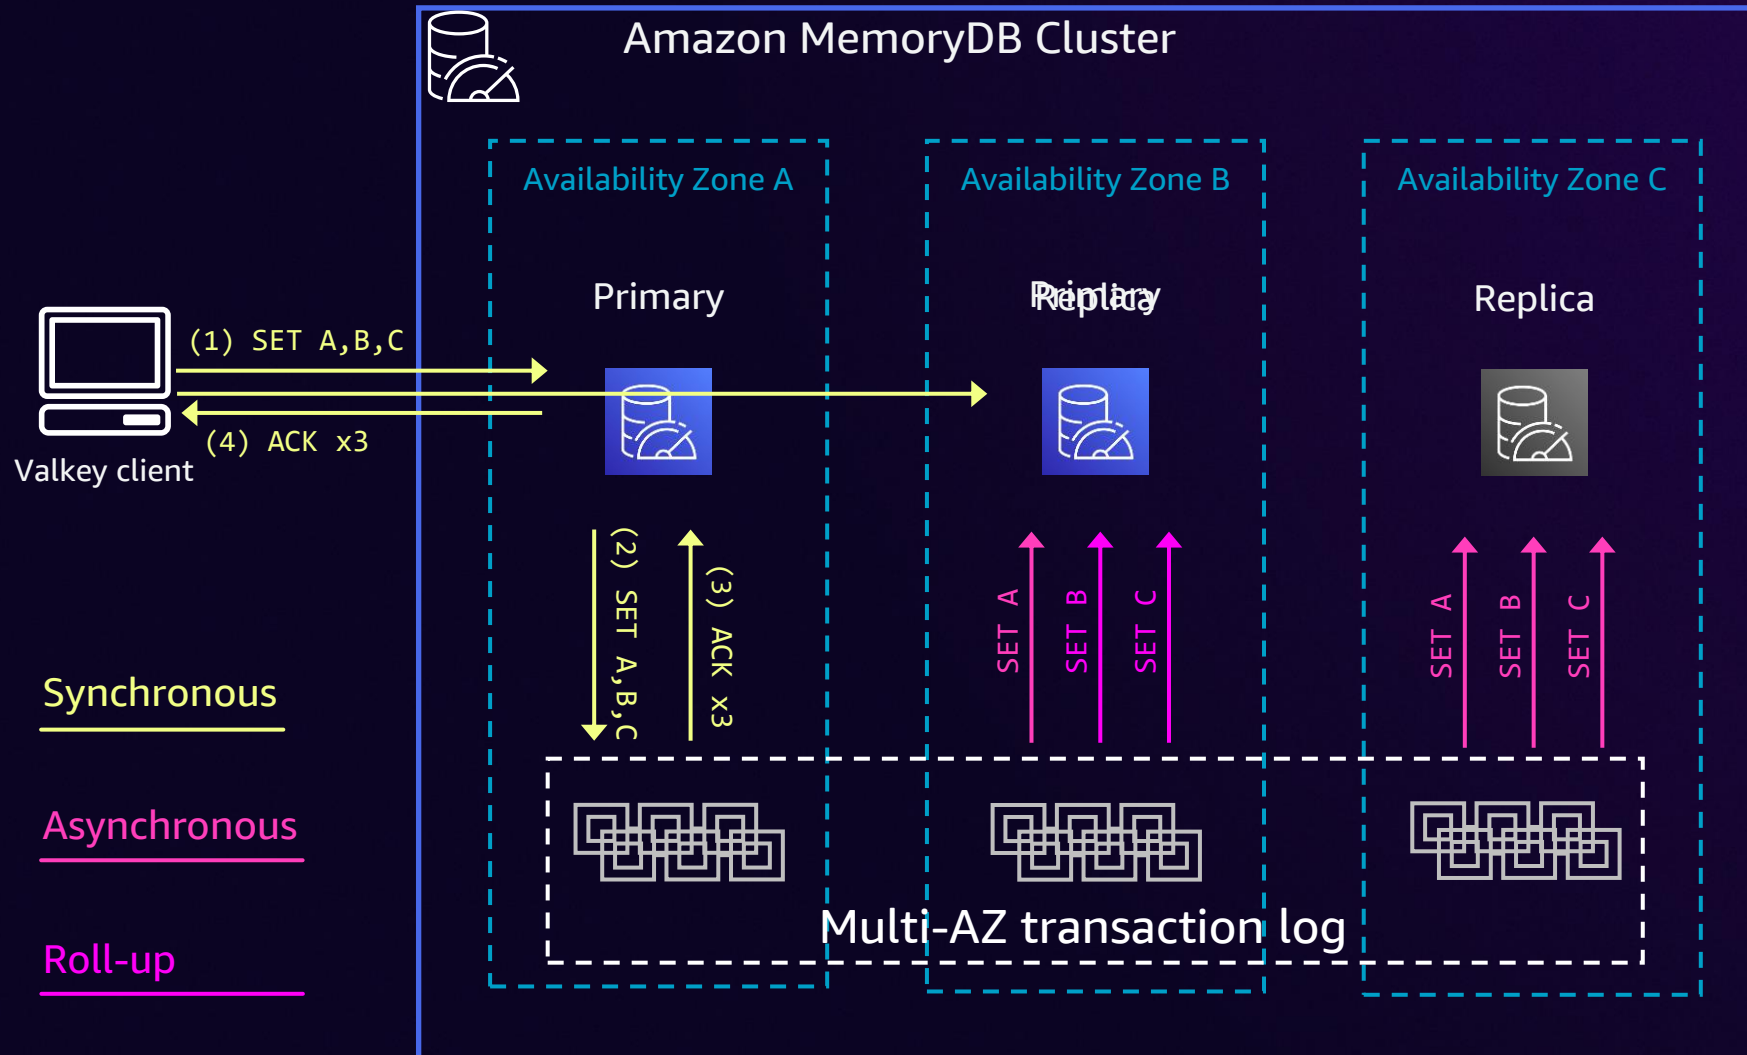
Data is replicated via the transaction log

Acknowledgement sent to client only after transaction is written and replicated across AZs

Data is consumed by replica nodes

Guaranteed delivery of writes to replicas

MemoryDB durability – Failover



Replica reads all transactions from the log, then is promoted to primary

Client is now redirected to newly promoted primary for all writes

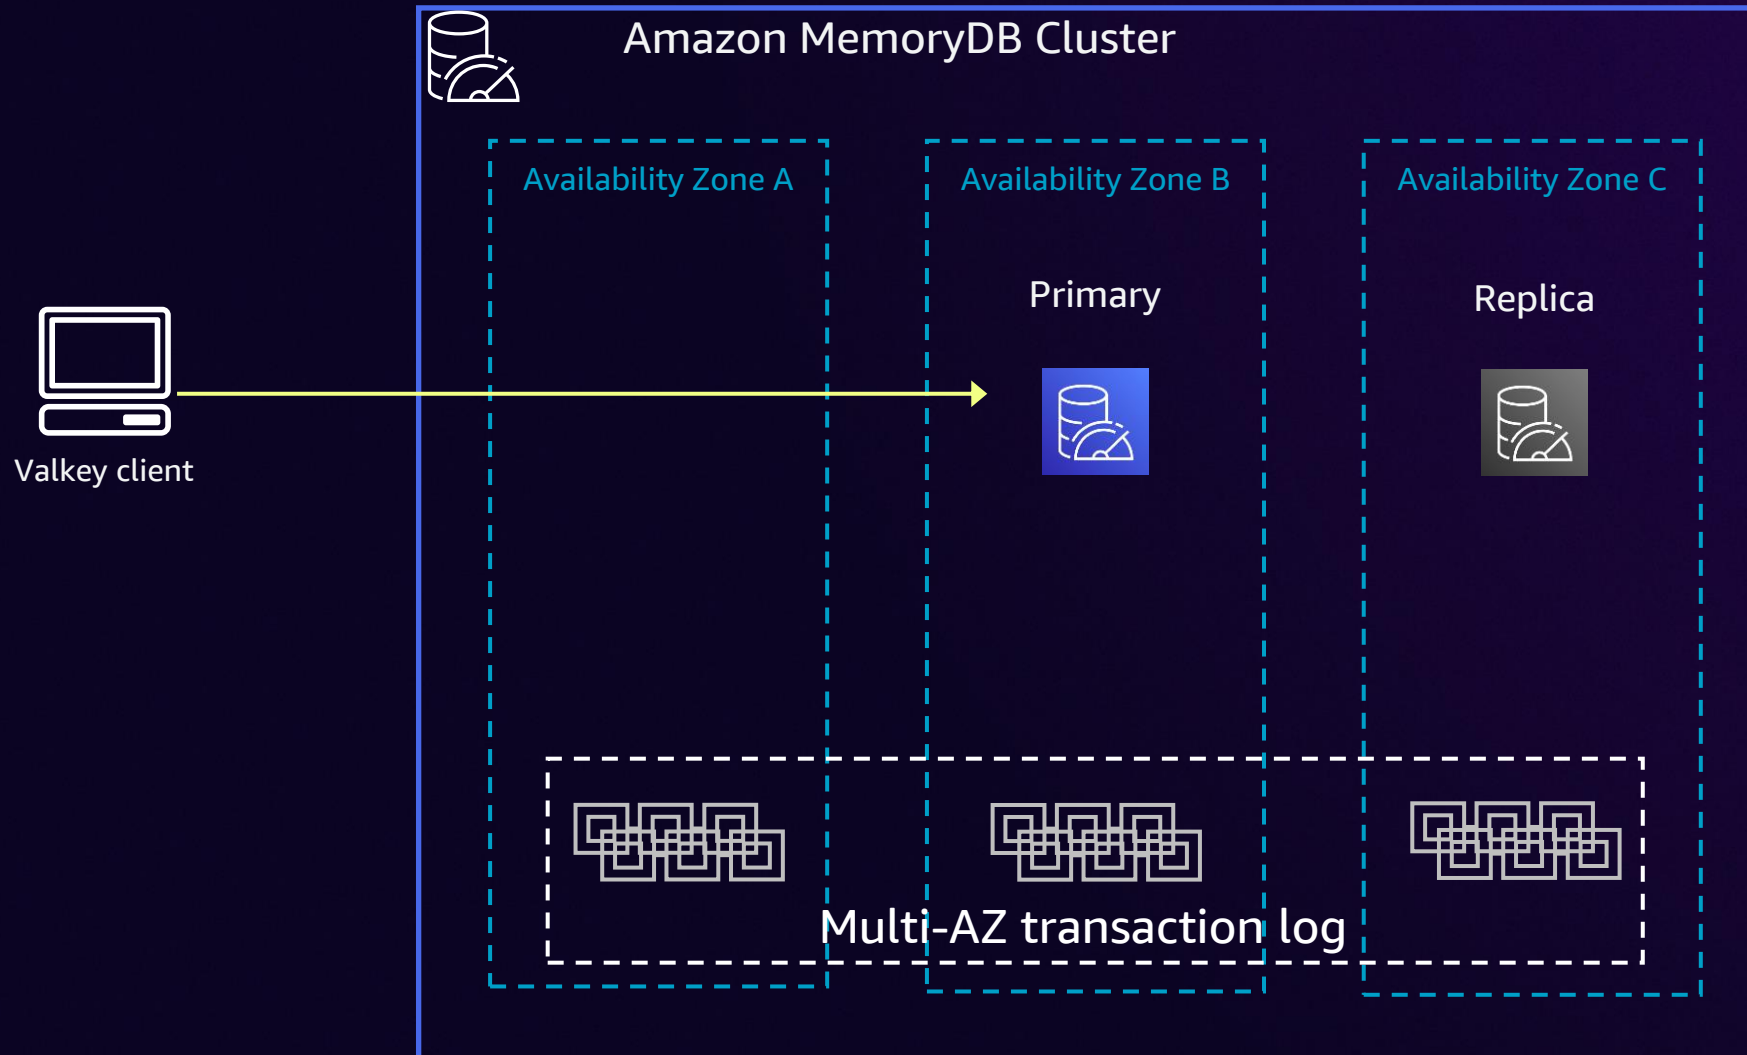
Replicas continue to consume from log asynchronously

Synchronous

Asynchronous

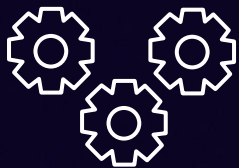
Roll-up

MemoryDB durability – Failover



Fast failover without losing data

Common use cases



Microservices

Performant datastore,
messaging bus



Web and mobile

User data stores,
session management,
geospatial indexing,
chat and message queues



Retail

Customer profiles,
inventory tracking,
fulfillment



Gaming

Player data stores,
session history,
leaderboards



Customer 360

Profiles, history, preferences,
contacts, likes/dislikes,
risk category



Banking and finance

User transactions,
fraud detection



Media and entertainment

User data stores,
real-time streaming



IoT

Streaming device data,
operational insights



Amazon MemoryDB Multi-Region

MemoryDB Multi-Region Active-Active



Design for 11 9s
Durability



μsec reads
< 10ms writes



Multi-Region
Active/Active
Asynchronous replication



99.999% SLA



Encryption in-transit
and at-rest



Scale horizontally
and vertically



Open source rich API:
Valkey / Redis OSS
compatibility



Compliance



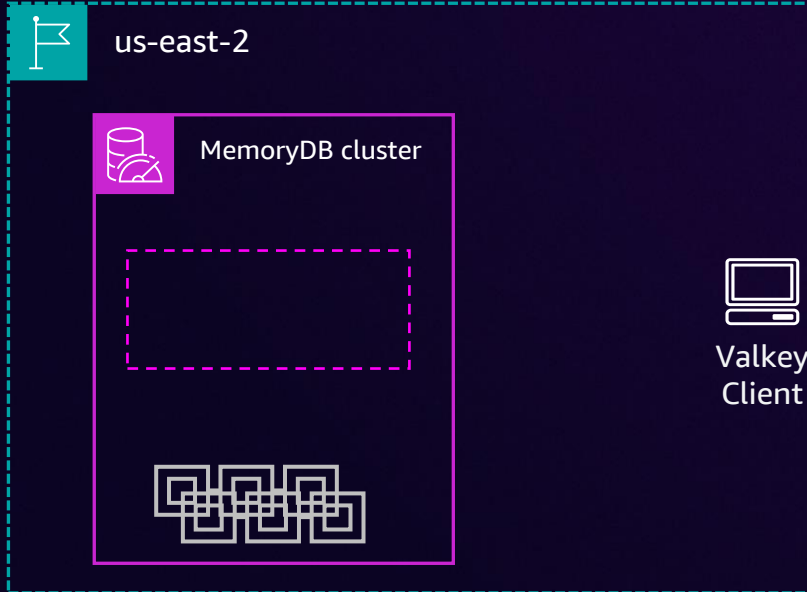
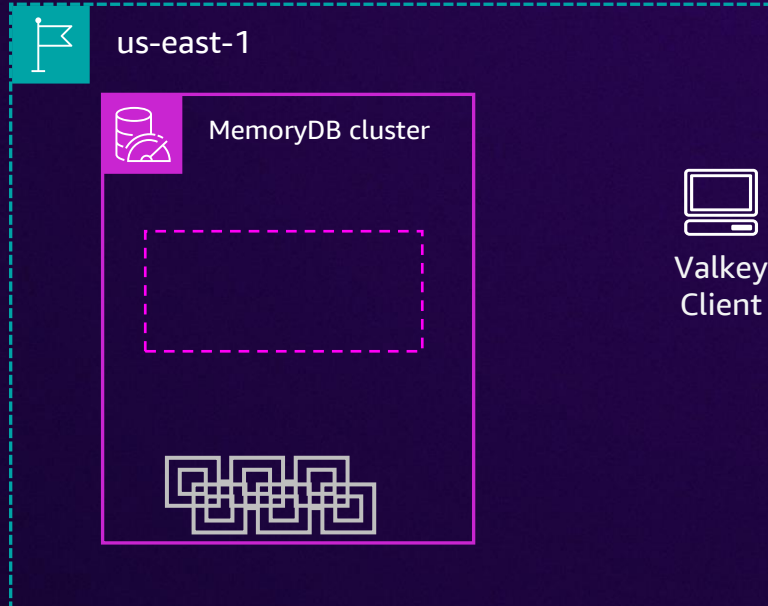
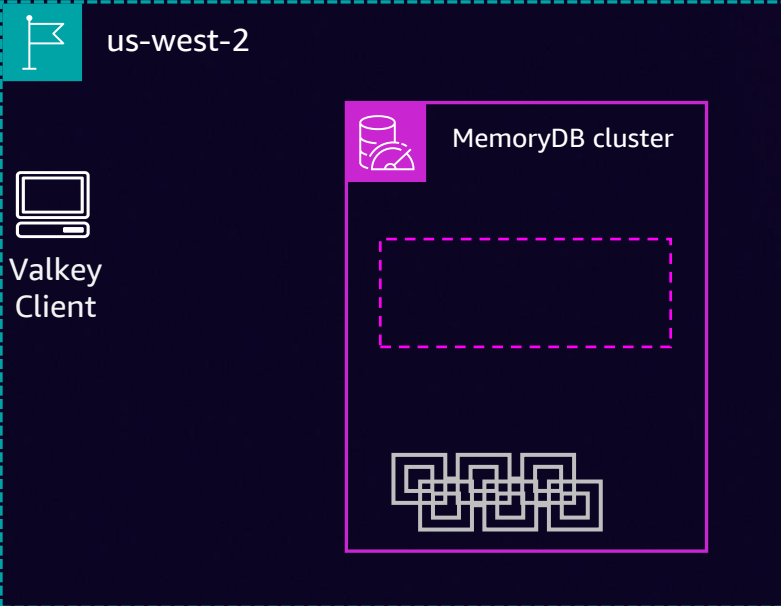
Automatic
backups



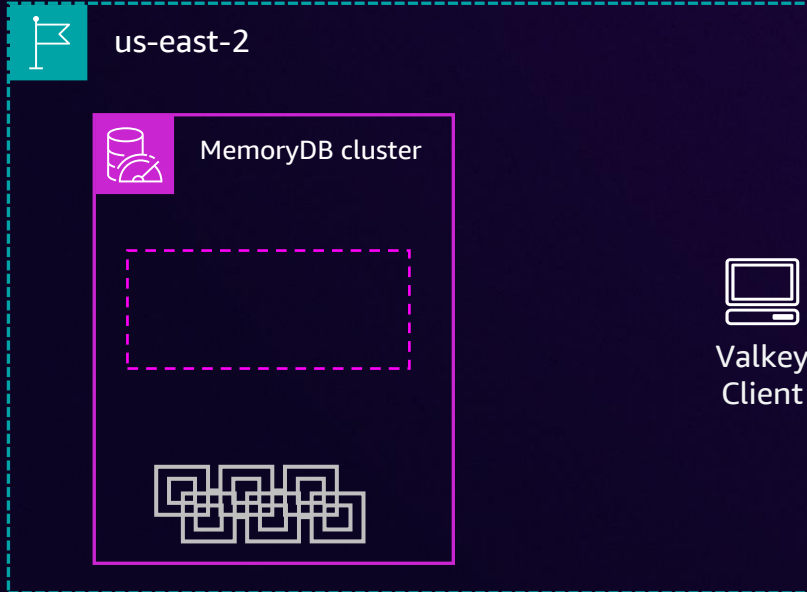
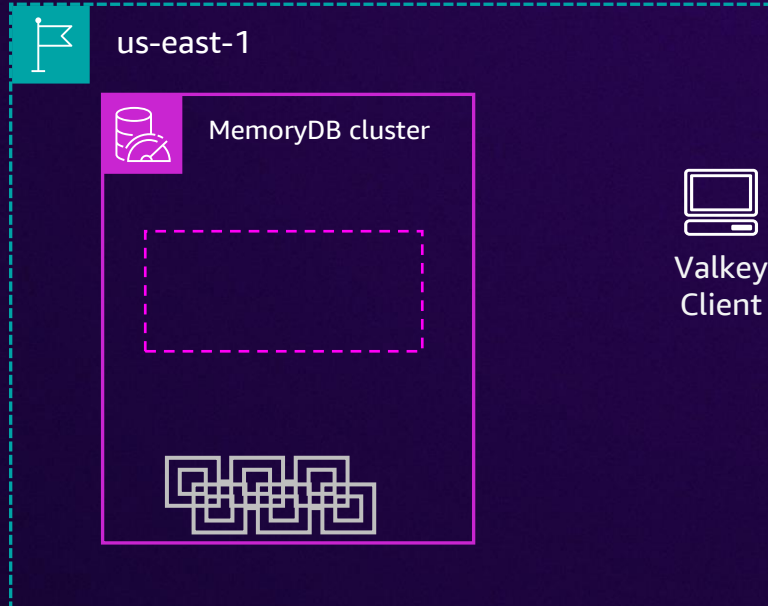
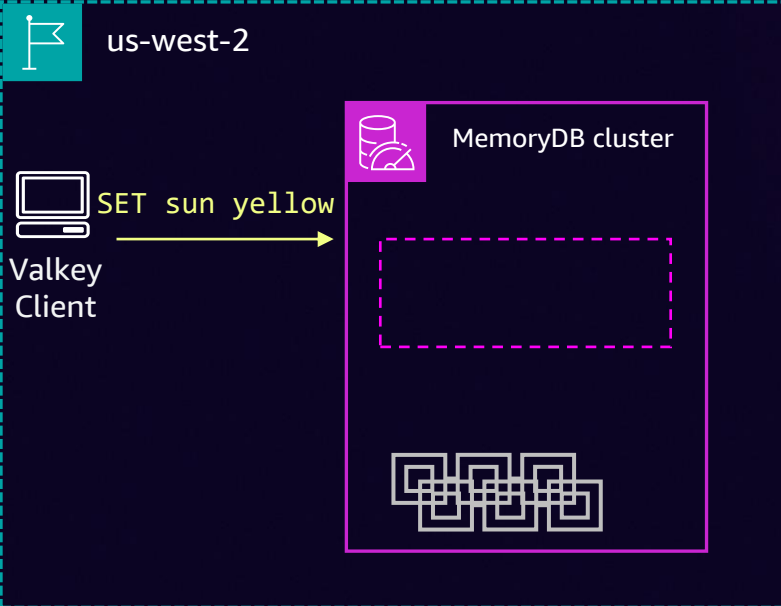
Multi-Region architecture



