
AWS
re:Invent
NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

DAT321

How Yahoo cost optimizes their in-memory workloads with AWS

Itay Maoz

General Manager
In-Memory DB Services
AWS

Siva Karuturi

Worldwide Specialist SA
In-Memory Databases
AWS

Maulik Shah

Sr. Principal Software Engineer
Yahoo

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

- Yahoo ad platform journey with Amazon ElastiCache
- Deep dive on Amazon ElastiCache data tiering
- Other AWS in-memory database services
- ElastiCache new features
- Q&A

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Yahoo ad platform journey with Amazon ElastiCache

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Yahoo

Our media, technology, and business platforms connect hundreds of millions of people around the world to the things they love. With apps and sites like Yahoo Sports, Yahoo Finance, Yahoo Fantasy, Yahoo Mail, and more and a powerful, open platform that connects brands and publishers to passionate fans, we're all about putting an exclamation point on what you're into.



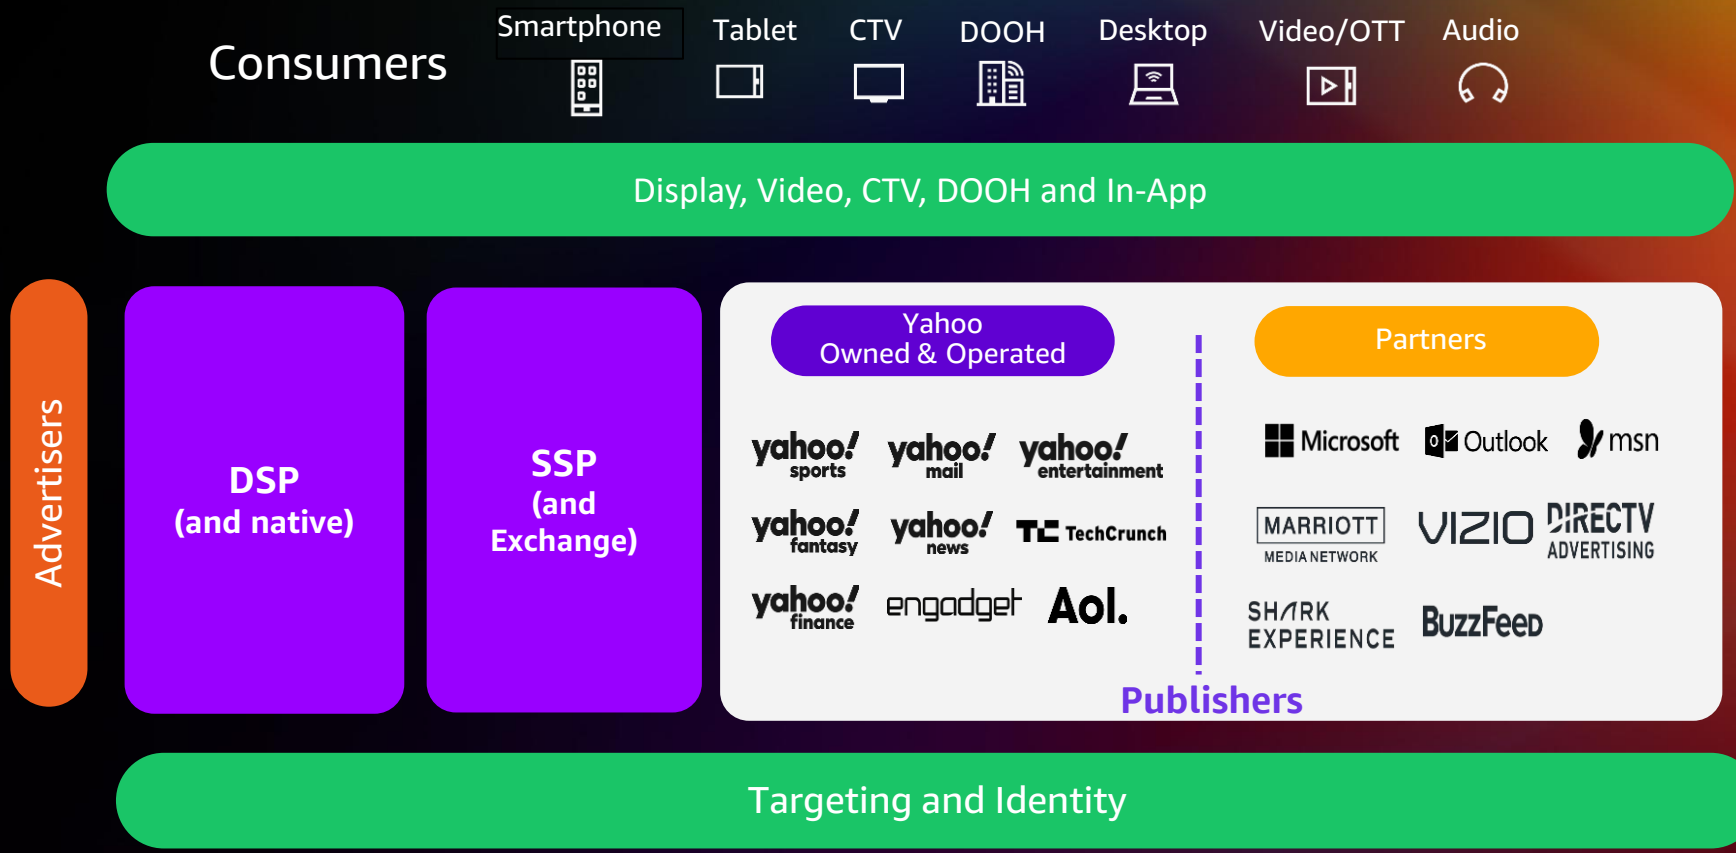






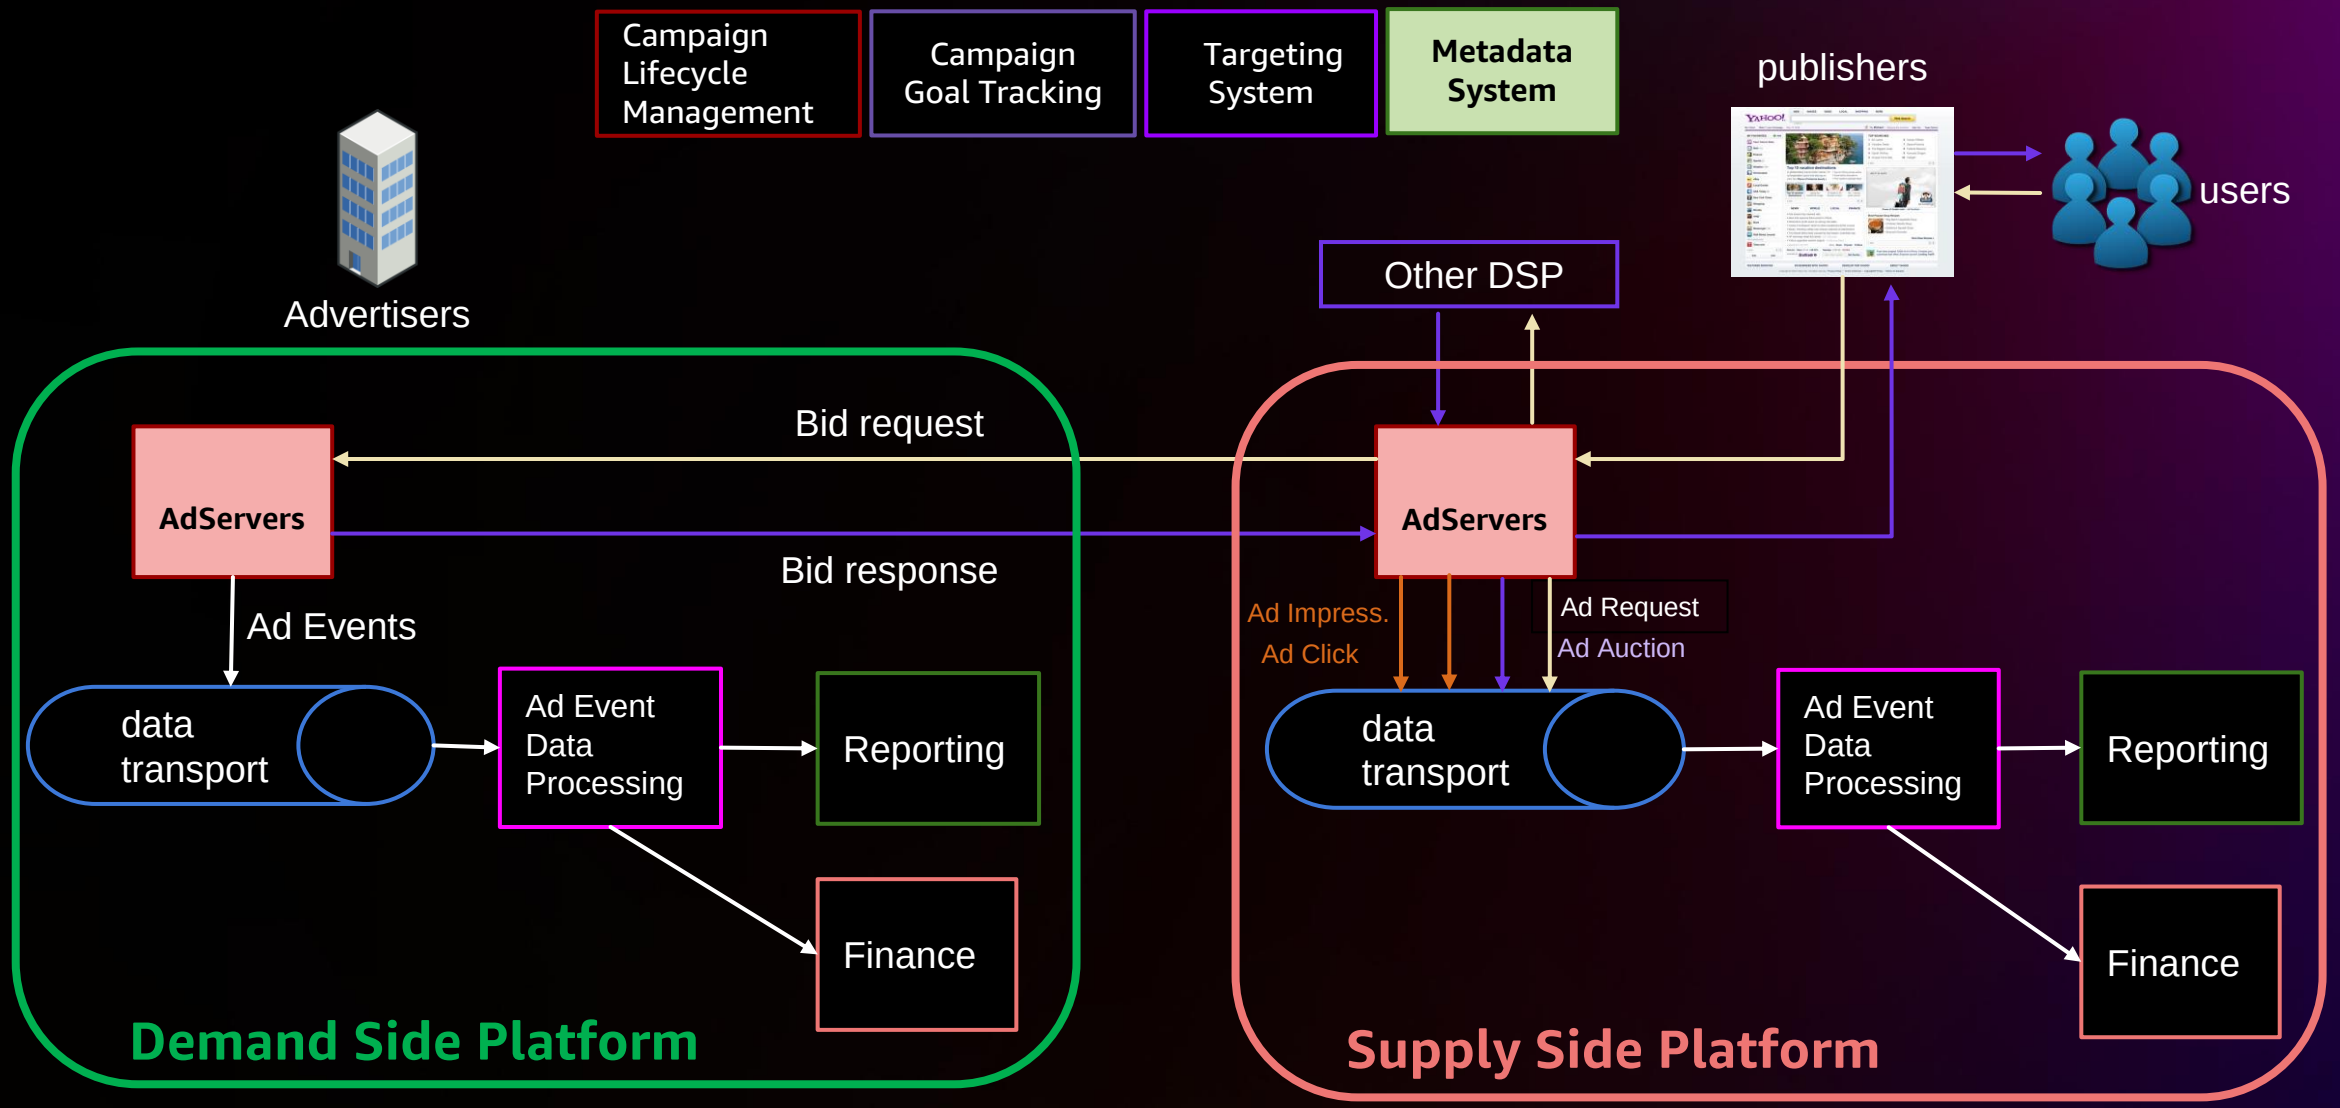
aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.
yahoo!