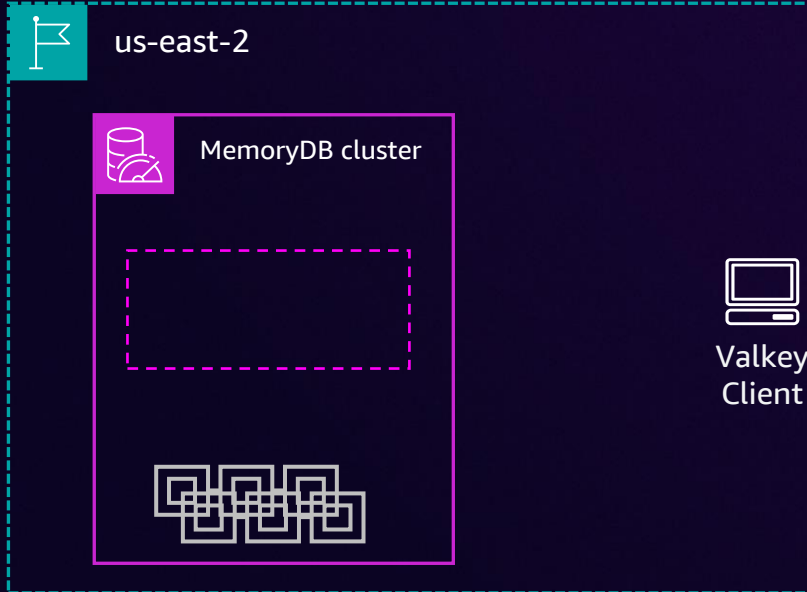
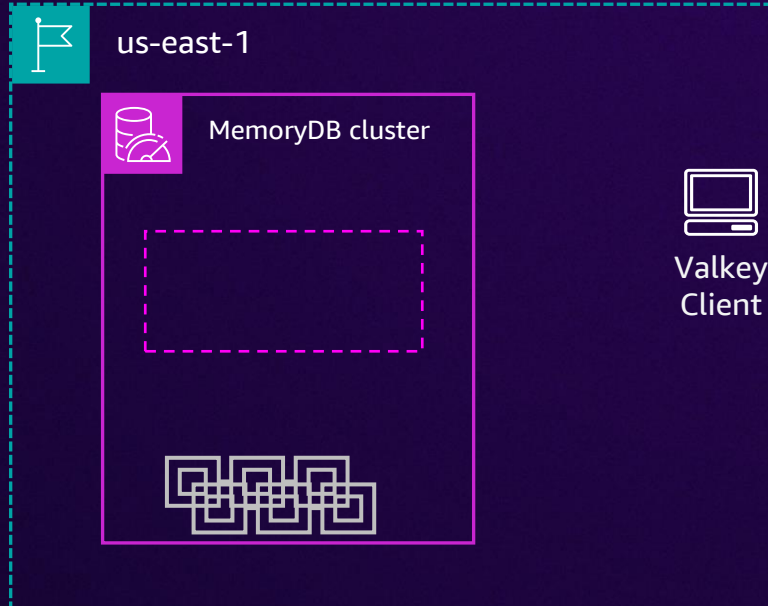
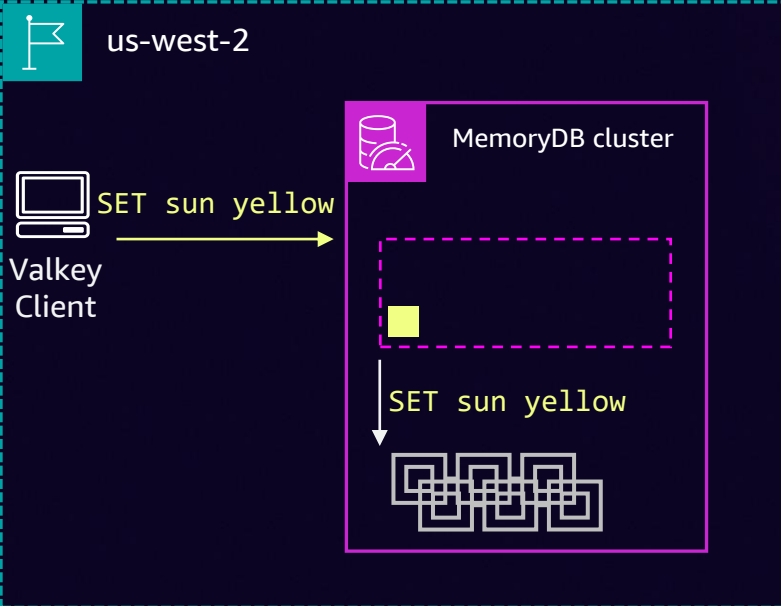
Multi-Region architecture



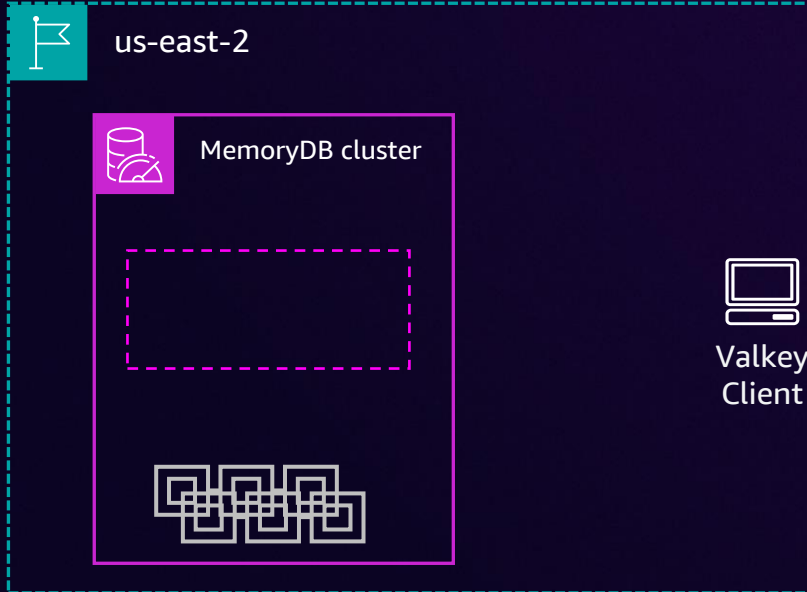
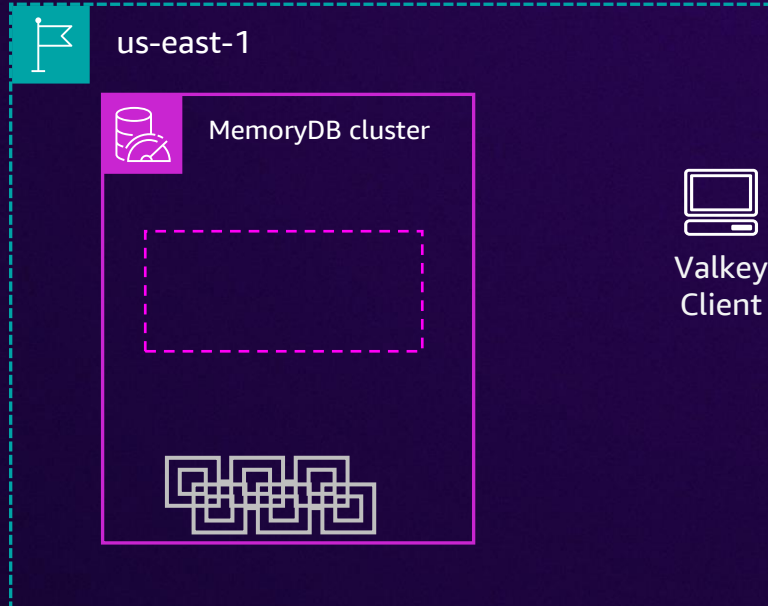
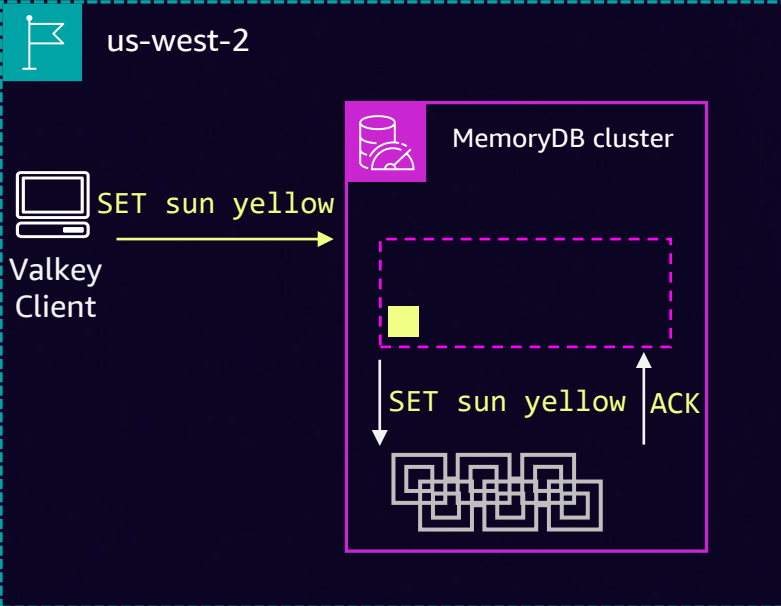
Multi-Region architecture



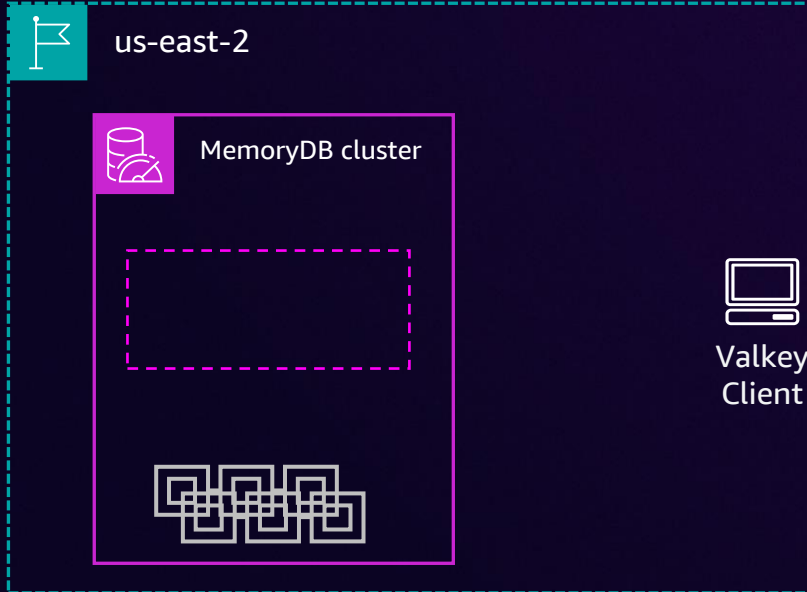
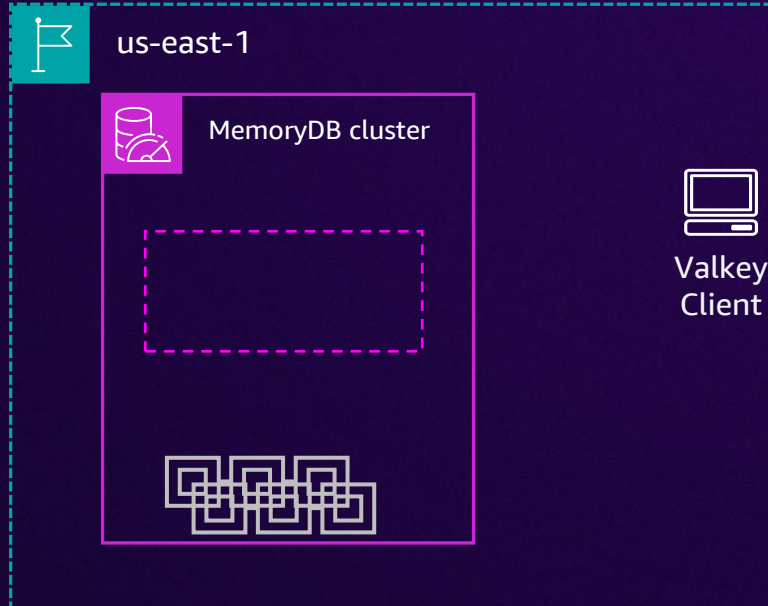
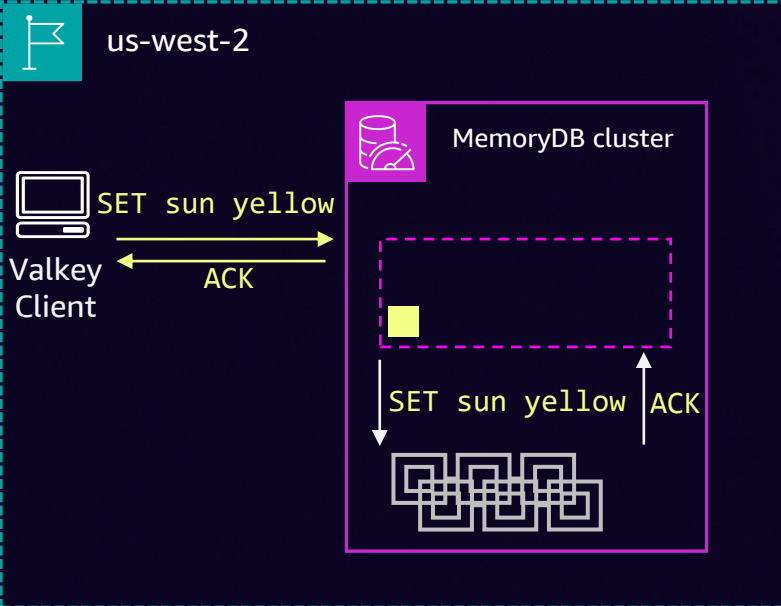
Multi-Region architecture



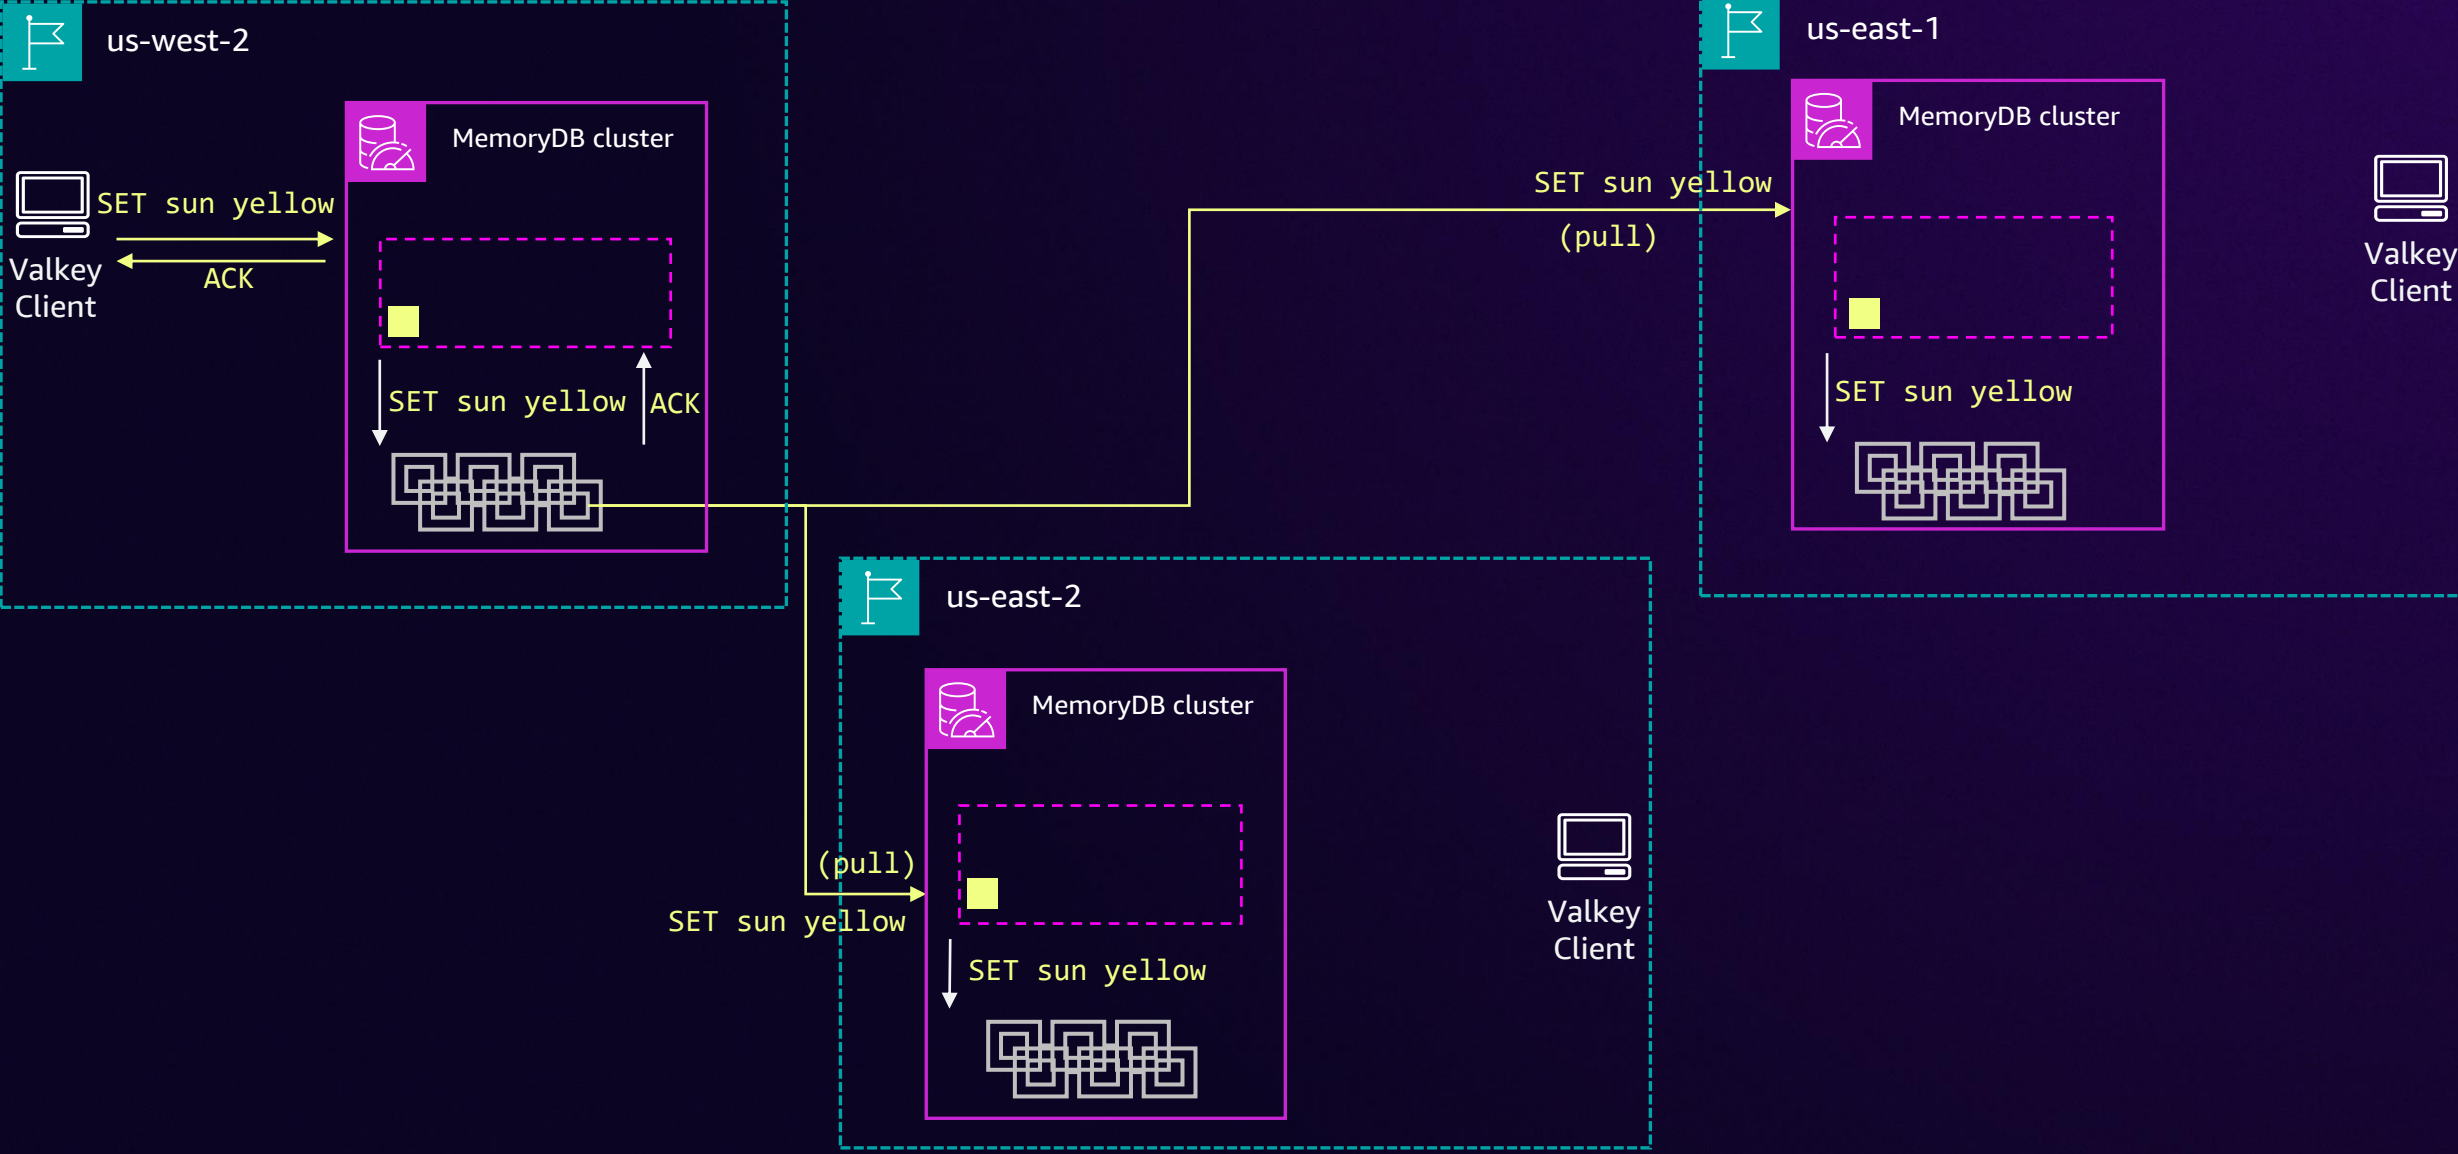
Multi-Region architecture



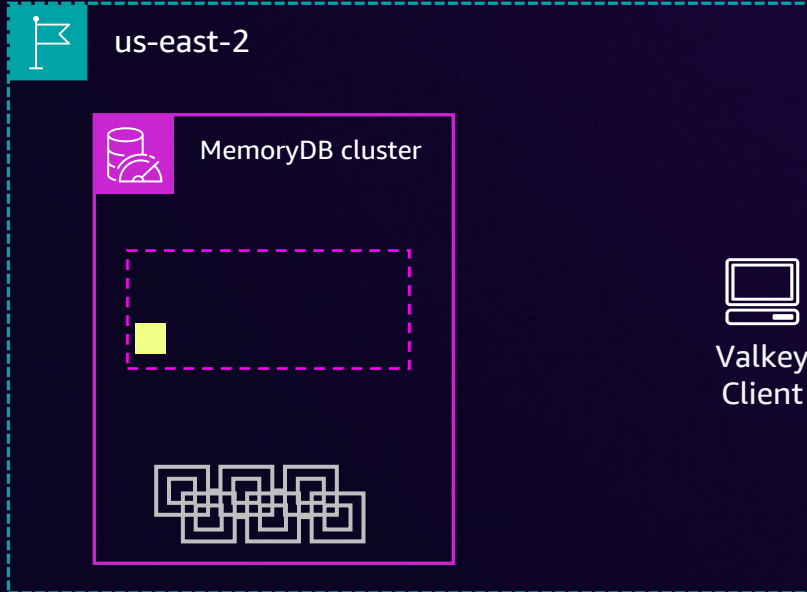
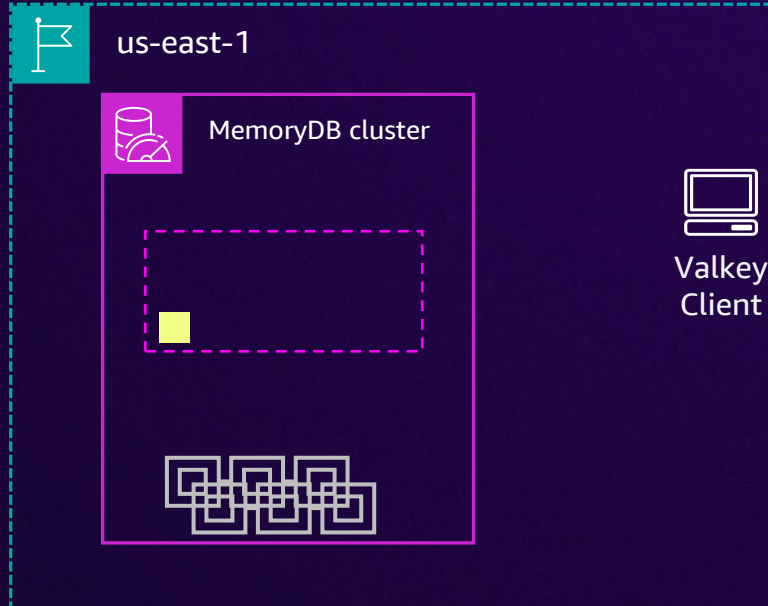
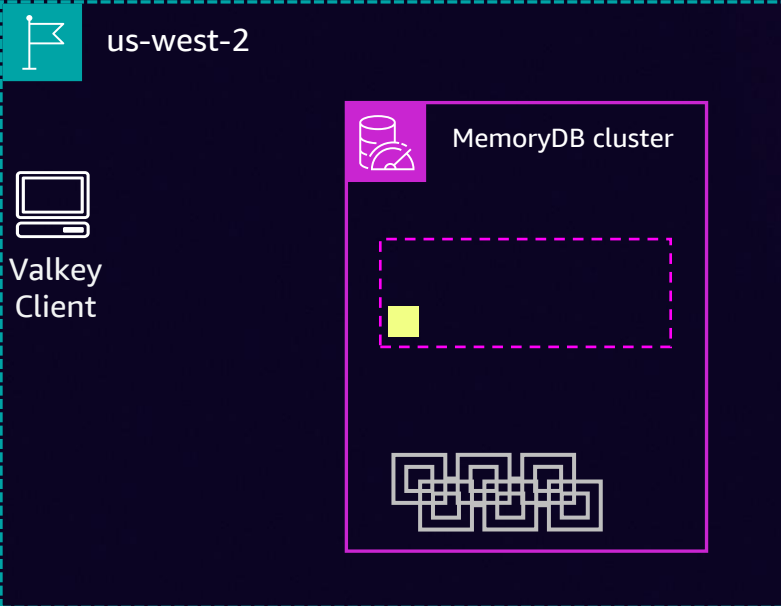
Multi-Region architecture



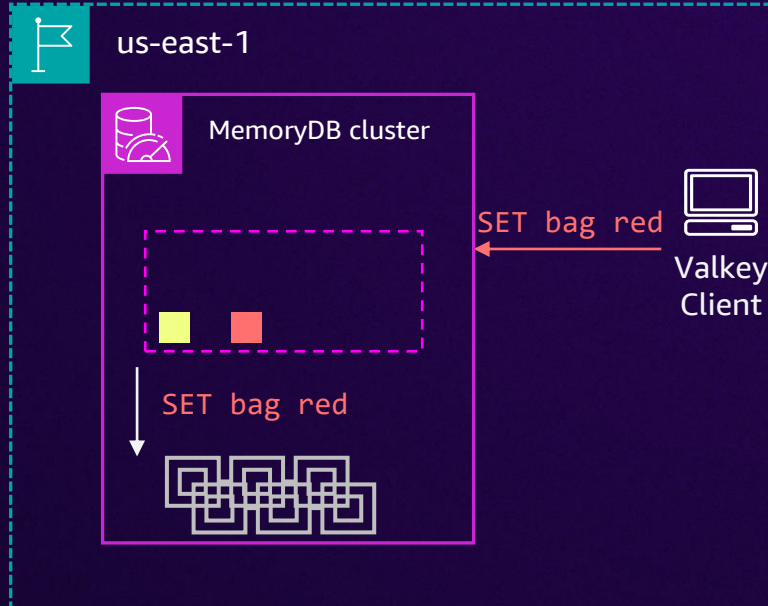
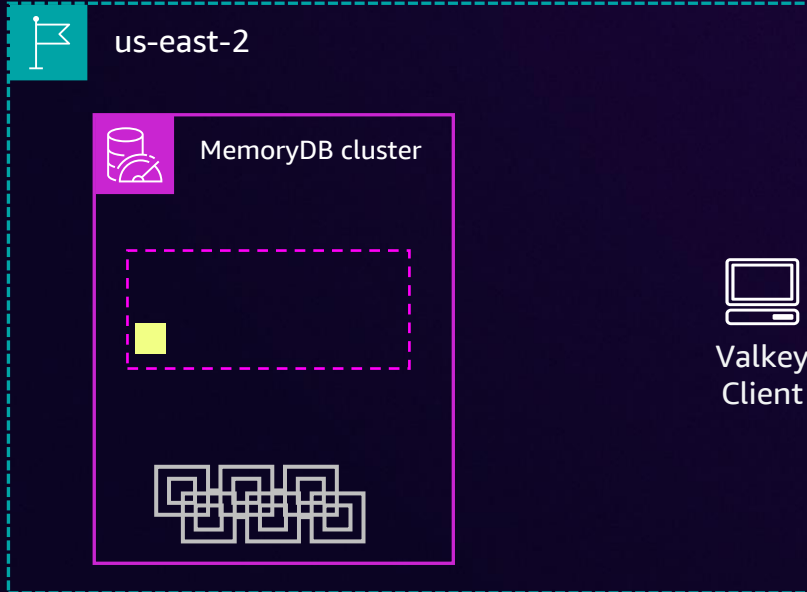
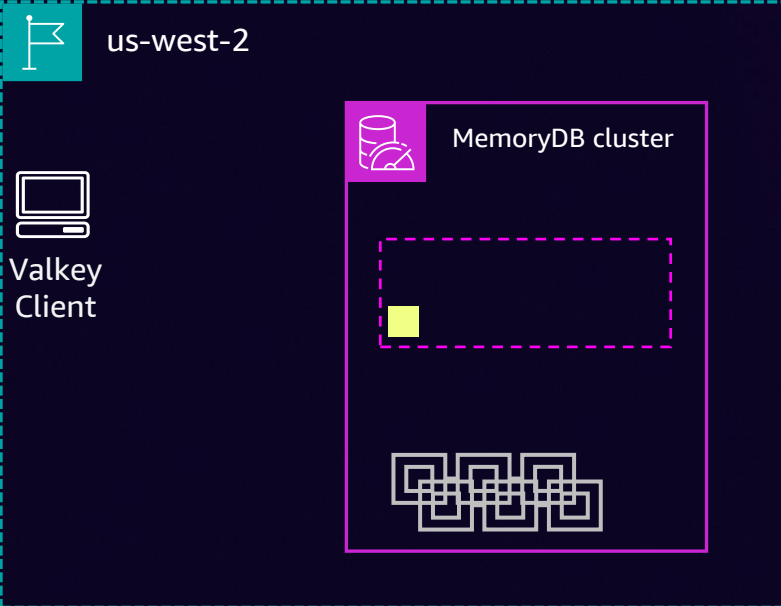
Multi-Region architecture



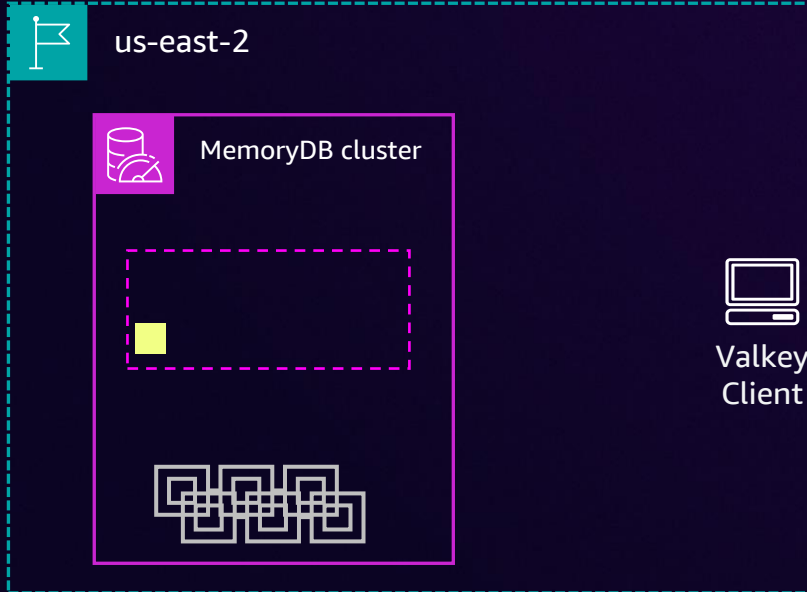
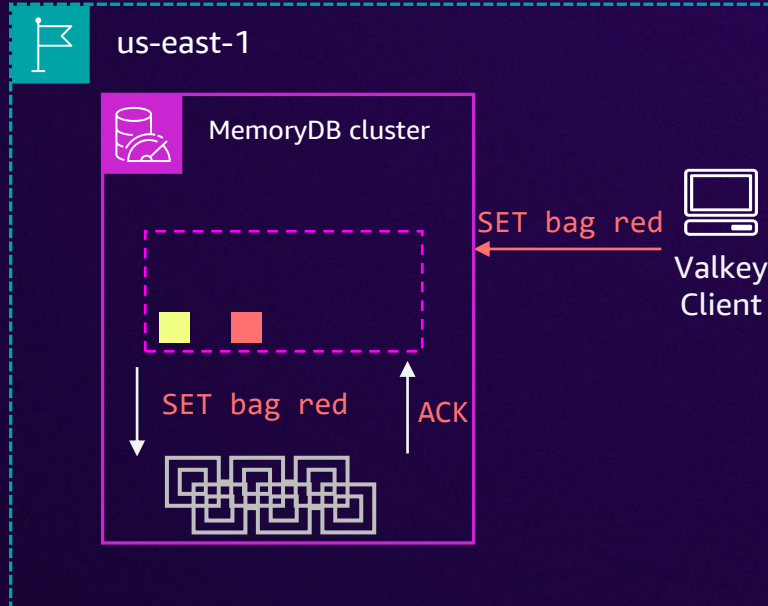
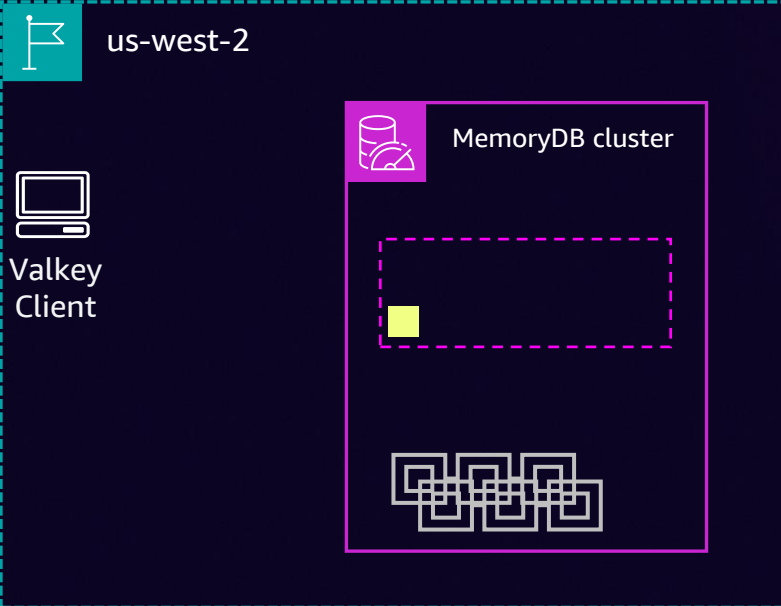
Multi-Region architecture



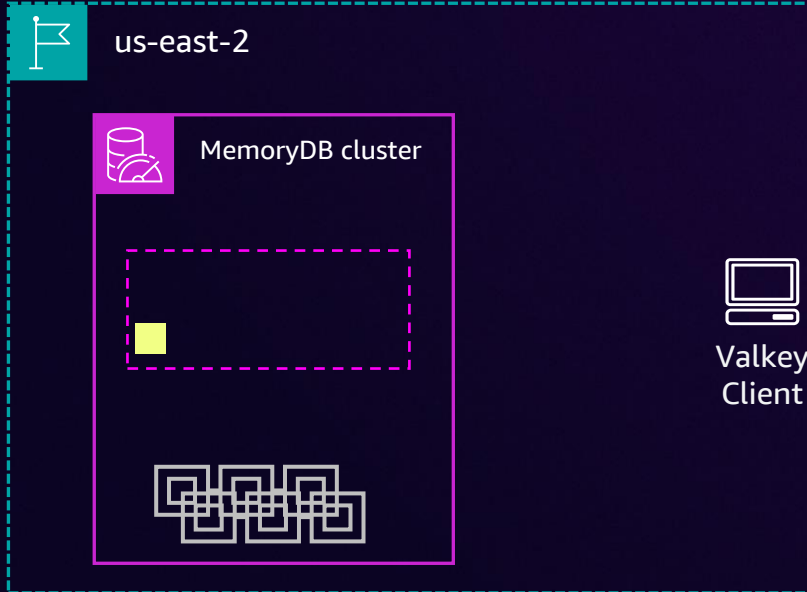
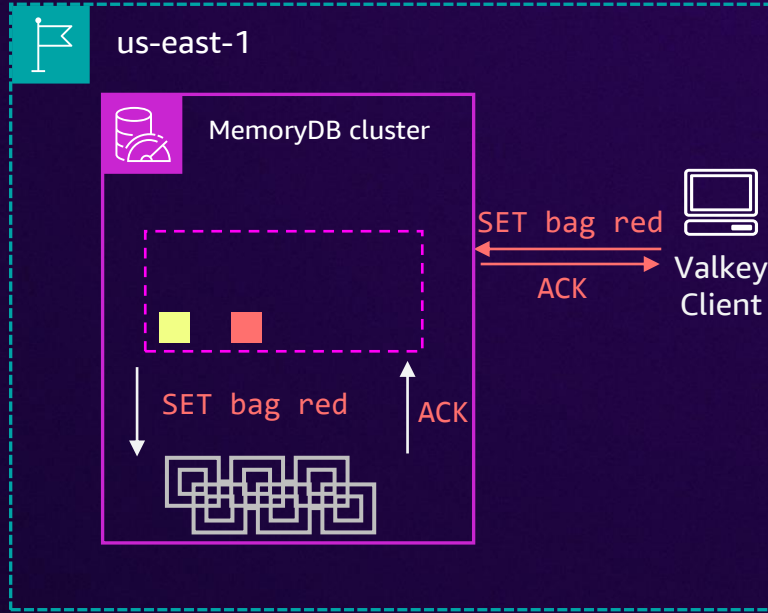
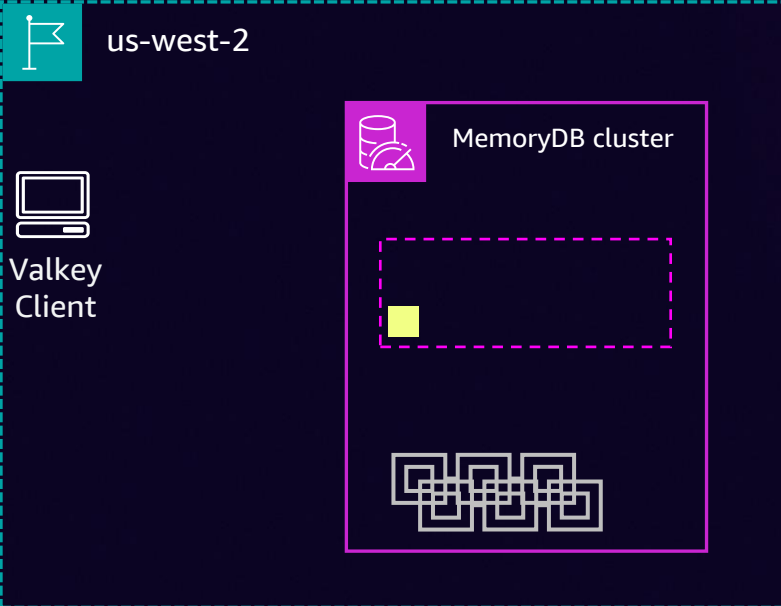
Multi-Region architecture