Advertising ecosystem at Yahoo



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

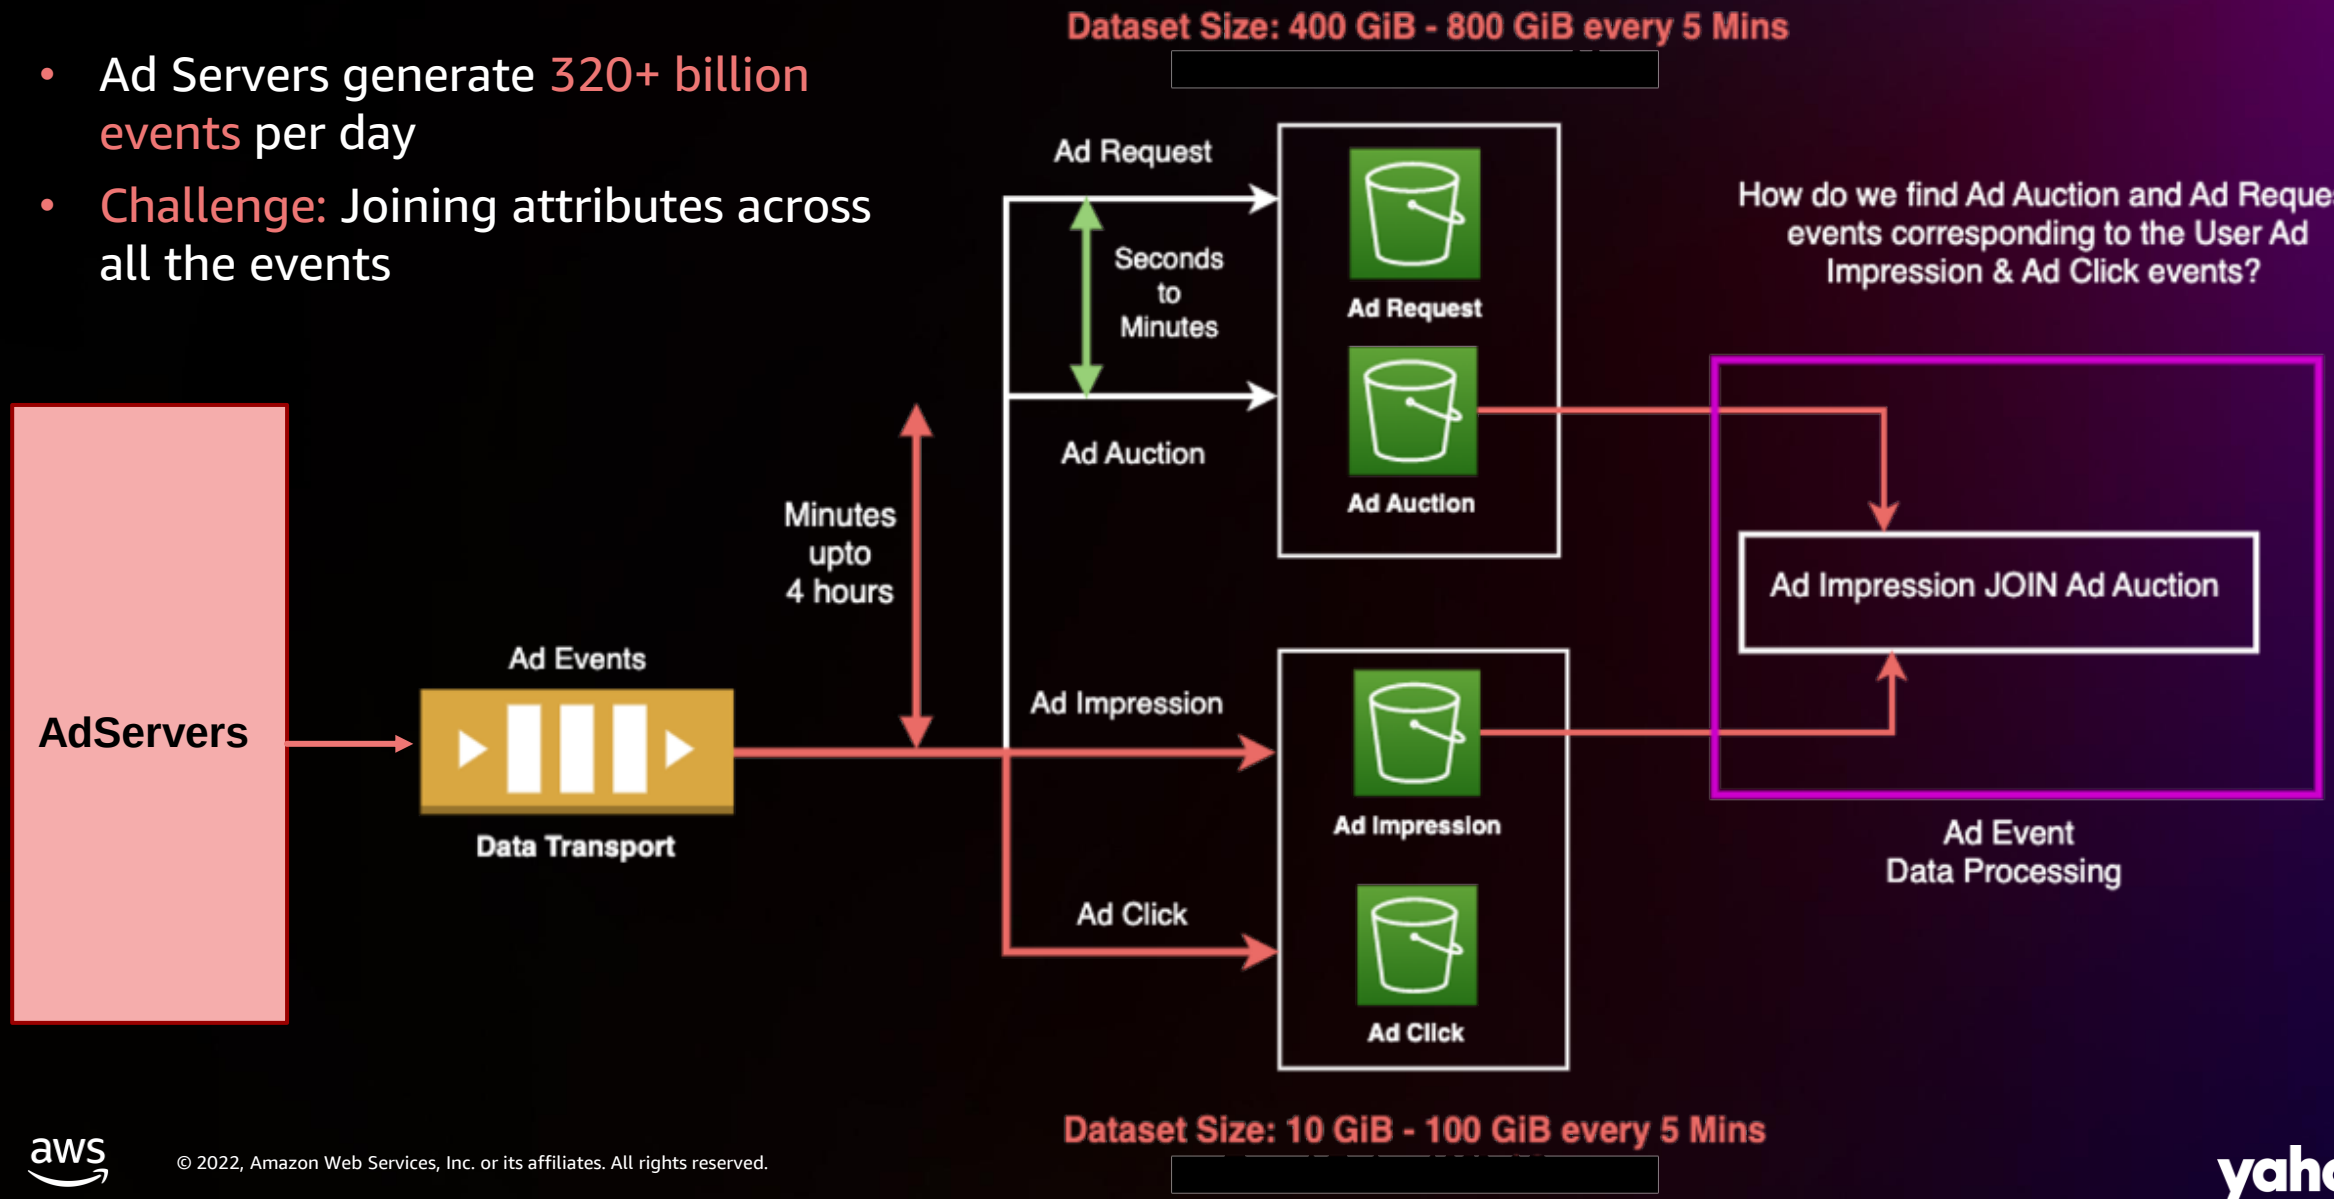
Advertising ecosystem at Yahoo



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Ad Events – Challenge

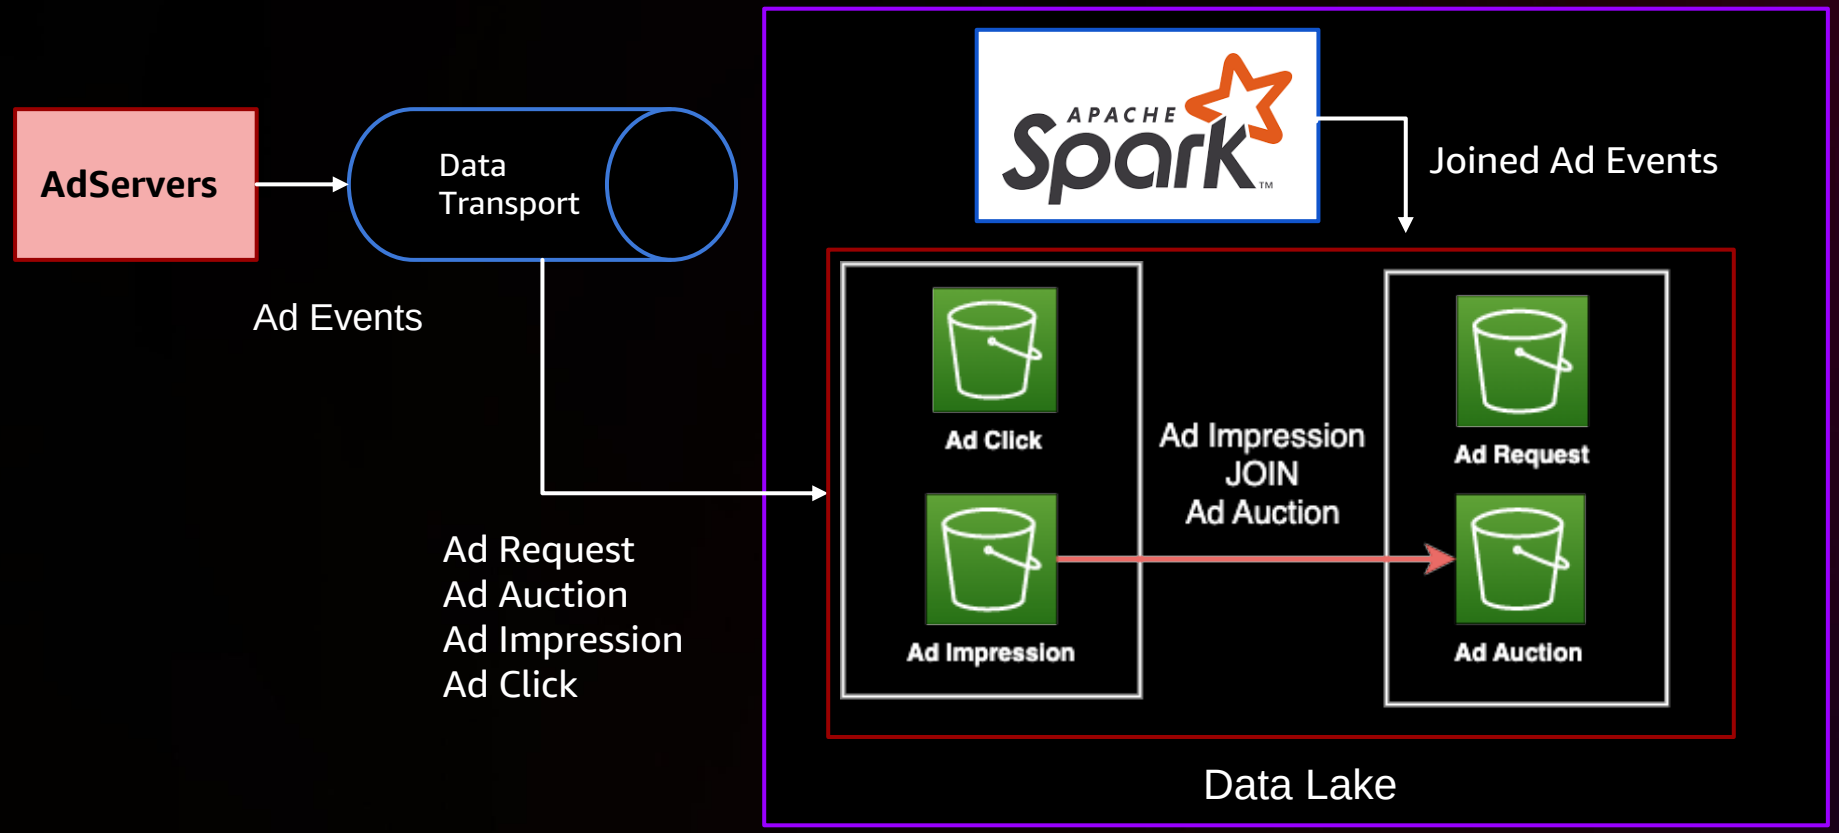
- Ad Servers generate 320+ billion events per day
- Challenge: Joining attributes across all the events