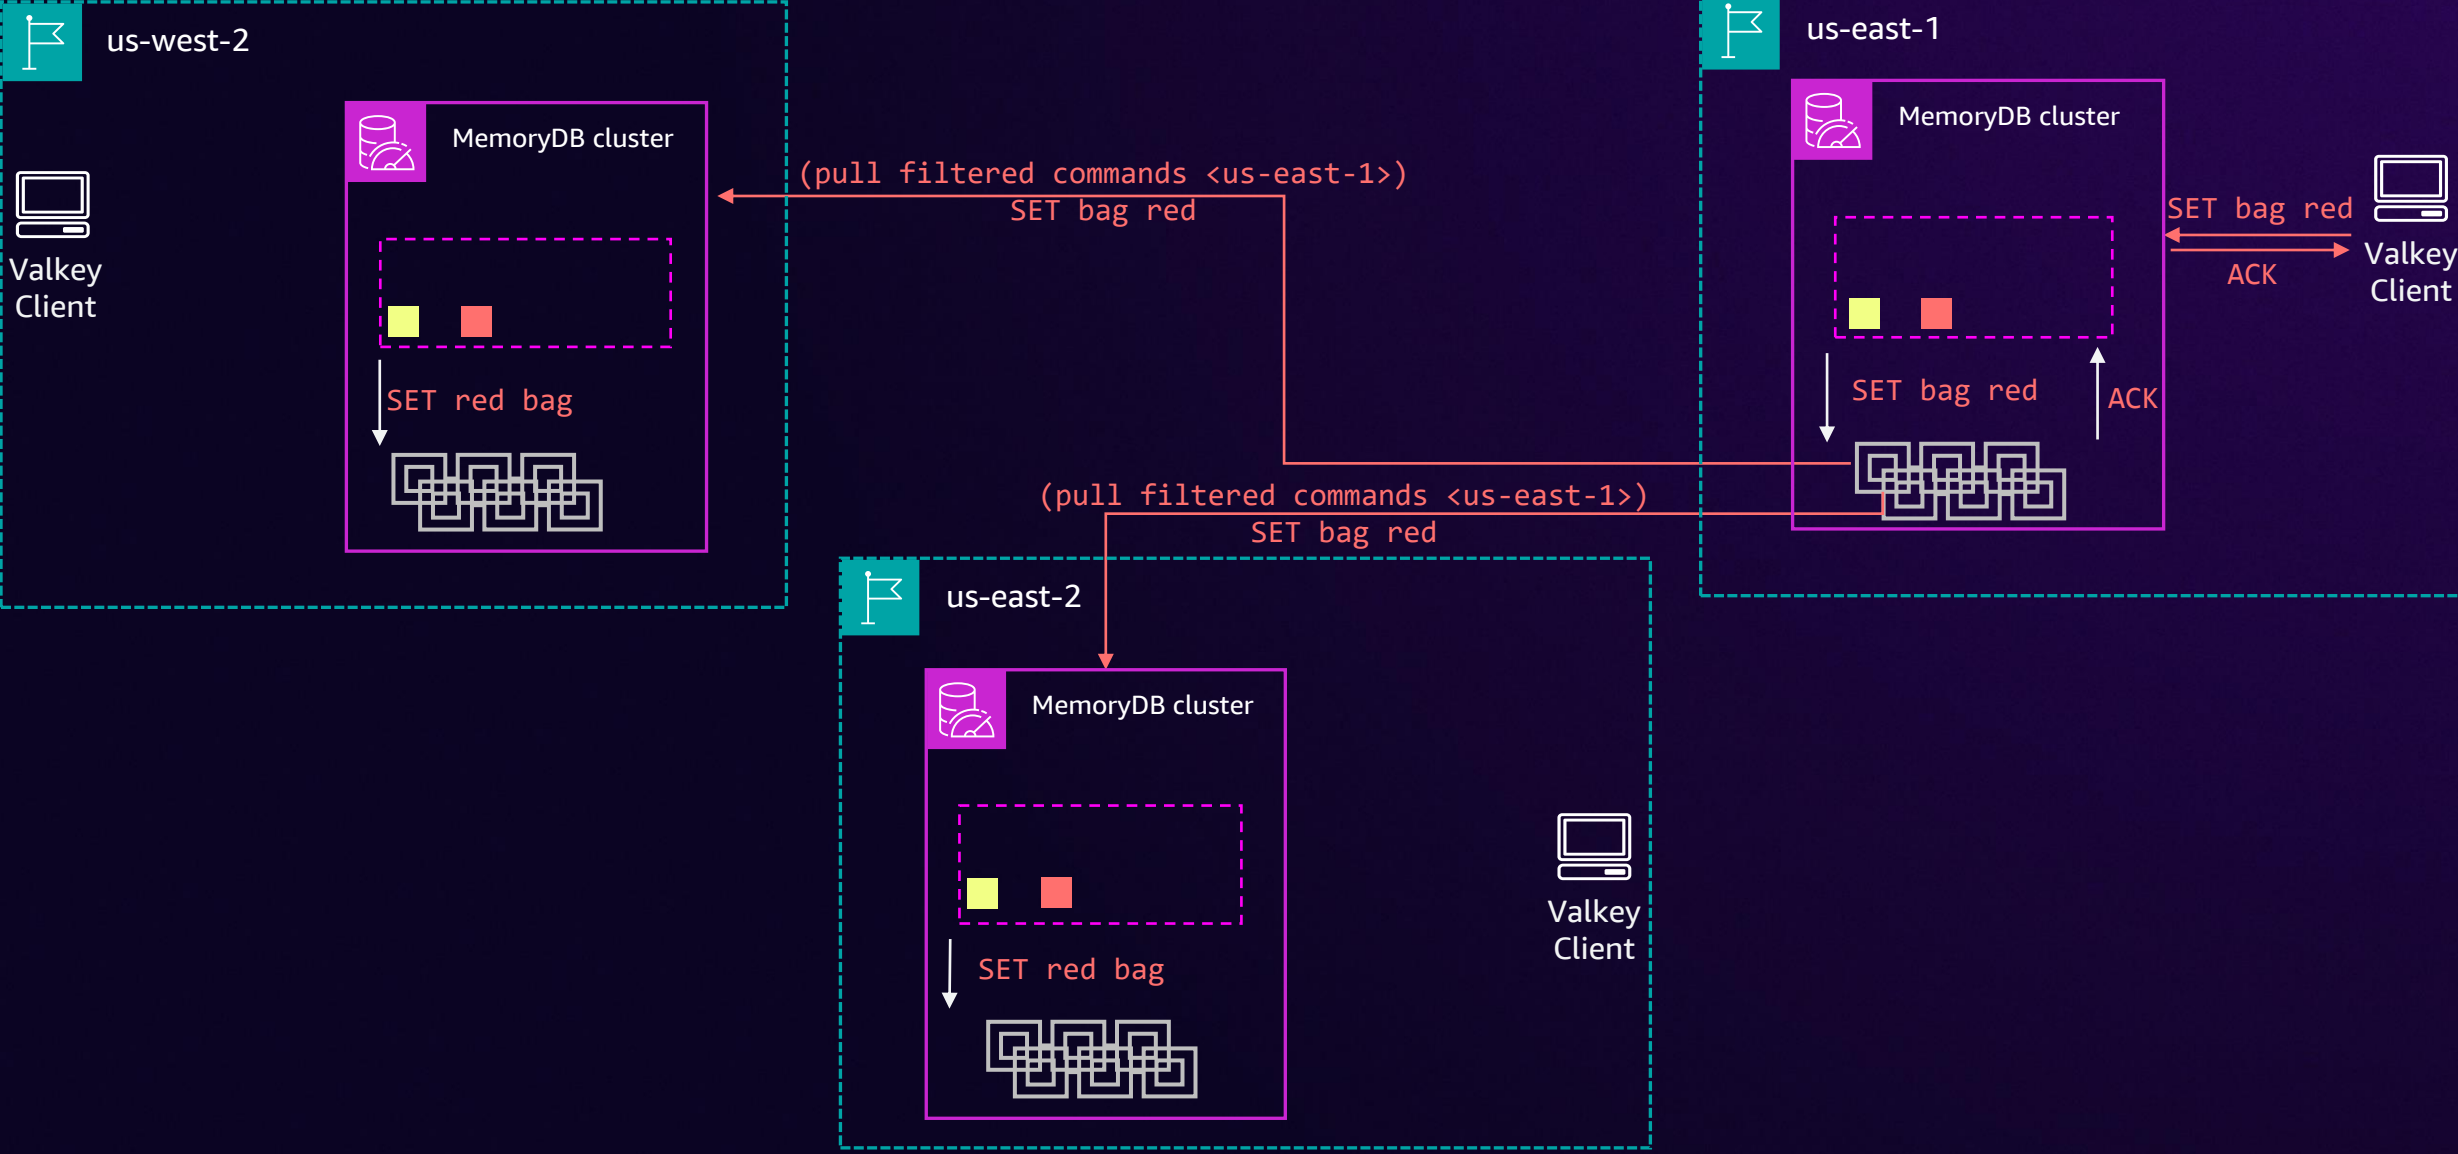
Multi-Region architecture



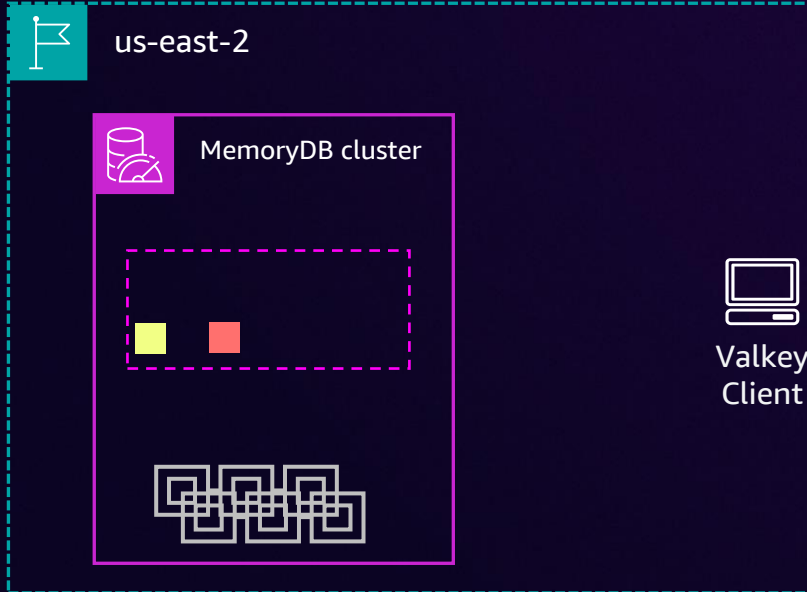
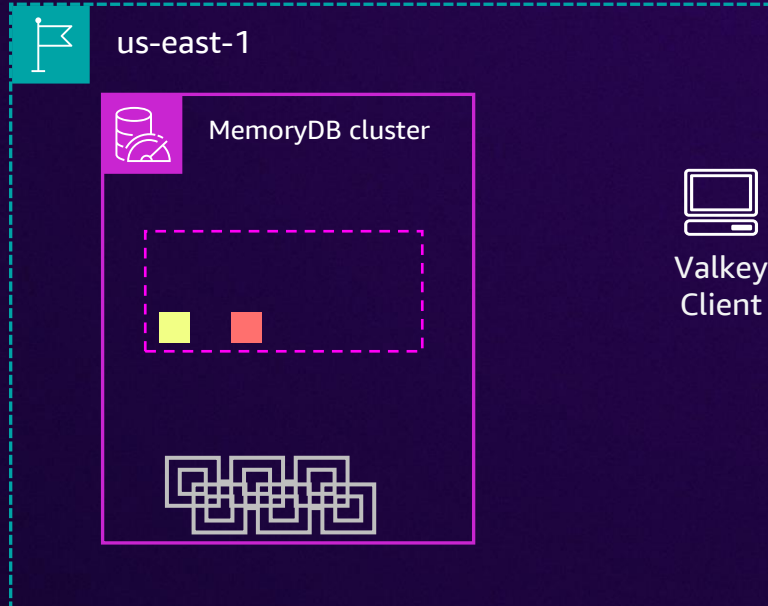
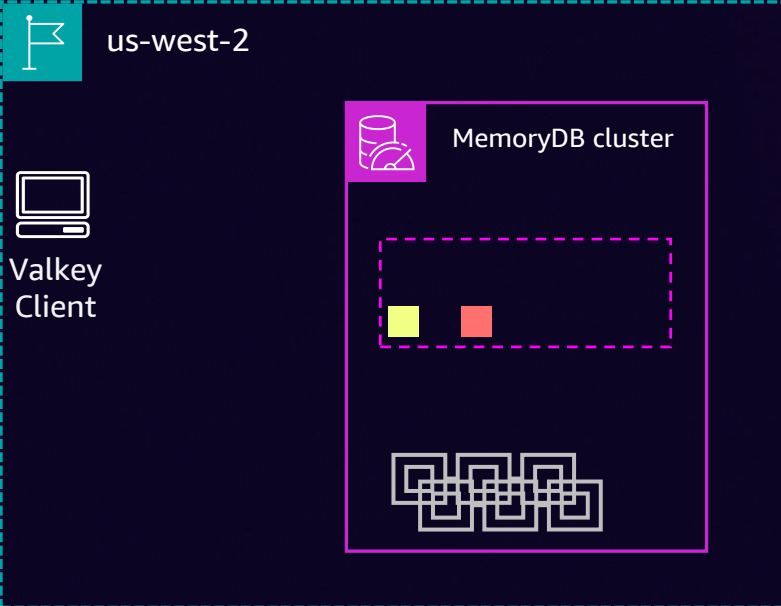
Multi-Region architecture



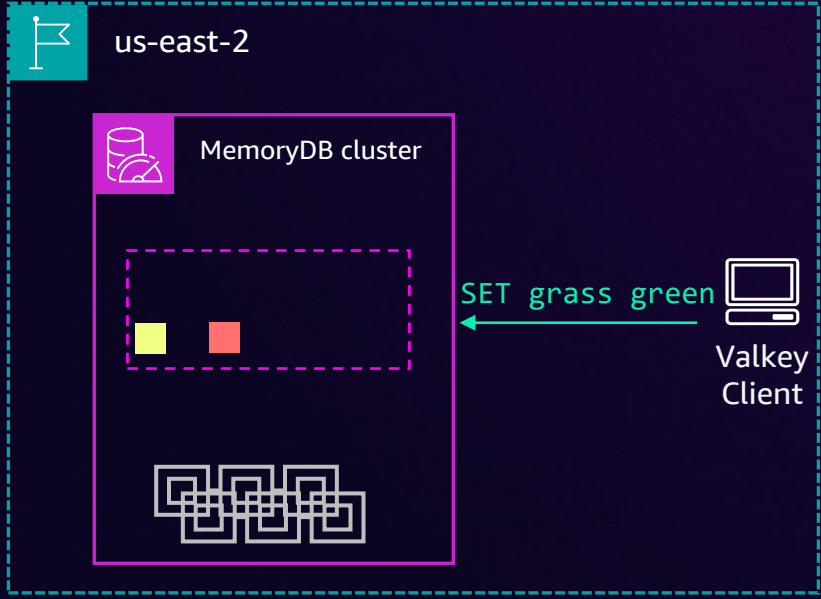
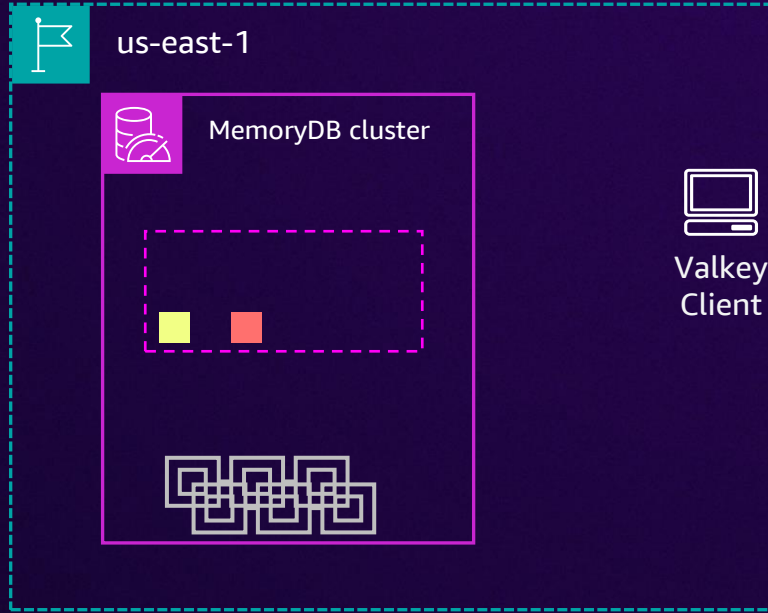
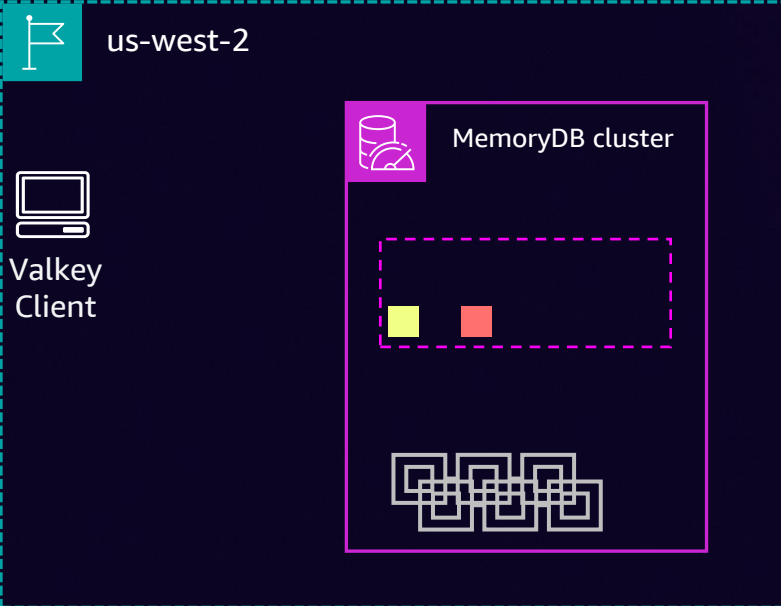
Multi-Region architecture