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Ad Events Join operations using MapReduce like frameworks

- Using built-in Join operators in Spark.
- Auction data size is 400GB – 800GB every 5 minutes
- 10GB – 100GB data join with 30+ TB, assuming 4 hour lookback
- Read 48x more data every 5 mins
- Sort Merge join is expensive with our data access pattern



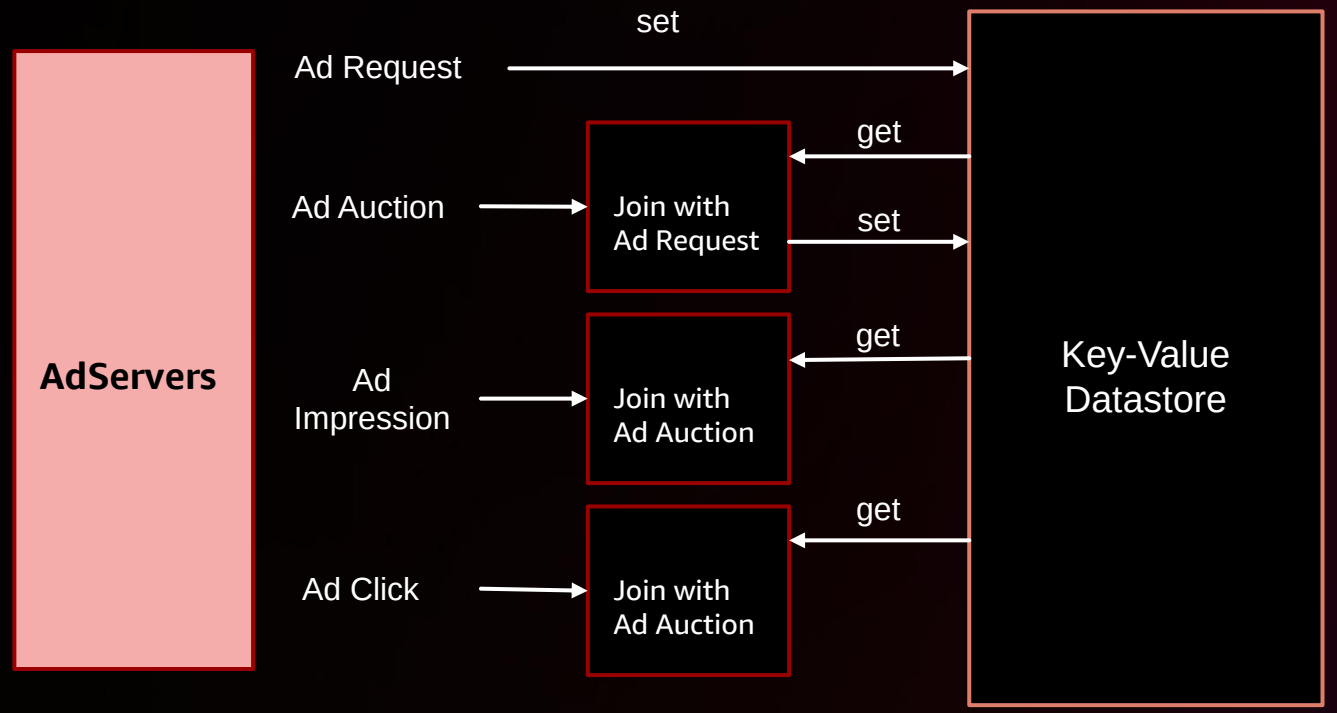
Auction data size
600GB (5 mins)
30 TB (4 hours)

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Ad Events Join operations with KV datastore

Better solution is to store the events in a Distributed Key-Value store

- Avoid repeated scans of the same dataset
- Random lookups are quick; only access required data
- Suitable for batch and streaming workload



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Ad Events – KV solutions considered

Distributed KV store/DB

- Amazon DynamoDB
  - Serverless, managed service
  - Durability/persistence guarantees
  - Cost model and higher cost for our use case



Amazon DynamoDB

- HBase
  - Amazon EMR provides it on HDFS and S3
  - On premises for batch and streaming
  - Maintenance (OS updates, tuning)



- Amazon ElastiCache
  - AWS fully managed service
  - In-memory key-value datastore
  - Uses Redis for streaming workload
  - Better Performance overall – sub-millisec response times

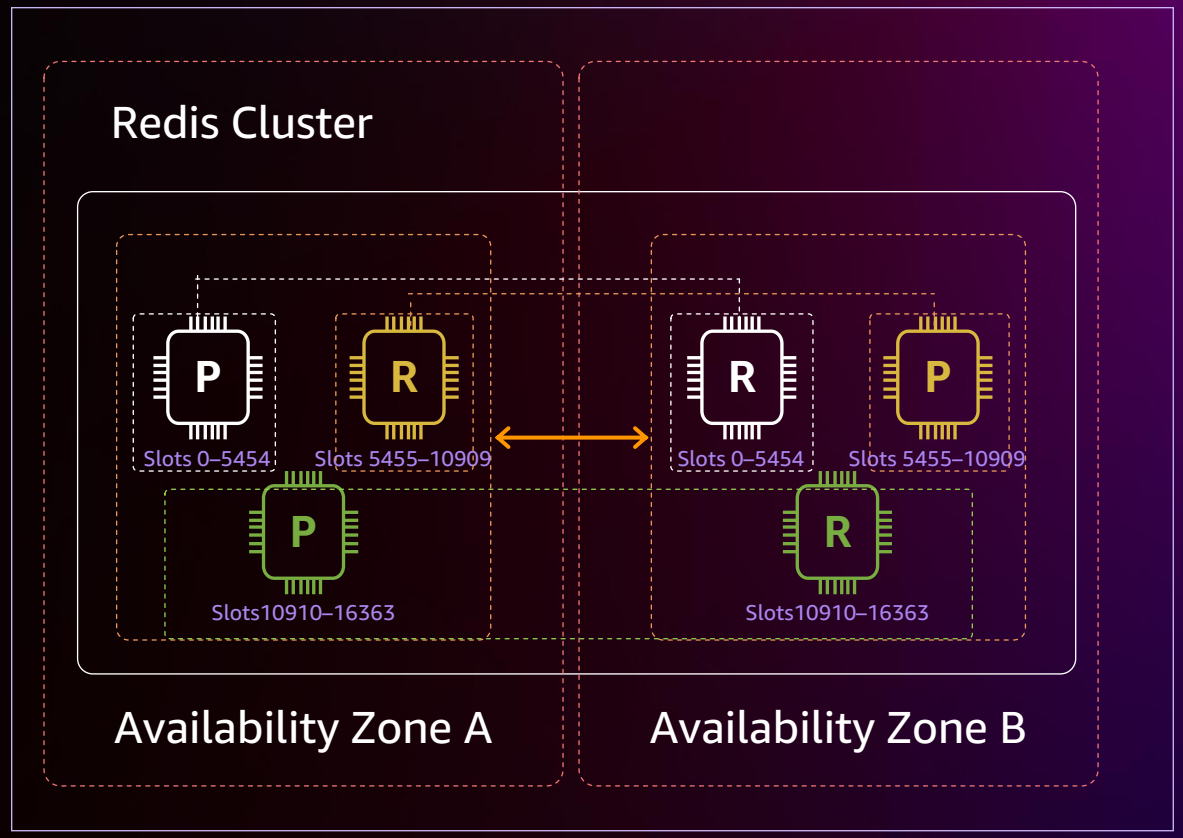


Amazon ElastiCache

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.
yahoo!

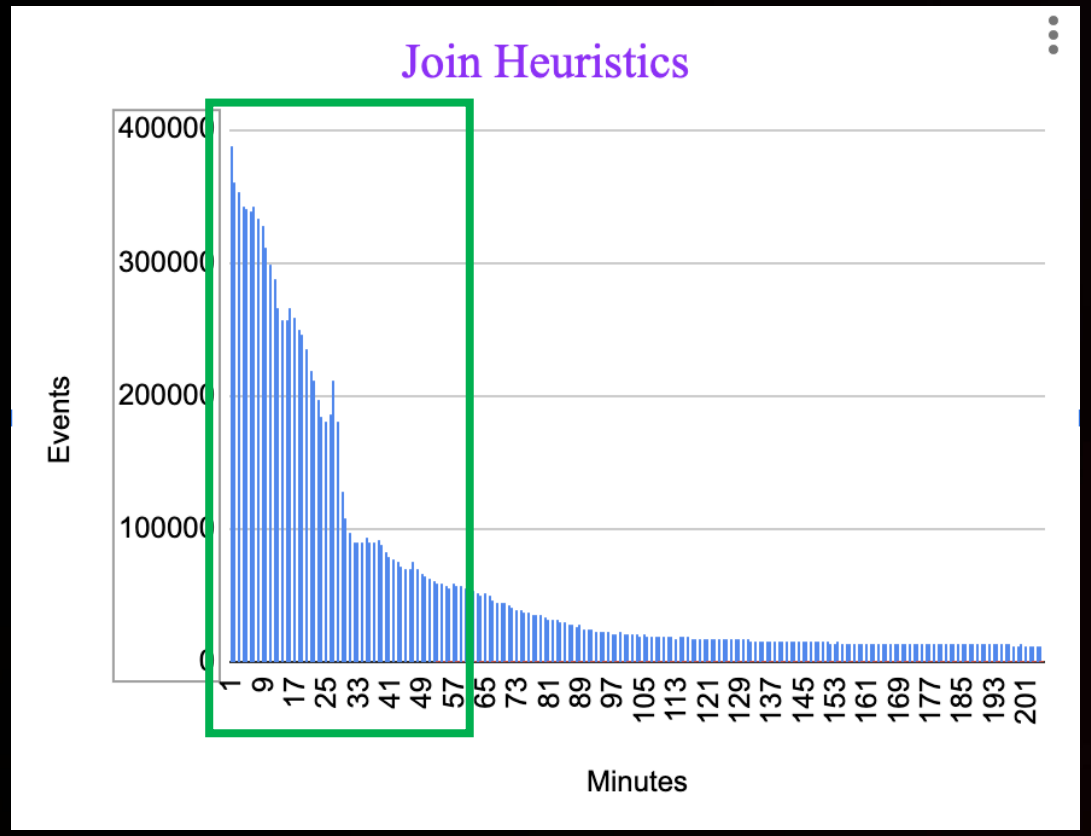
Ad Events processing with Amazon ElastiCache

- Setup
  - Redis 6.2.6, R6g.2xlarge nodes (52.82 GiB RAM)
  - Cluster mode
  - No eviction
  - Lettuce Client
  - Encryption in-flight and rest
- Data Load and Test
  - Loaded 1 hour of Data in Cache
  - Join performance for 5m data
- ElastiCache Cluster sizing
  - Needed approximately 650 nodes to cache 4 hours data (30+TB)



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

ElastiCache – Optimizing cluster footprint



Stats for 4-hour lookback window

Join Window	% Events Joined	% of cache
1 hour	95%	25%
4 hours	5%	75%

- Caching entire dataset results in large cluster footprint
- Started gathering Join Heuristics
- 95% of our joins access only 25% of Cache (1 hour of data)
- 75% of caching cost is for remaining 5% of data access
- Can we offload 3 hours of cold data to disk-based storage without sacrificing Application Performance?