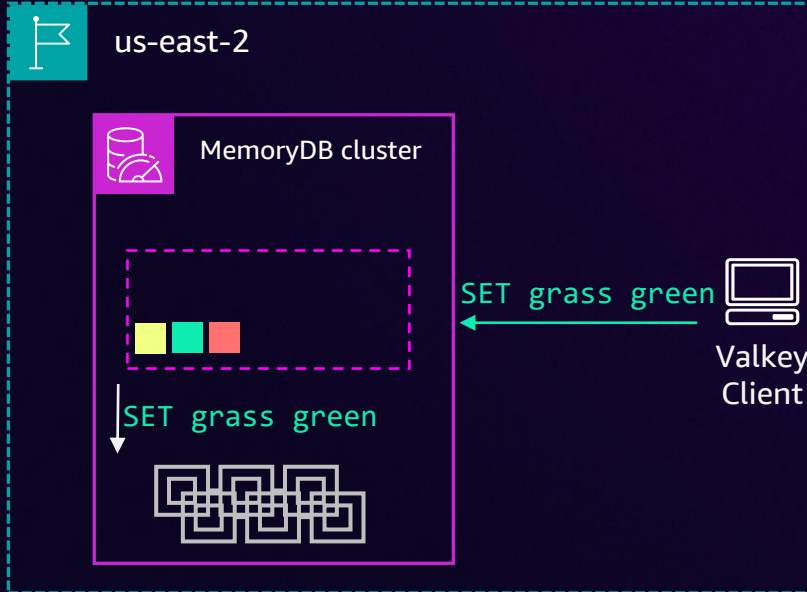
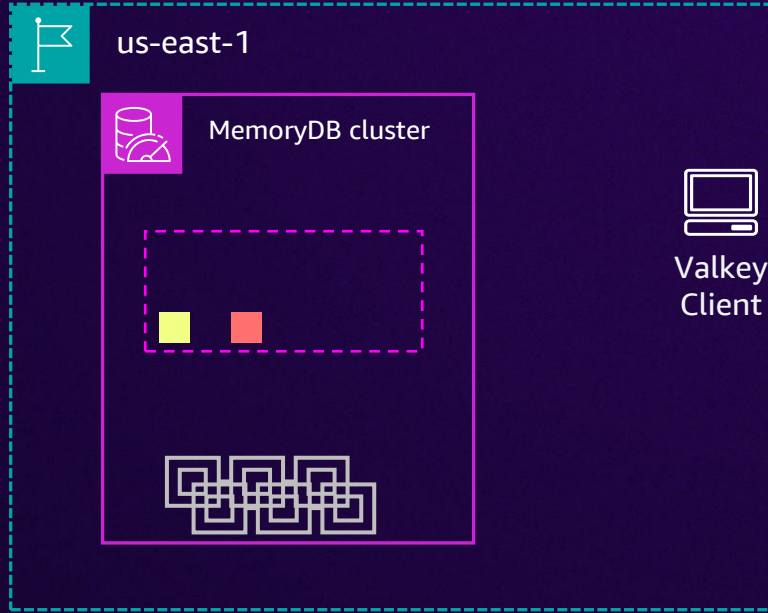
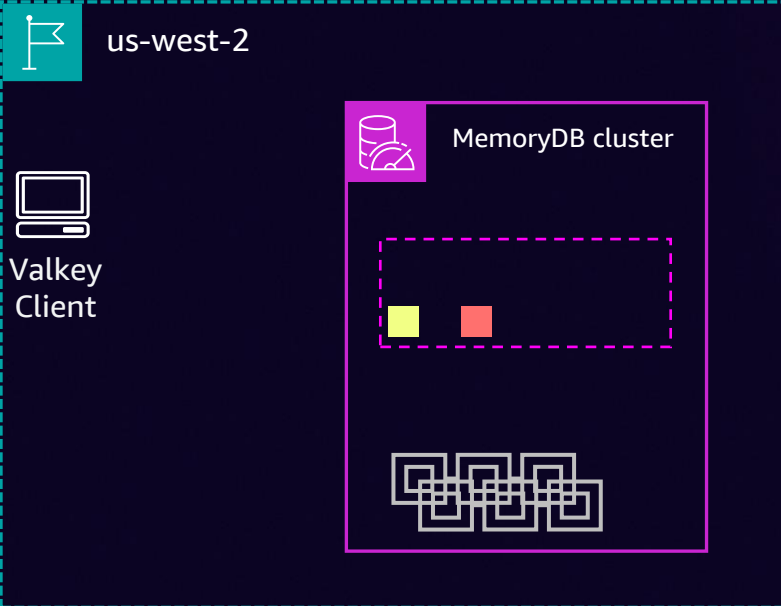
Multi-Region architecture



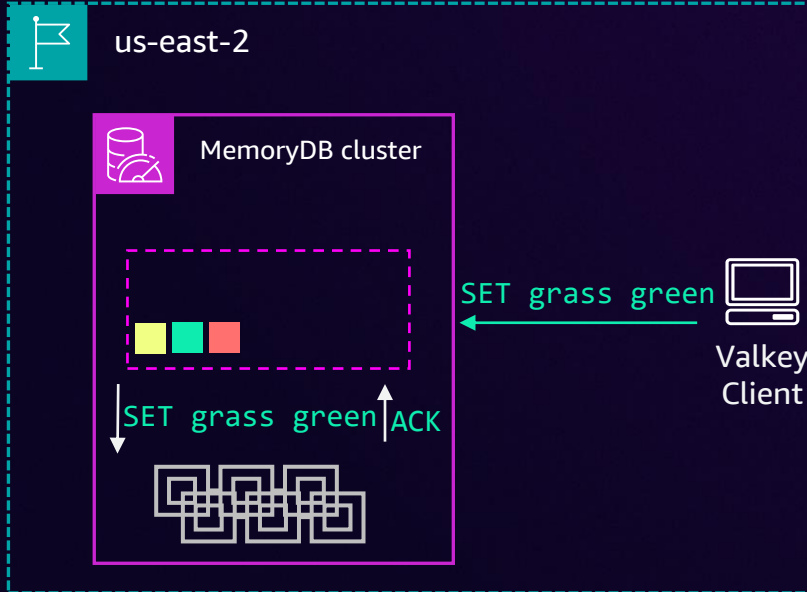
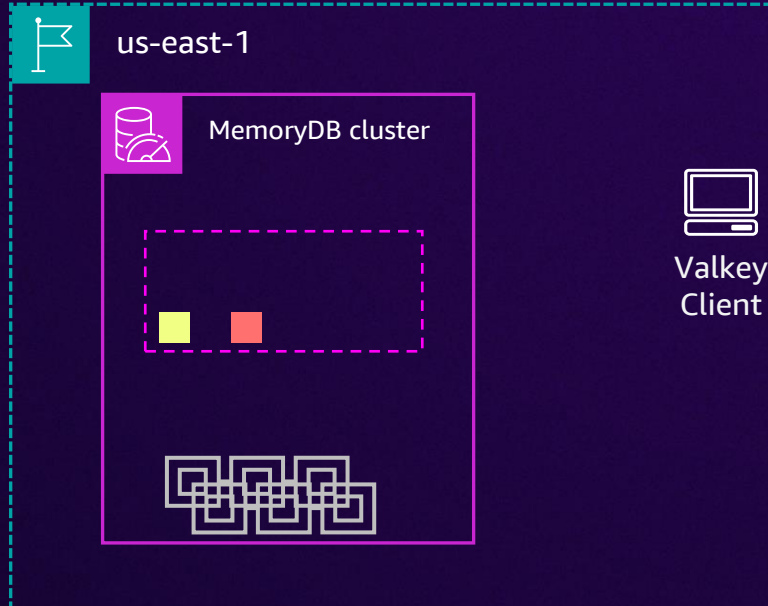
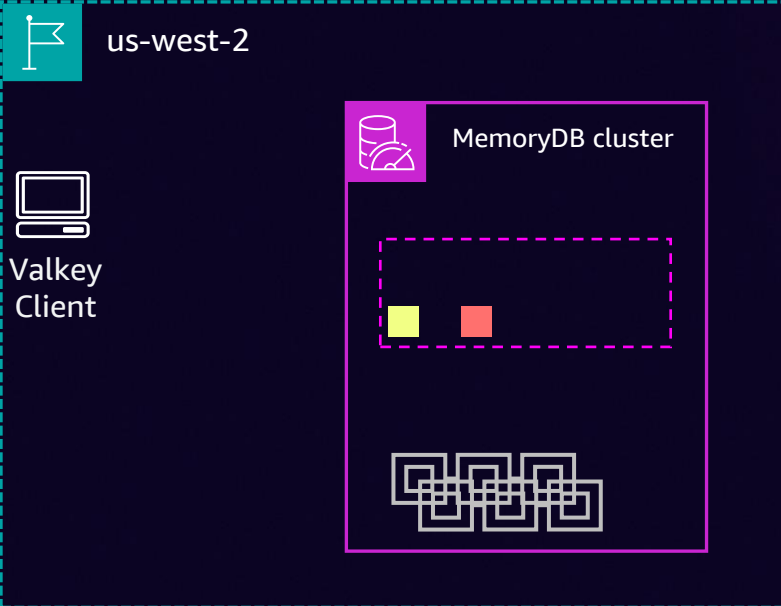
Multi-Region architecture



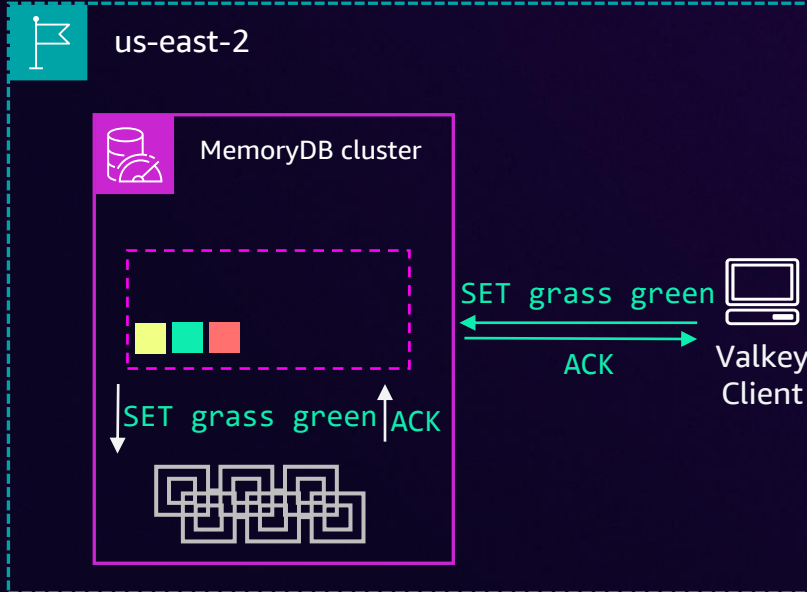
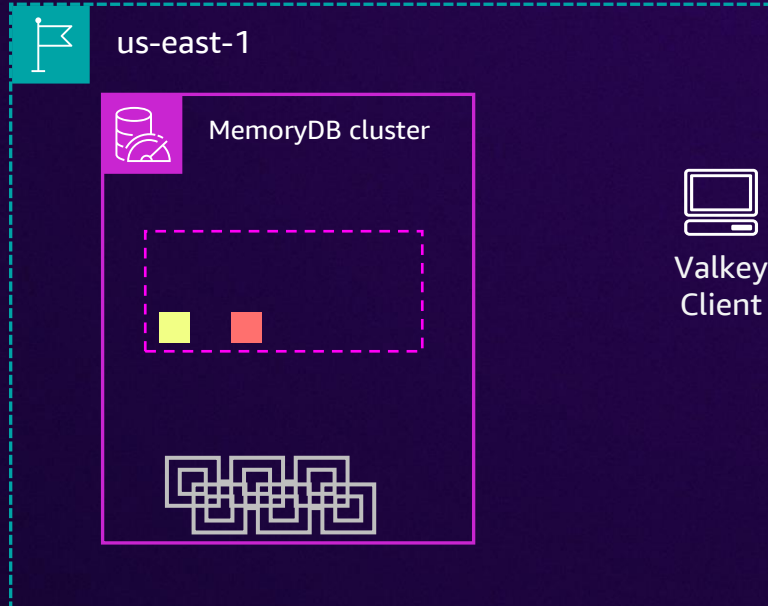
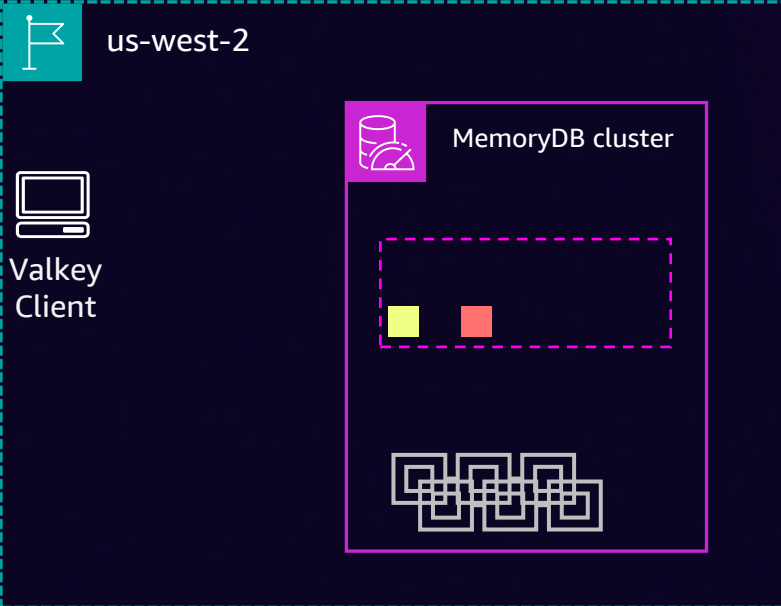
Multi-Region architecture