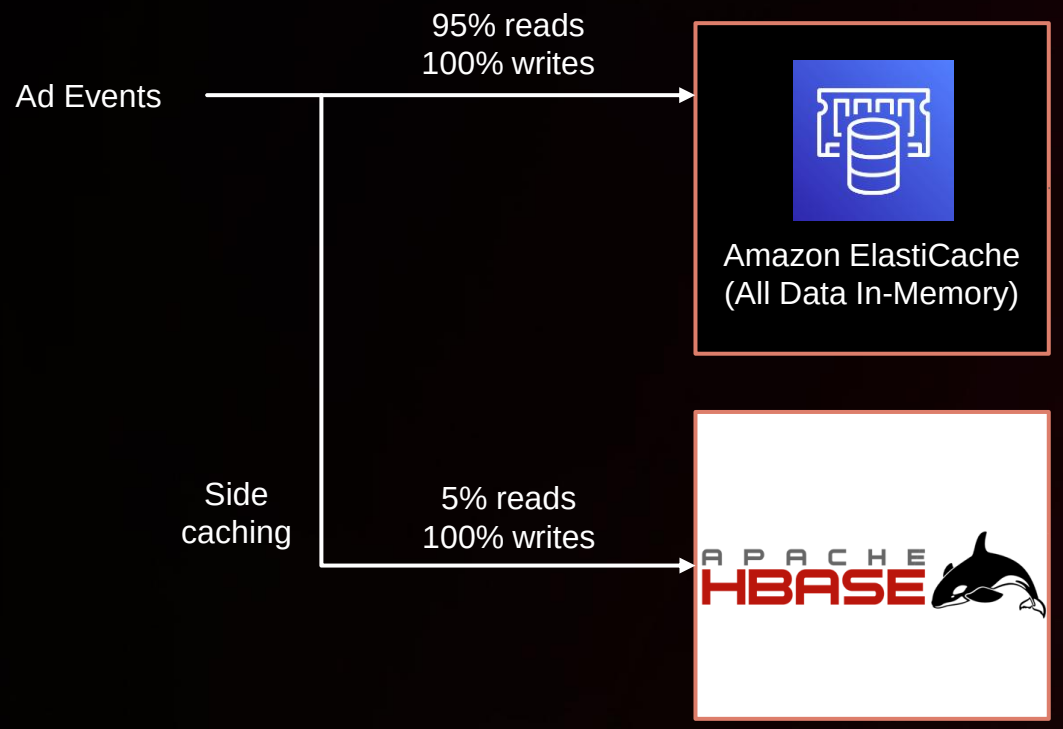
aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.
yahoo!

Ad Events processing

Cost optimization solution considered

Storage options

- Side caching Amazon ElastiCache with Hbase
- Extra writes, maintenance, complex read logic