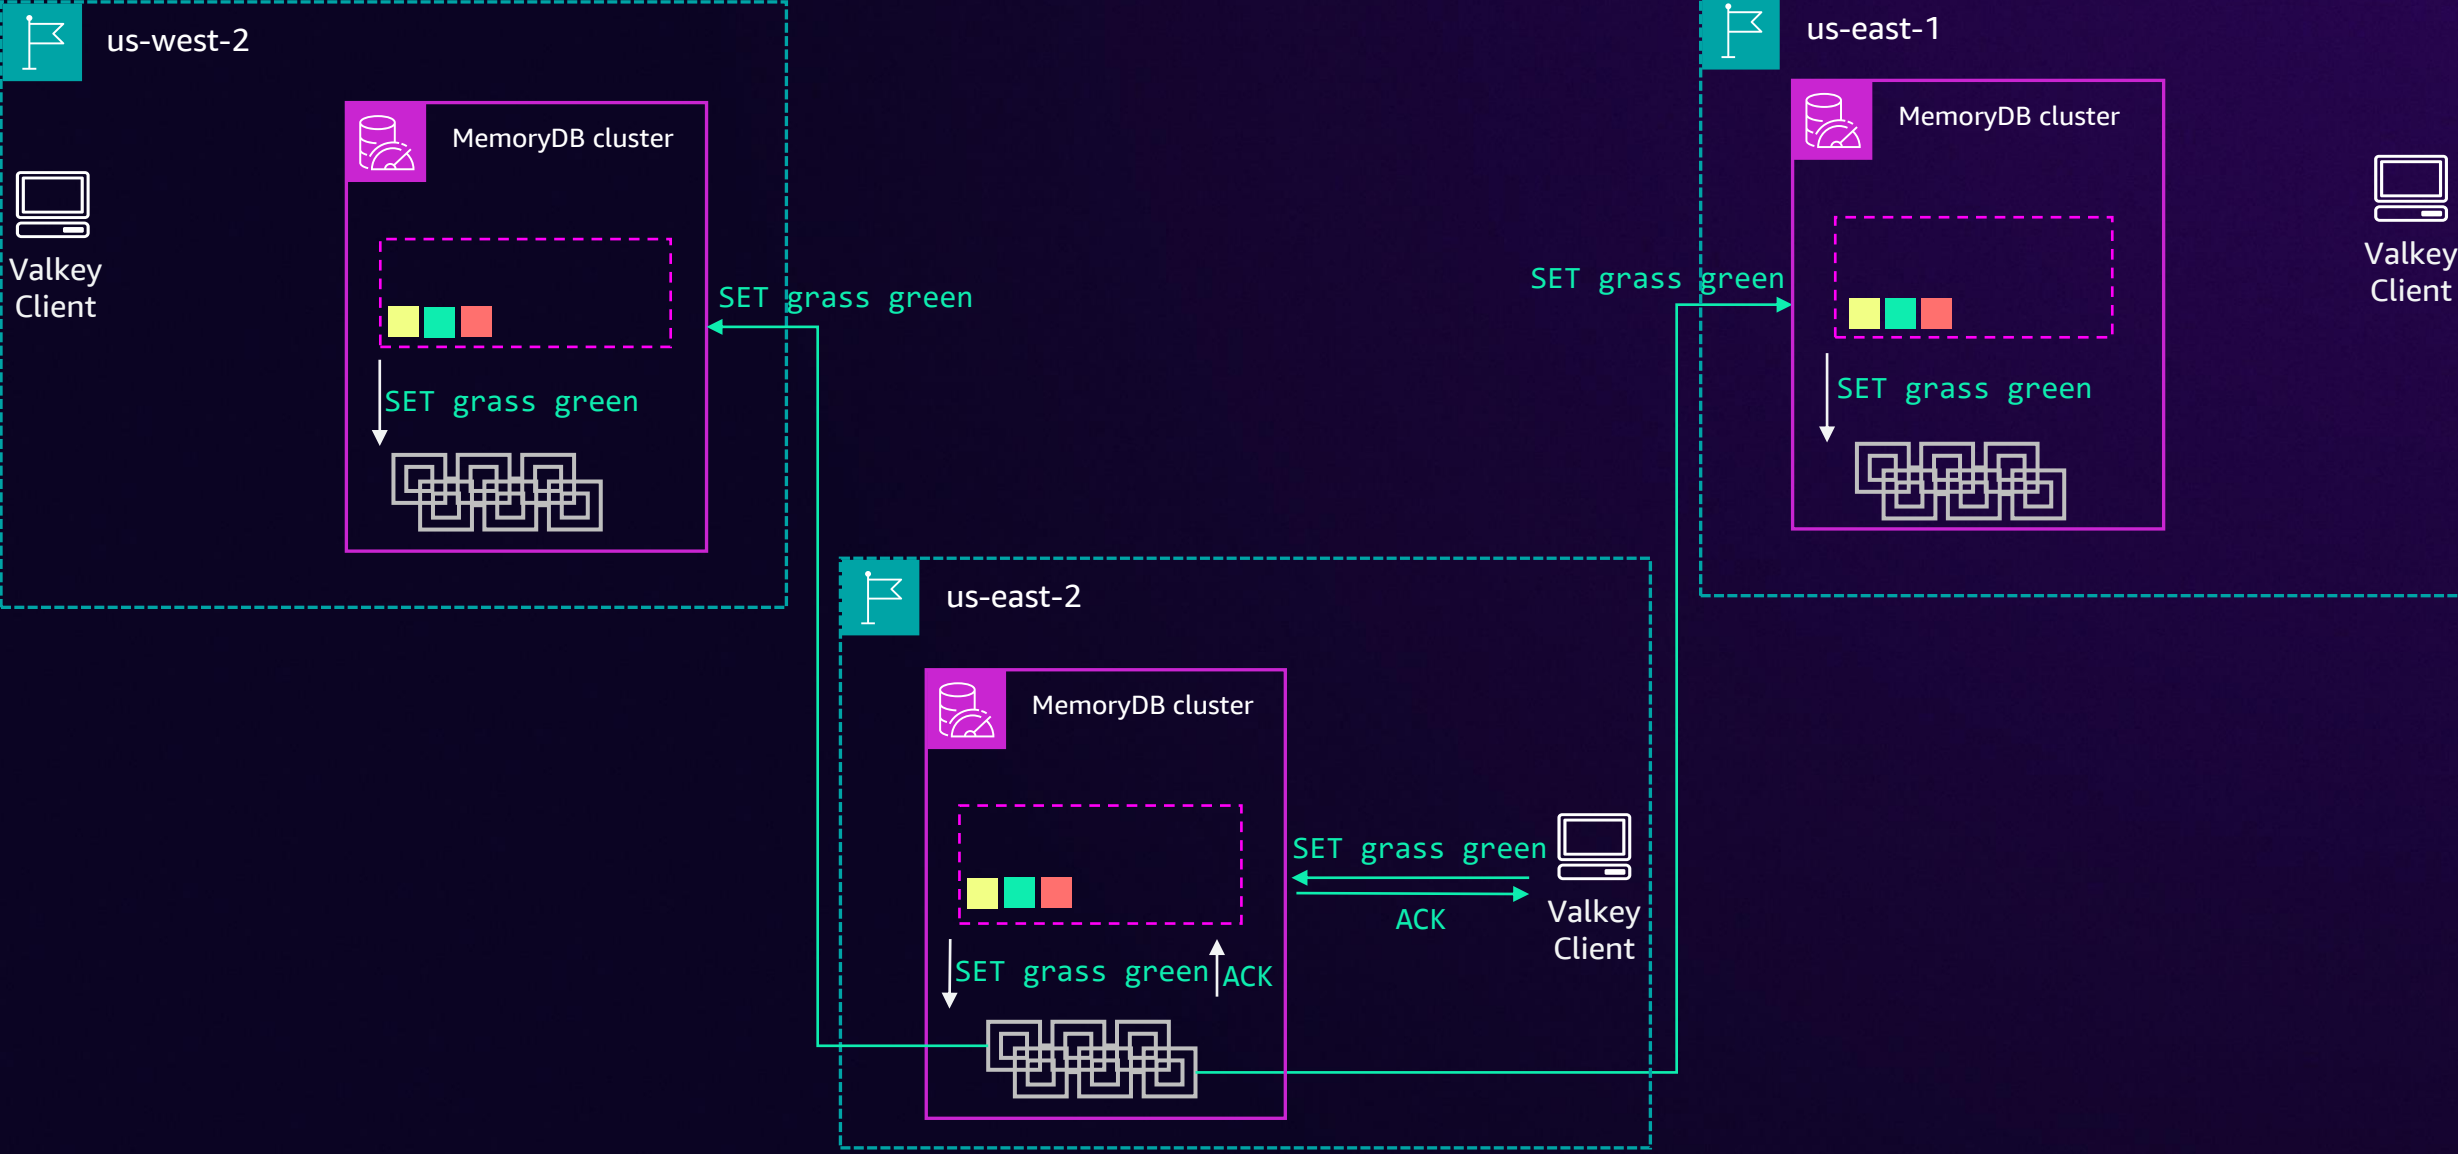
Multi-Region architecture



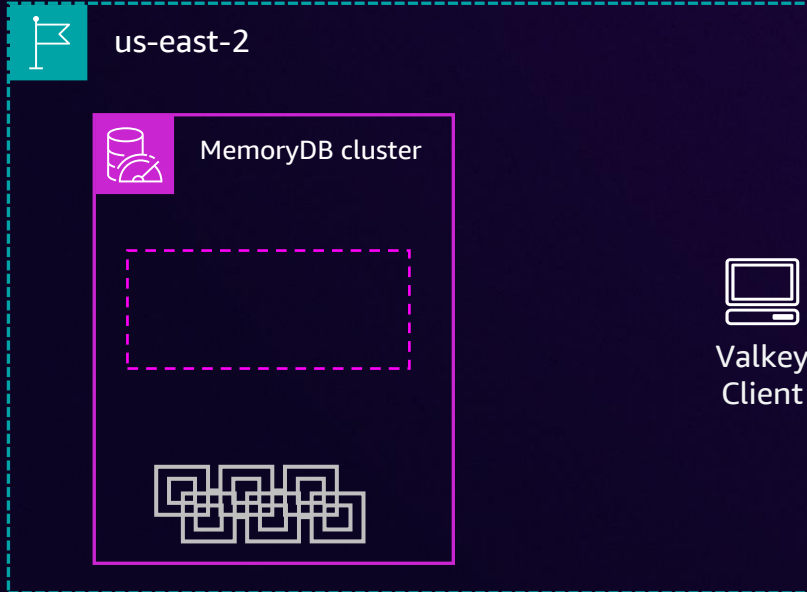
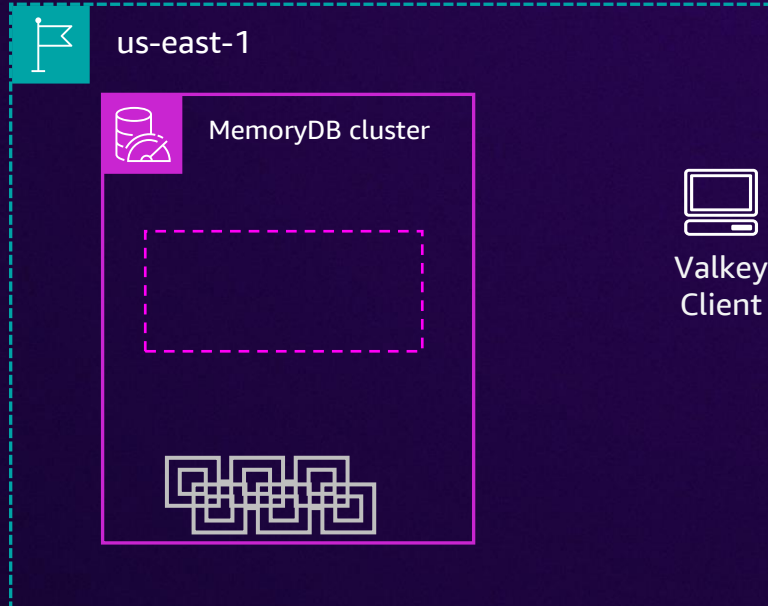
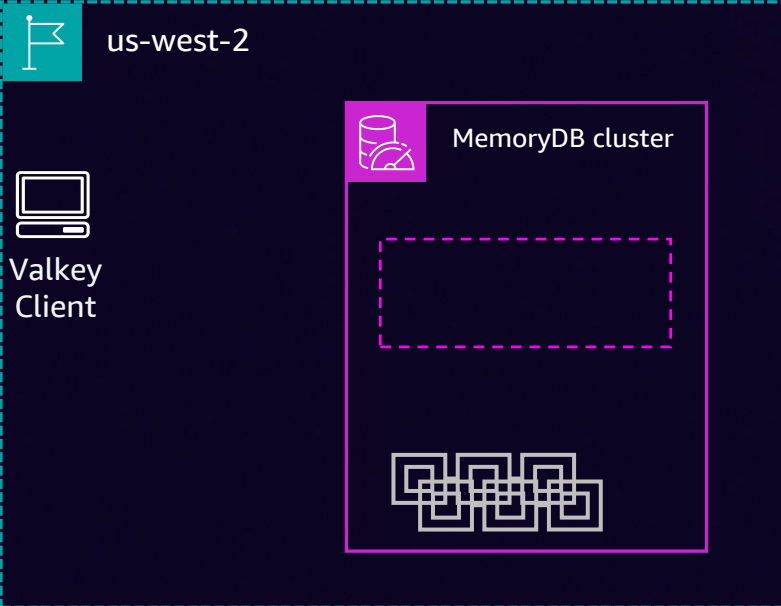
Multi-Region architecture



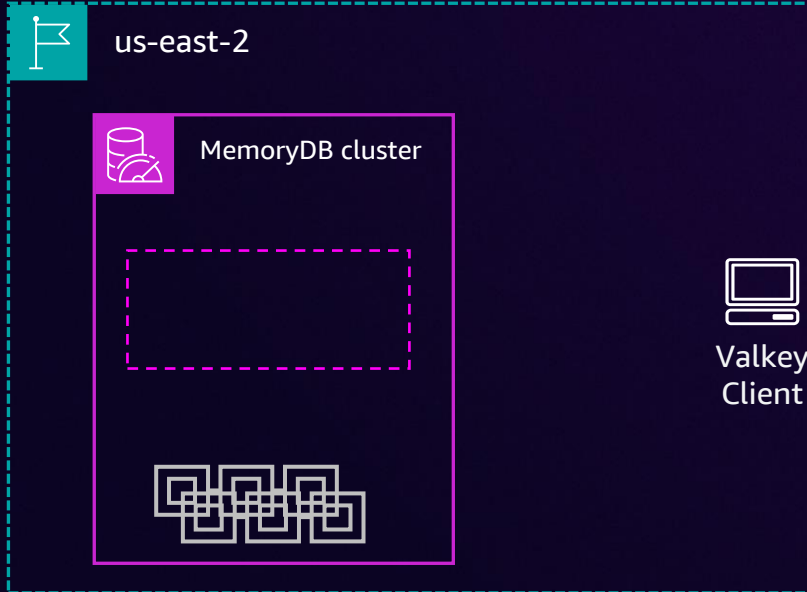
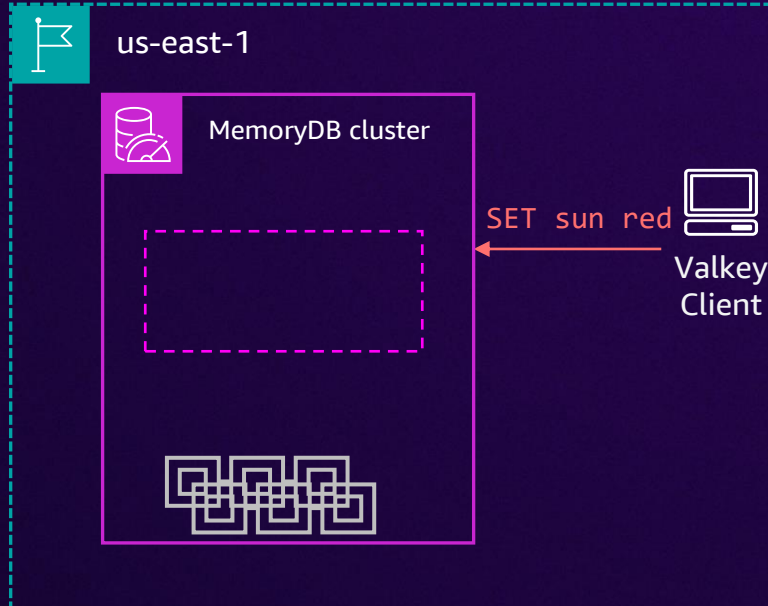
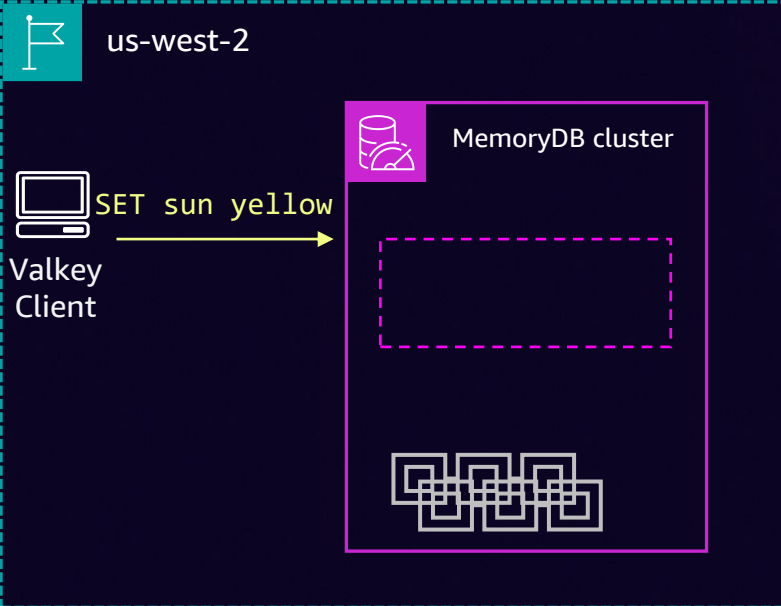
Multi-Region architecture



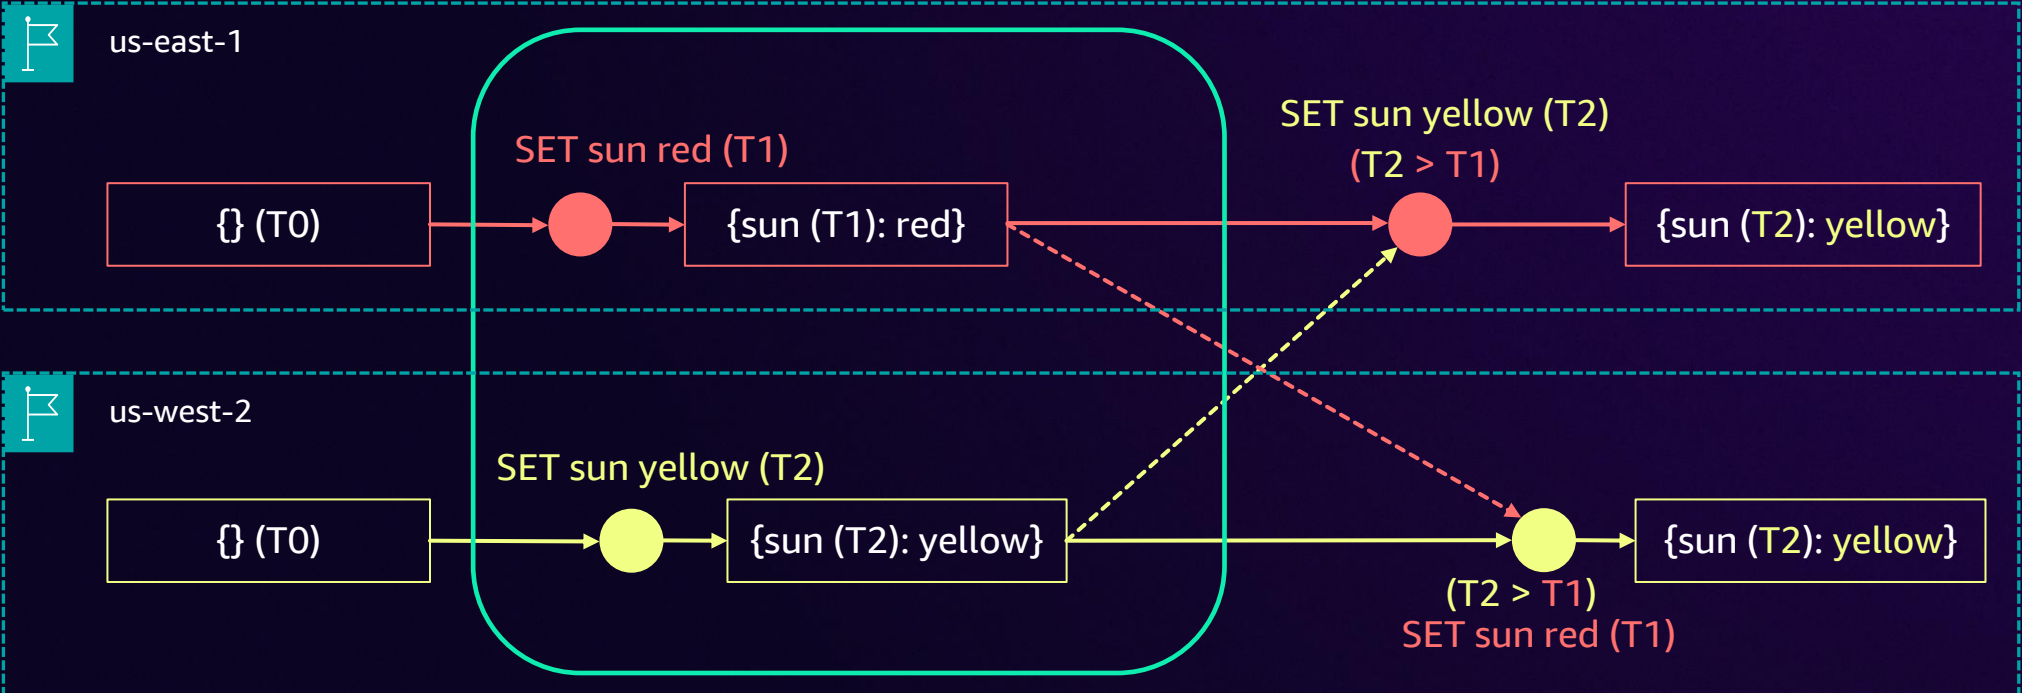
Multi-Region architecture



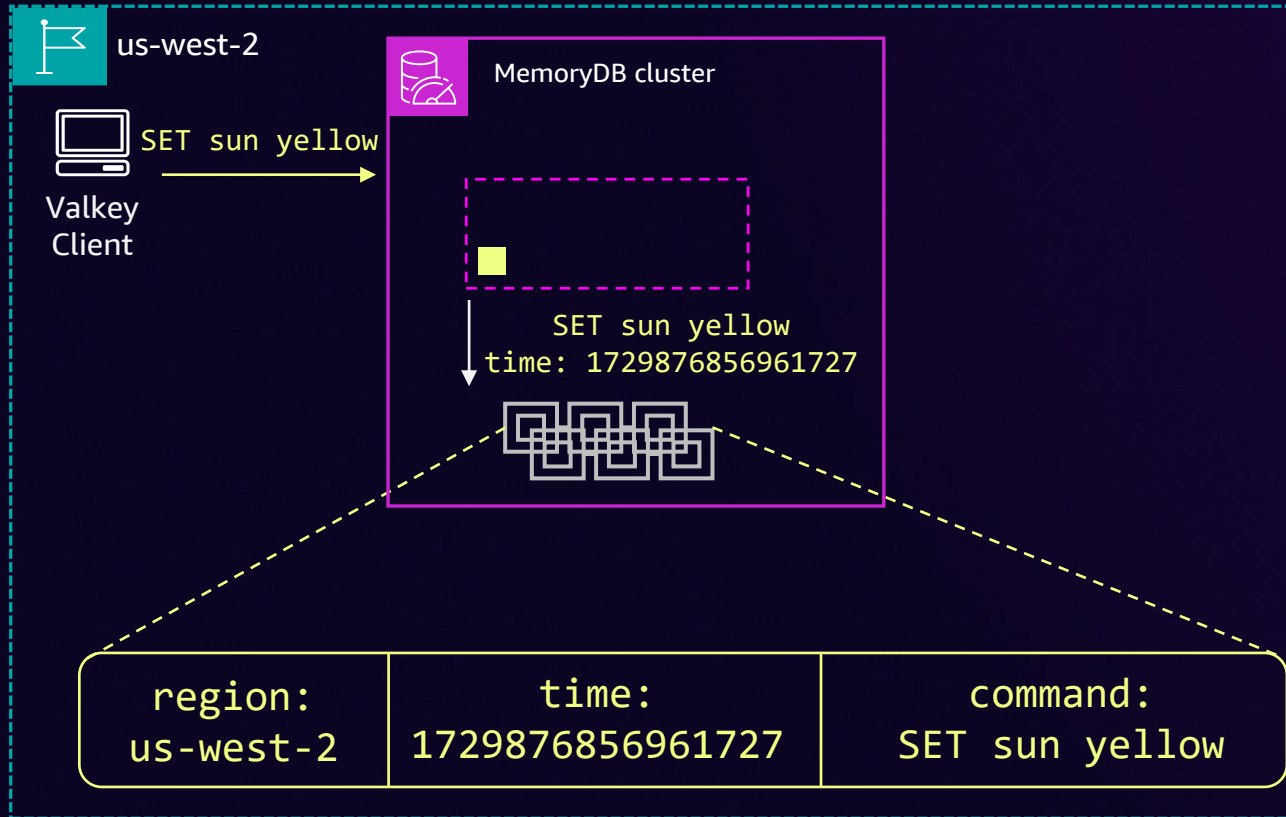
Multi-Region architecture



What if there are conflicts?