25% of Overall Data

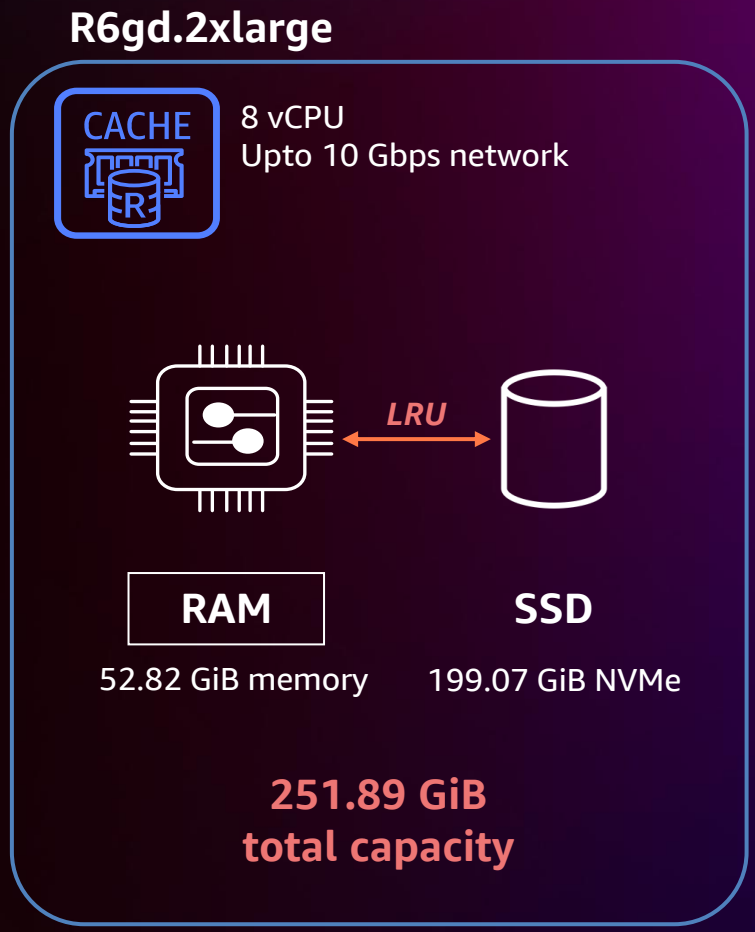
All Data

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Ad Events processing

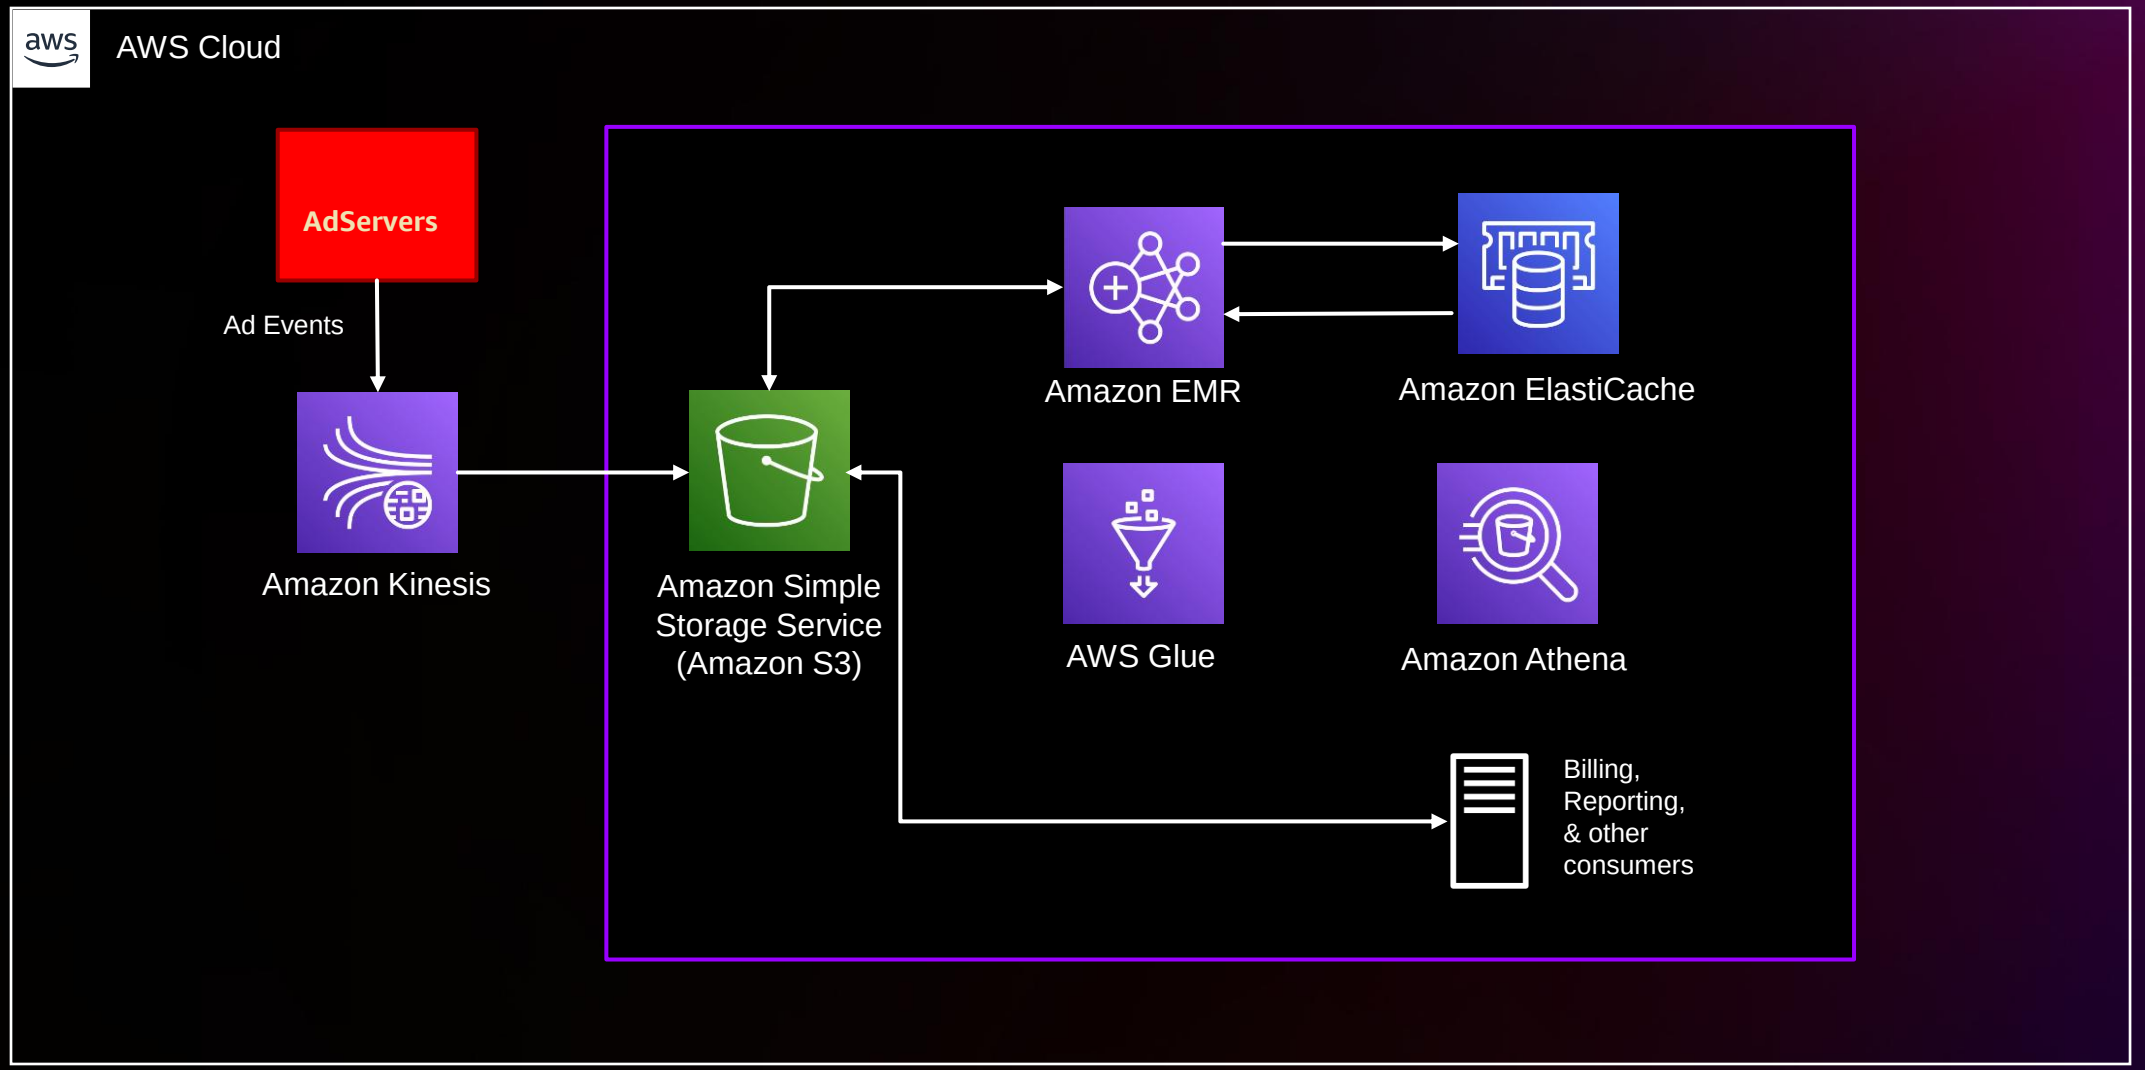
Amazon ElastiCache with data tiering

- R6g**d**.2xlarge, 52.82GiB RAM, 199.07GiB SSD
- Cluster 150 nodes, 4x reduction in resources
  650 nodes to 150 nodes
- 50% reduction in Cost
- Allowed us to easily and cost effectively scale to 10s of TiB capacity
- No performance impact to data processing workload!



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.
yahoo!