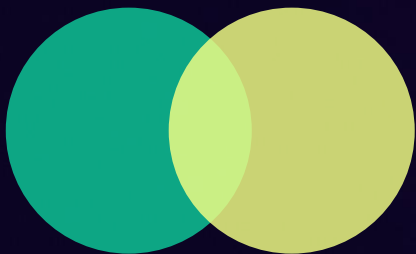
Conflict resolution strategy – Last writer wins



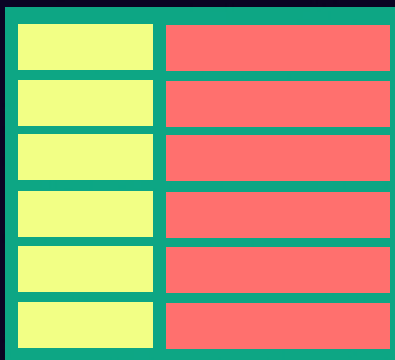
Key	Value	Region	Timestamp
sun	yellow	us-west-2	1729876856961727

Commands are labeled with μ s timestamps and the Region of origin

Valkey is more than a key-value store

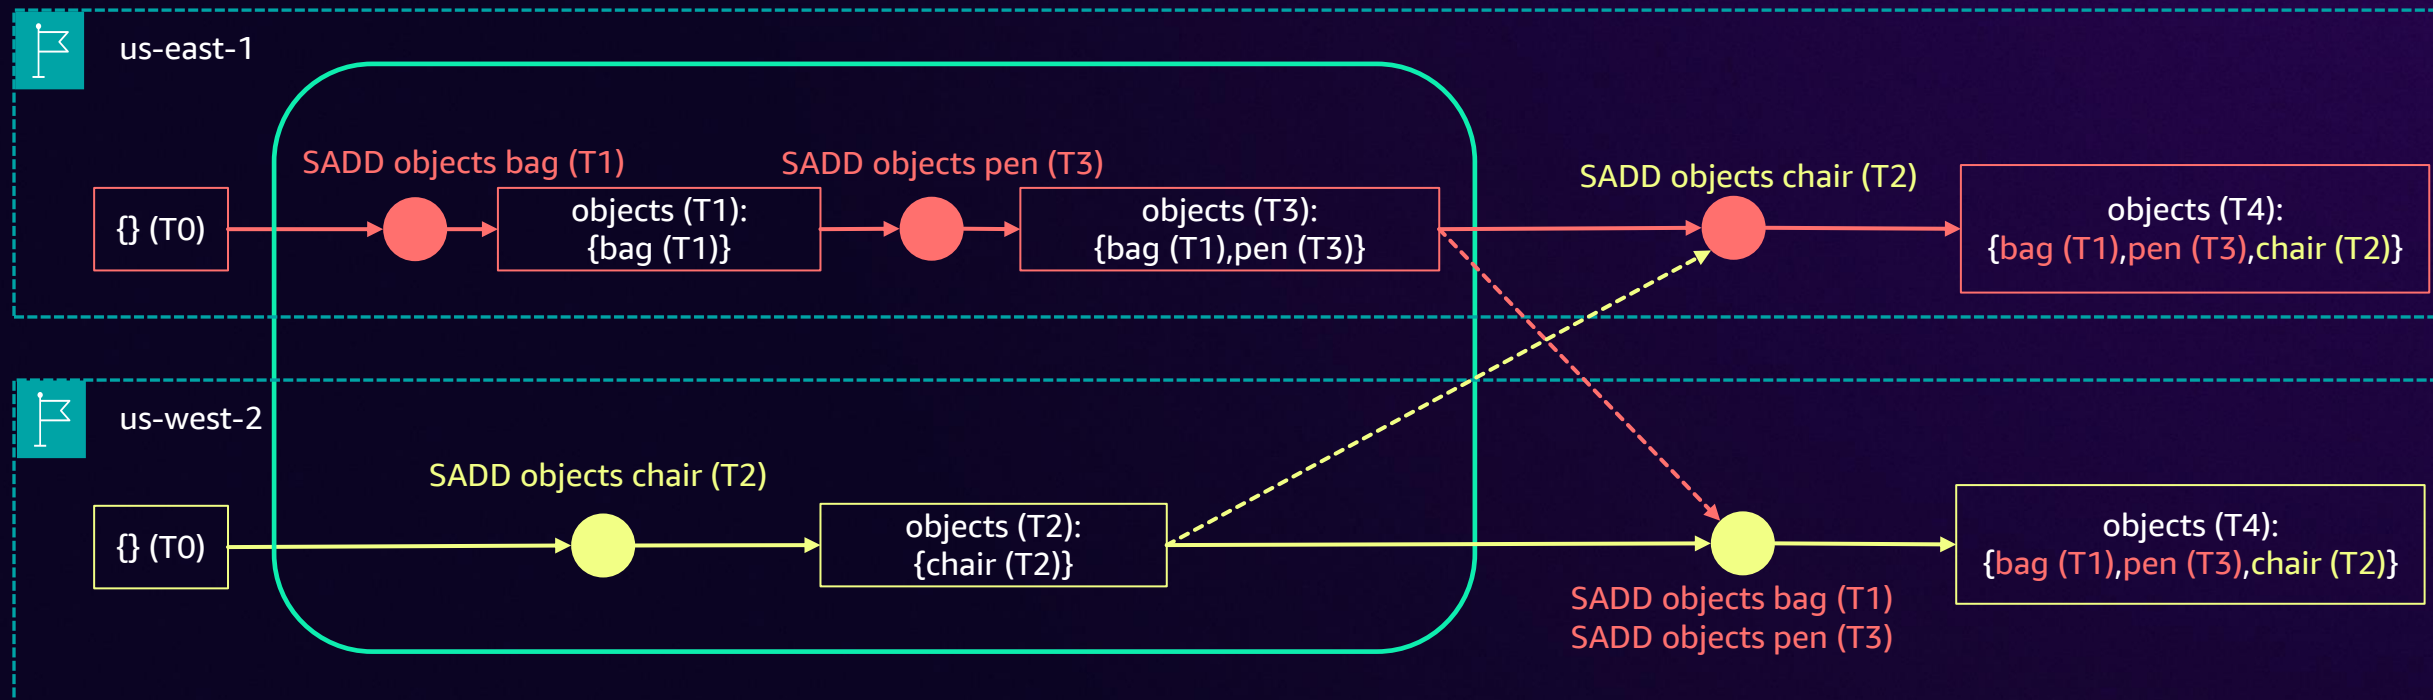


Sets/Sorted sets: Leaderboard, count unique items



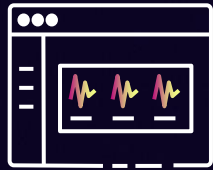
Hashmaps: Session stores

Is "last writer wins" sufficient?

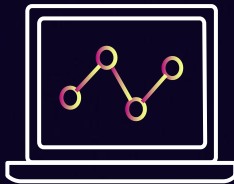


SADD: Valkey command to add an element into a set

CRDT – Conflict-free replicated data type



Allow multiple Regions to write concurrently even on the same key



Writes are propagated asynchronously with the original timestamps (μ s)



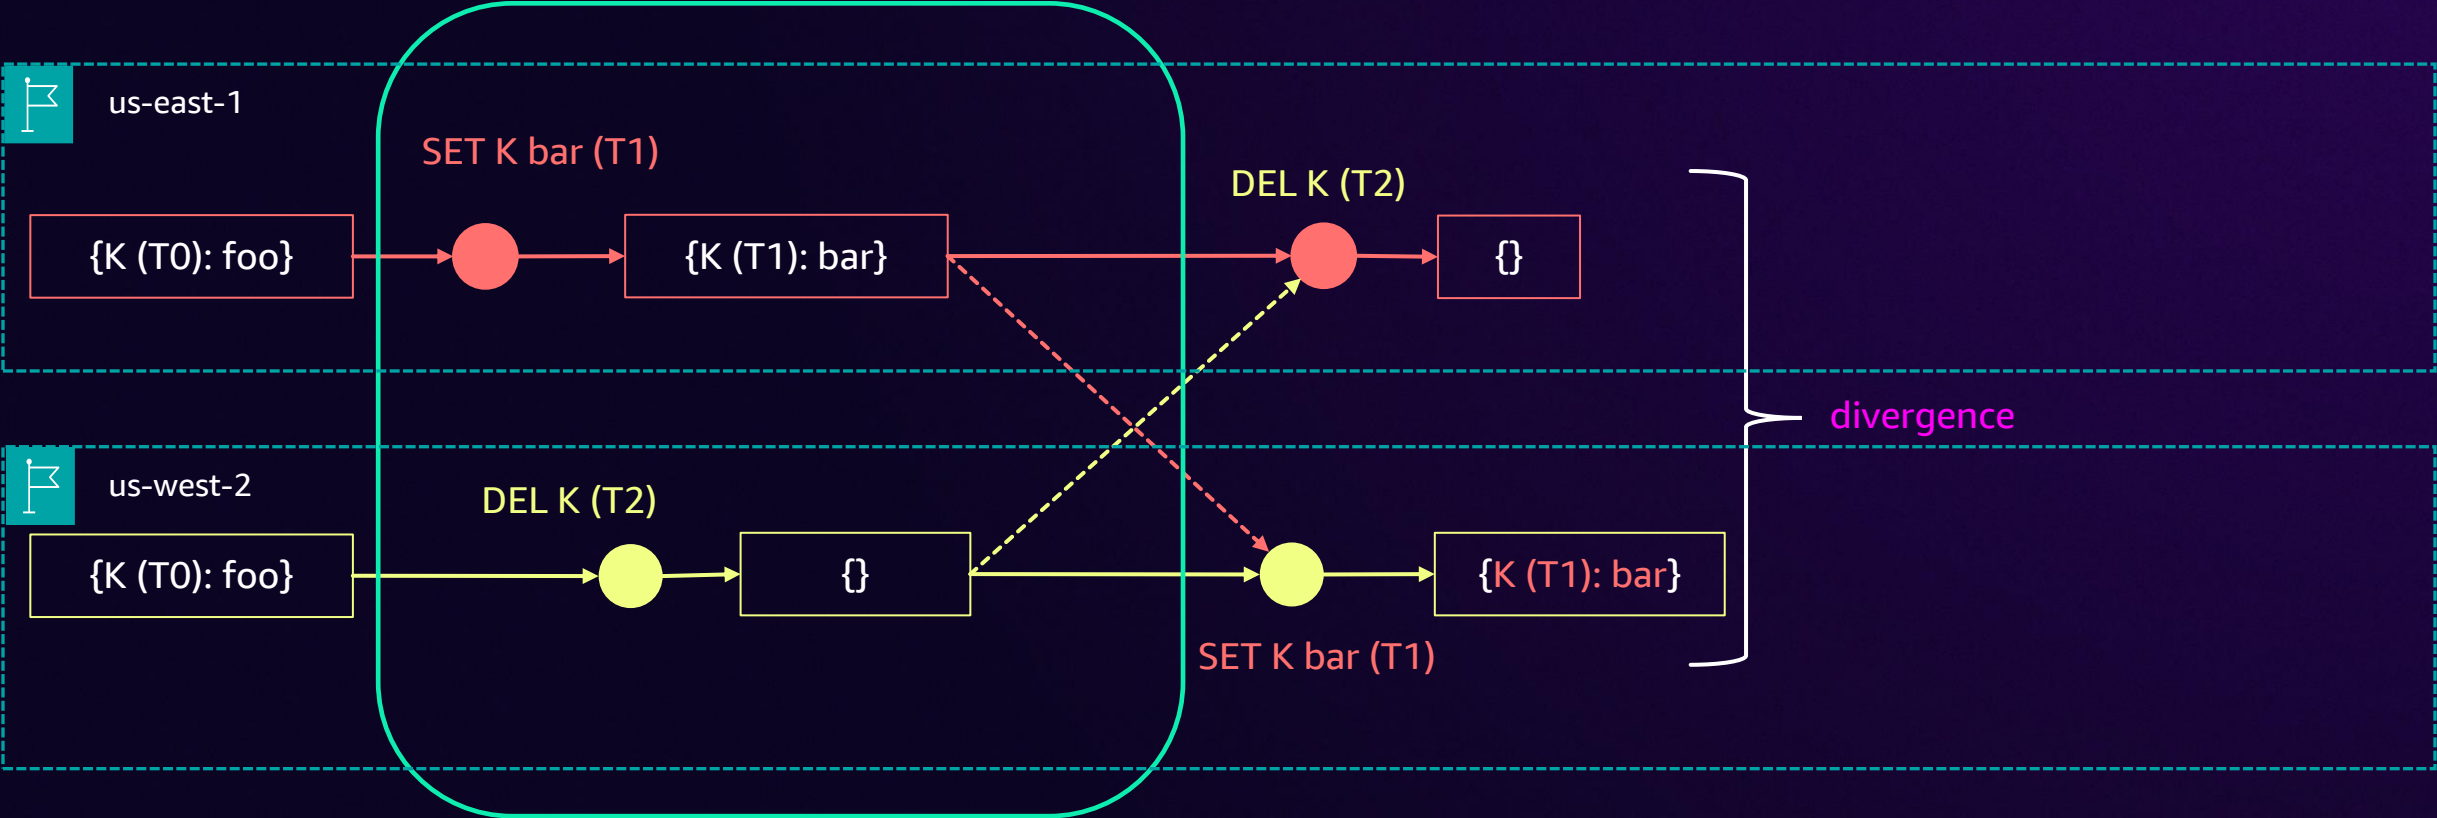
Concurrent writes are merged in a consistent way

LWW – timestamps are tracked at the key and subkey level

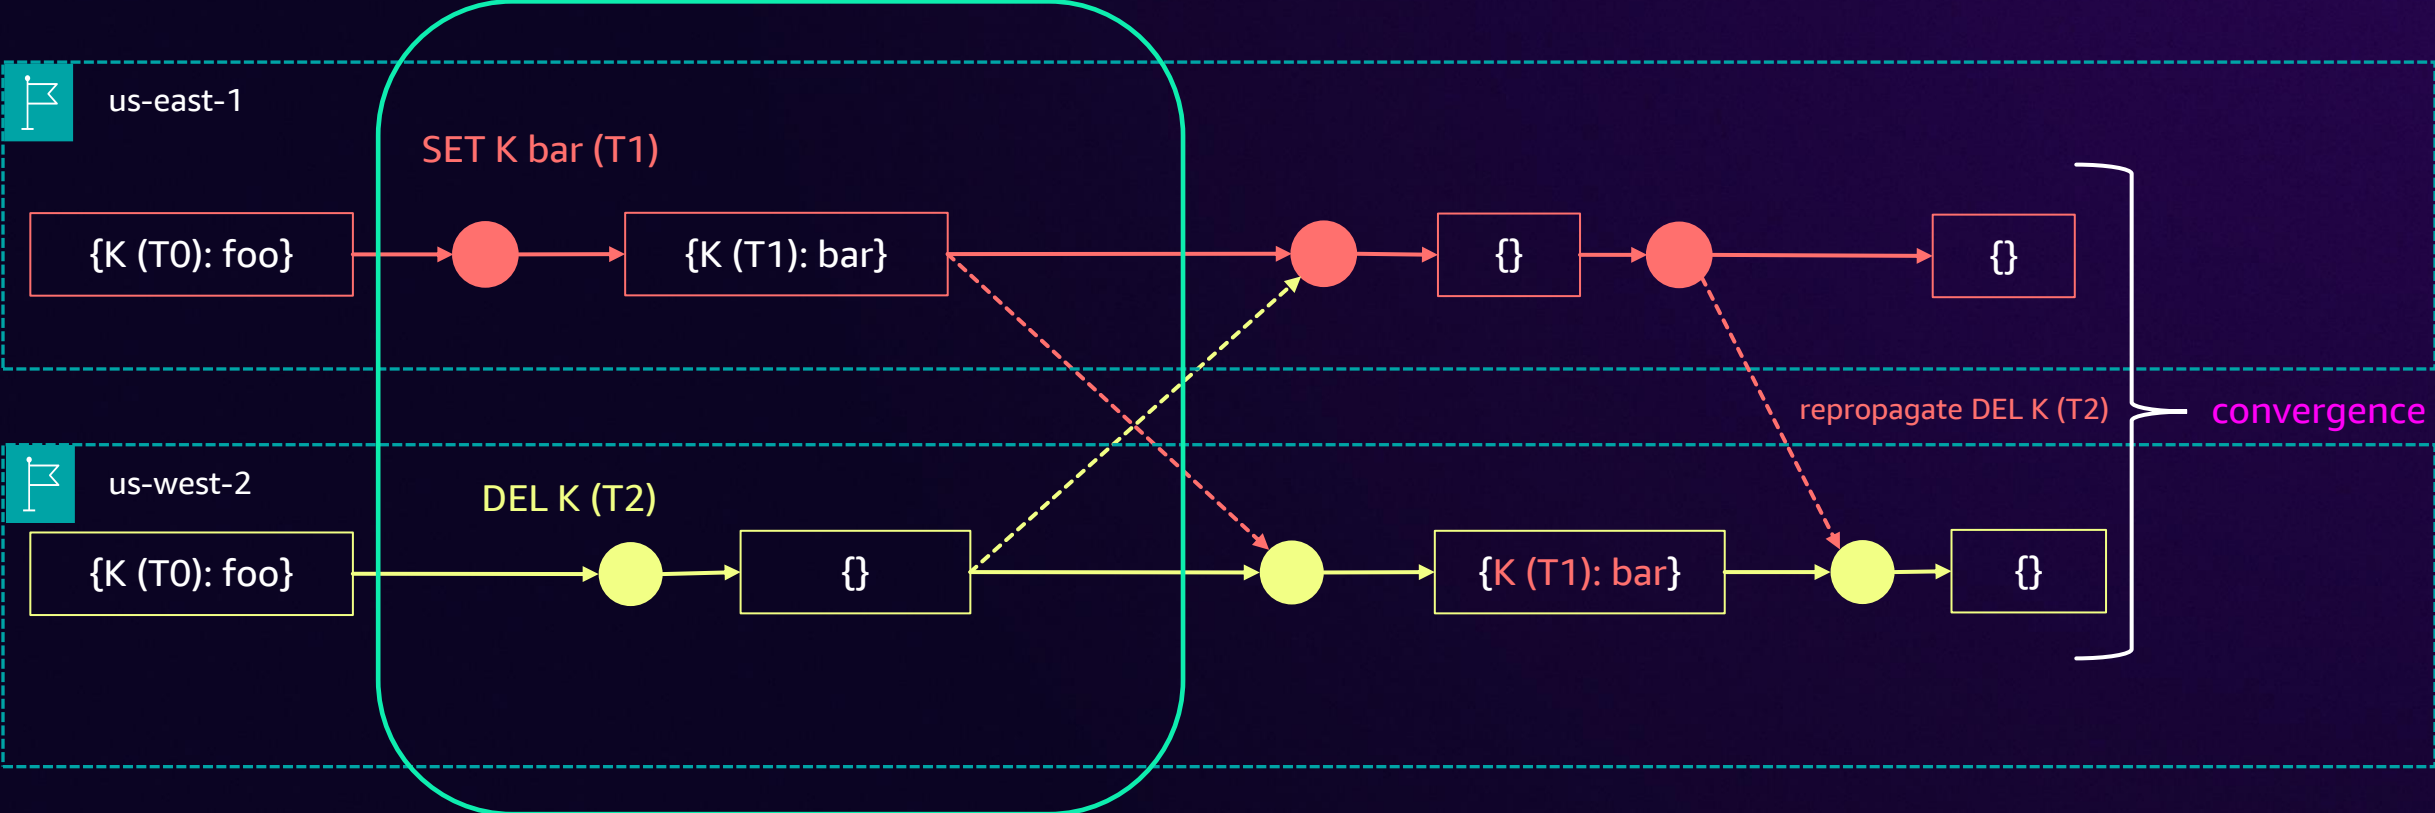


Analyzed with formal modeling technique using the P language

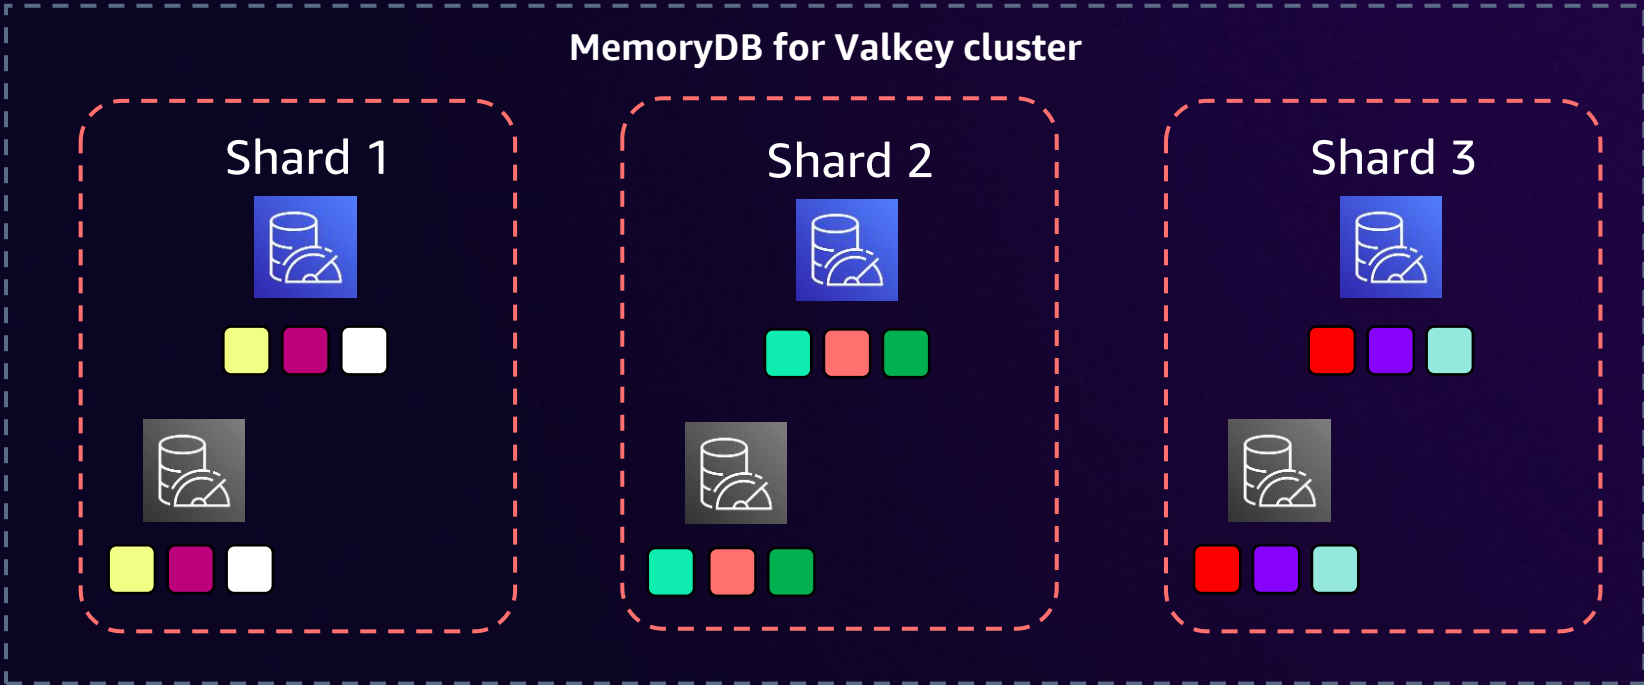
The delete problem






The delete problem



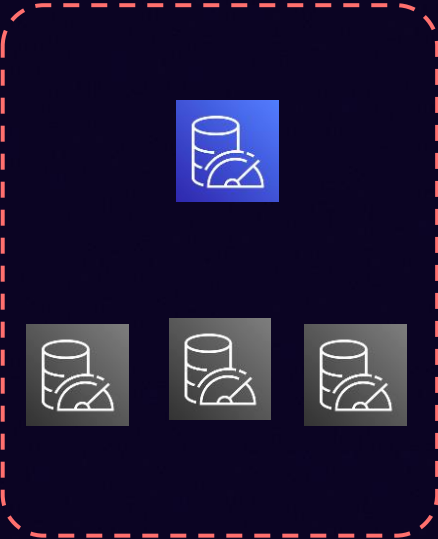
Flexible scaling



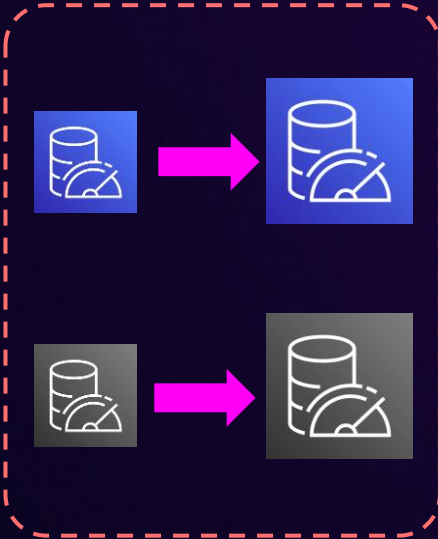
-  Primary node
-  Replica node
-  Slot, a group of data items

Scaling options

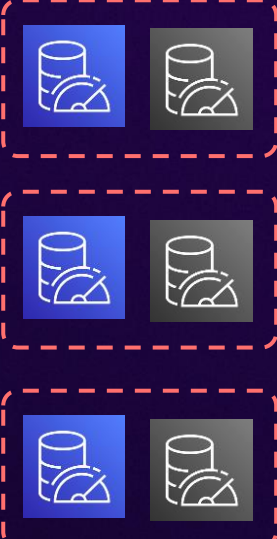
Scaling # replicas



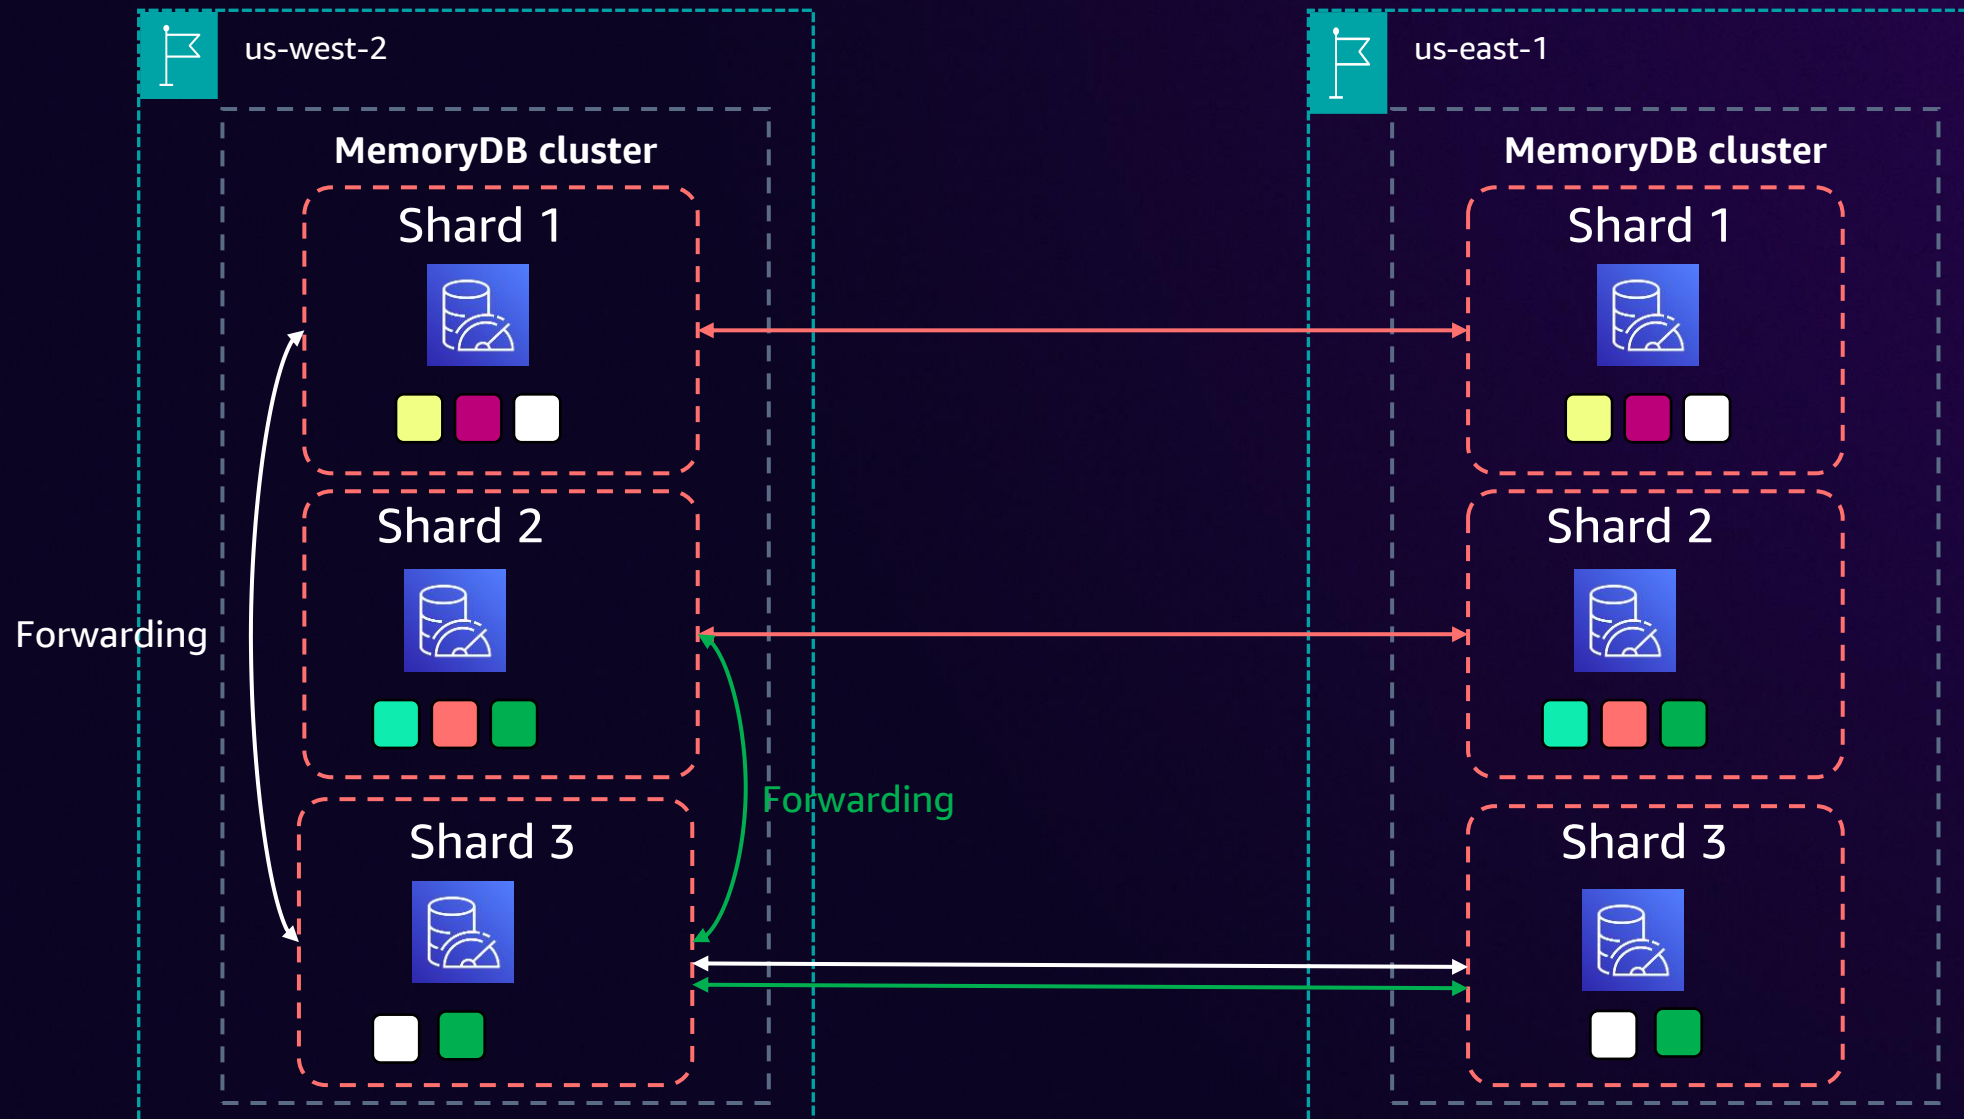
Scaling instance (scale-UP)



Scaling # shards (scale-OUT)



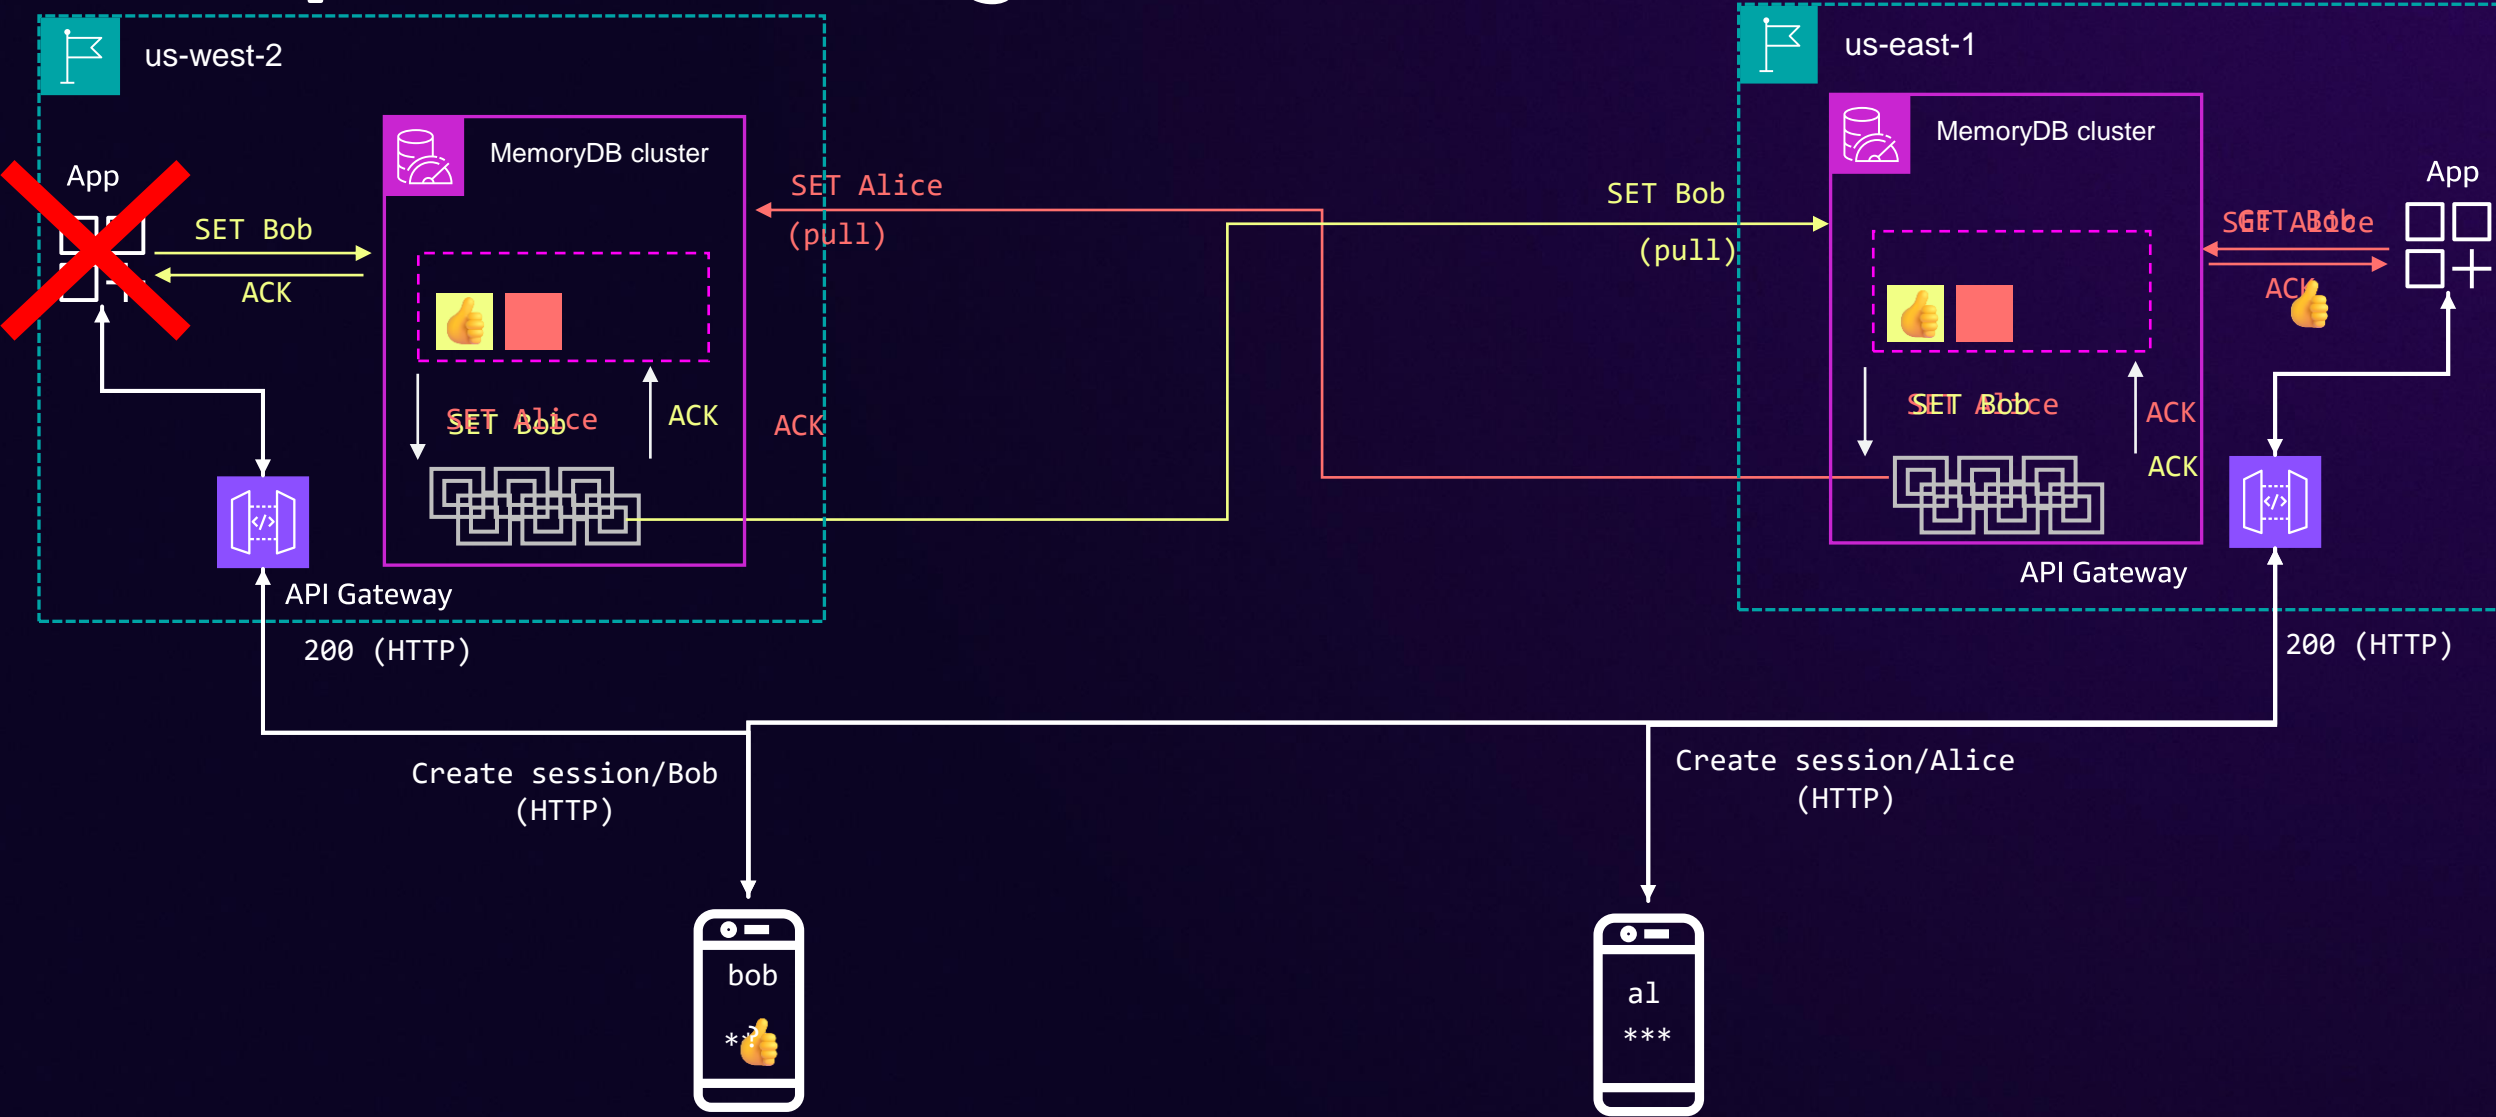
Horizontal scaling MemoryDB Multi-Region



Example – Multi-Region session store



Example – Multi-Region session store



Getting started



Getting started

Multi-Region cluster settings Info

Creation method
Choose from the options for creating your new cluster.

Cluster type

Single-Region cluster
Create a cluster in the current AWS Region.

Multi-Region cluster
Create a multi-Region cluster that spans multiple AWS Regions.

Cluster creation method

Easy create
Use recommended best practice configurations. You can also modify options after you create the cluster.

Create new cluster
Set all of the configuration options for your new cluster.

Restore from snapshots
Use an existing RDB file to restore a cluster.

Configuration
Select one of these options to configure the node type and default configuration of your cluster.

Production
db.r7g.xlarge
26.32 GiB memory
Up to 12.5 Gigabit network performance

Dev/Test
db.r7g.large
13.07 GiB memory
Up to 12.5 Gigabit network performance

Multi-Region cluster info
Configure the name and description of your multi-Region cluster.

Name
The name of the multi-Region cluster.

The name is required, can have up to 40 characters, and must begin with a letter. It should not end with a hyphen or contain two consecutive hyphens. Valid characters: A-Z, a-z, 0-9, and -(hyphen)

Amazon MemoryDB > Clusters

Clusters (1) Info View details View metrics Actions Create cluster

	Name	Description	Status	Node type	AWS Regions	Shards	Total nodes
<input type="radio"/>	ldgnf-test-jguyader	testdescription	Available	db.r6g.large	2 regions	1	-
<input type="radio"/>	test-jguyader-member-2	-	Available	db.r6g.large	eu-central-1	1	-
<input type="radio"/>	test-jguyader-member-1	-	Available	db.r6g.large	us-east-1	1	2

- MemoryDB Multi-Region cluster is available today in 12 Regions
- More Regions coming soon

Thank you!



Please complete the session survey in the mobile app