Ad Events data processing with AWS services

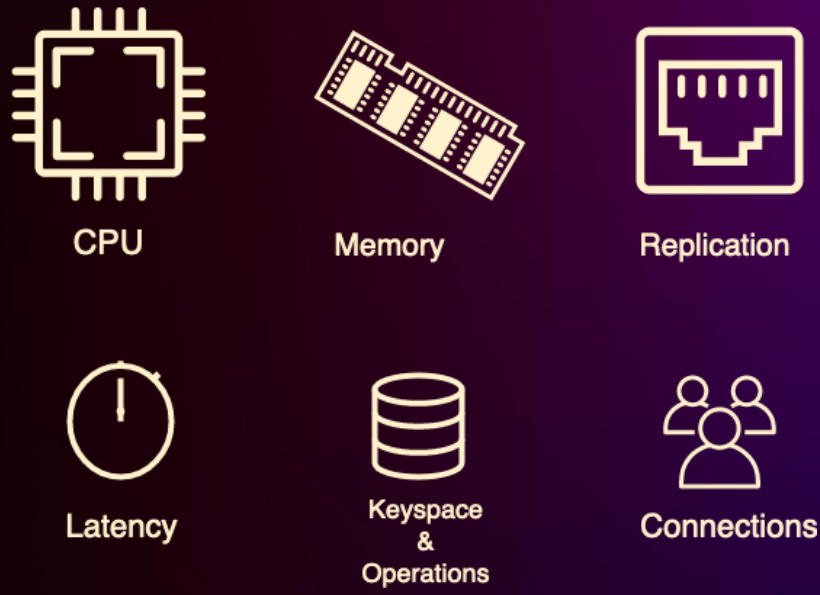


© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Production metrics

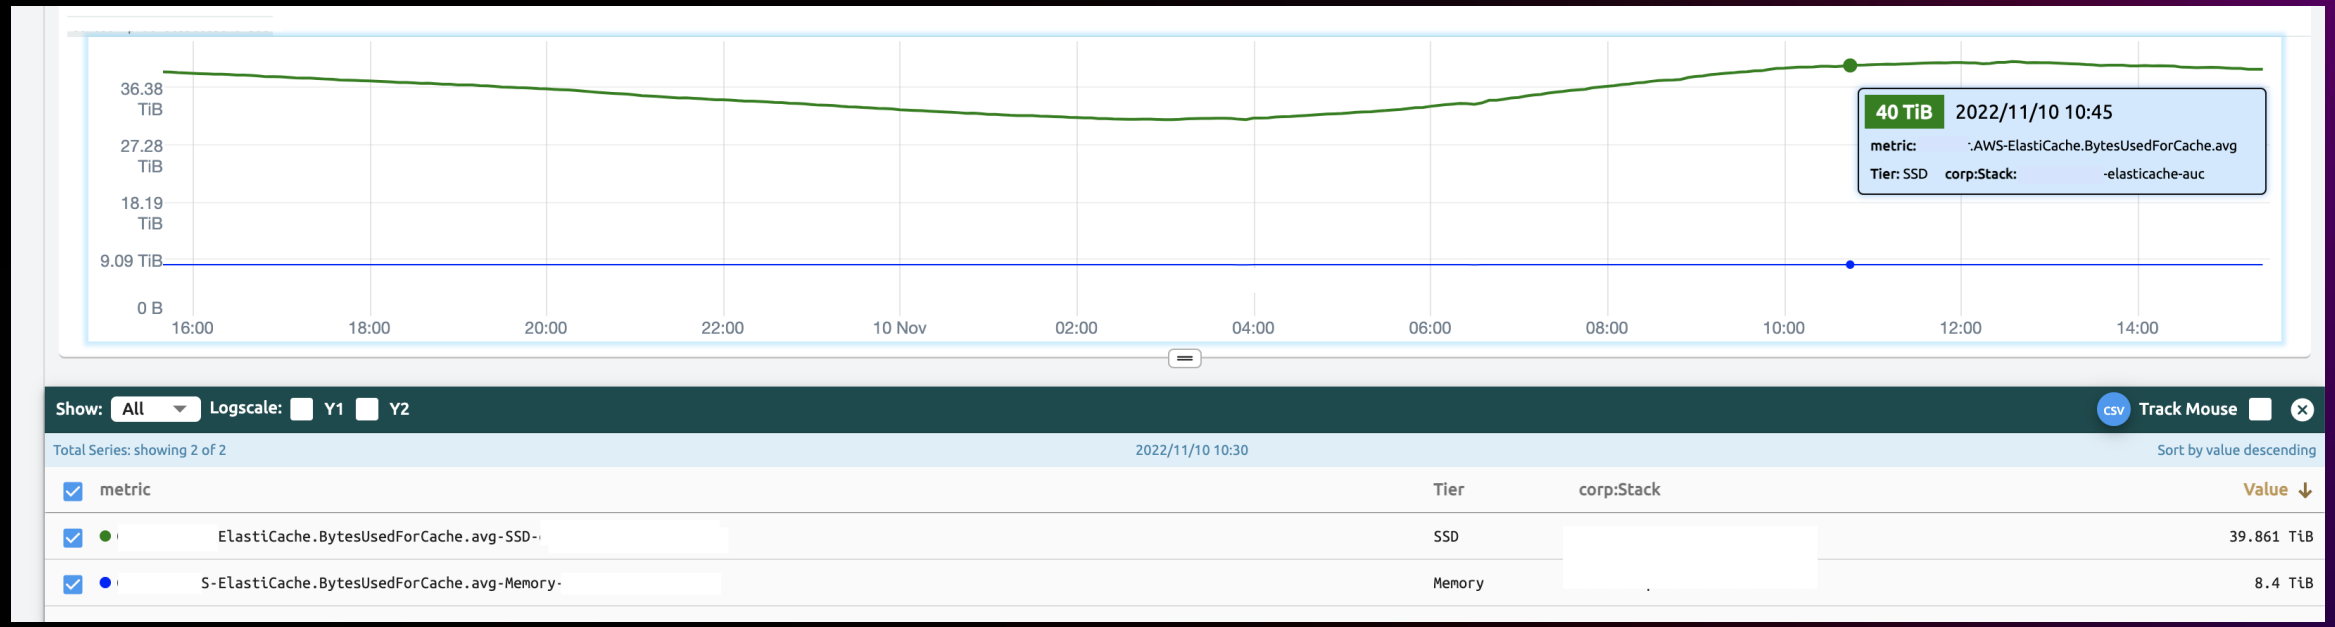
- Alerting and scale-out operations
  - EngineCPUUtilization
  - BytesUsedForCache
  - SwapUsage

- Monitoring peak traffic / Backlog processing
  - CurrConnections, NewConnections
  - BytesWrittenToDisk, BytesReadFromDisk
  - GetTypeCmdsLatency, SetTypeCmdsLatency
  - GetTypeCmds, SetTypeCmds



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.
yahoo!

ElastiCache – Storage tiers (SSD and memory)



- Auction data size is 400GB – 800GB every 5 minutes
- Total cluster size ranges between ~20 TB (400*12*4hrs.) and ~40TB (800*12*hrs.)

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.
yahoo!

Lessons learned

- Have persistent connections
- Share connections between threads
- Encryption in-transit can impact performance
- Implement exponential backoff with retries

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.
yahoo!

Future work

- Further cost reductions by payload optimizations
- Implement auto scaling solution with data tiering
- Implement streaming workload with Amazon ElastiCache

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.
yahoo!

Yahoo ad platform – Summary

- Advertising ecosystem overview
- Request flow and Ad events
- Data processing join use case with Amazon ElastiCache
- Cost savings with ElastiCache data tiering
- Metrics and lessons learned

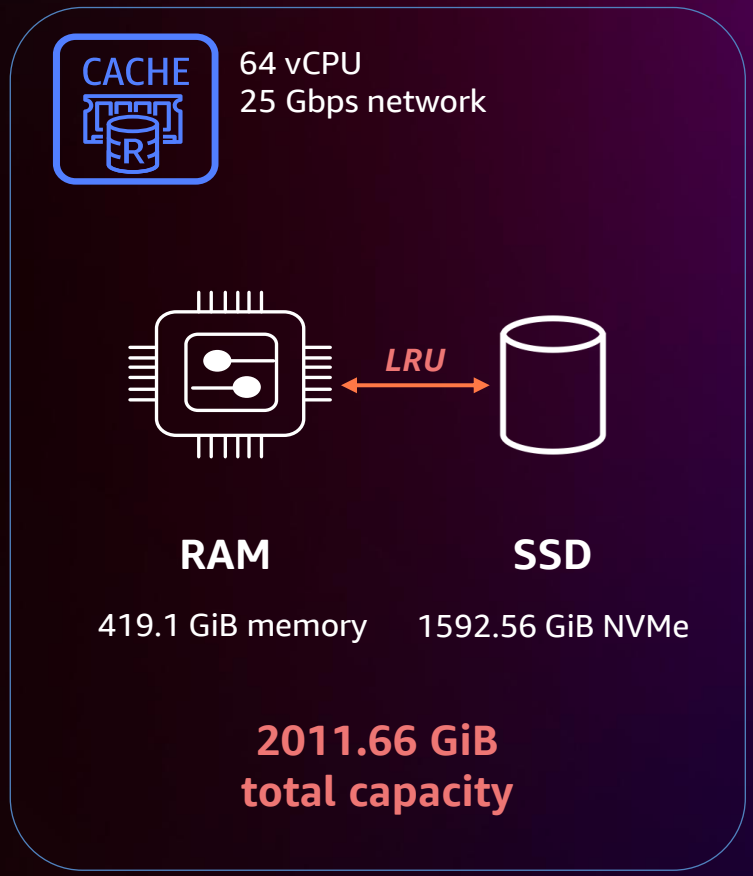
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data tiering with Amazon ElastiCache for Redis

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon ElastiCache data tiering with R6gd nodes

- Expands the per-node data capacity by moving least recently used (LRU) items from memory to SSD when the memory usage in Redis reaches “maxmemory” limit
- Ideal for workloads that access up to 20% of their dataset regularly
- Save over 60% per GB at full capacity utilization
- No application changes required; minimal performance impact
- Clusters scale up to 1 PiB in size



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Cost-effectively scale with R6gd

with data tiering

	R6g.xlarge	R6gd.xlarge
vCPU	4	4
RAM (GiB)	26.32	26.32
SSD (GiB)	0	99.33
Total capacity (GiB)	26.32	**125.65**
Hourly node price*	$0.411	$0.781
Price per GiB-hour*	$0.016	$0.006

4.8x capacity

over 60% savings!

Offered in xlarge, 2xlarge, 4xlarge, 8xlarge, 12xlarge, and 16xlarge sizes

*Displayed pricing for on-demand nodes in the US East (N. Virginia) Region

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data tiering node specs

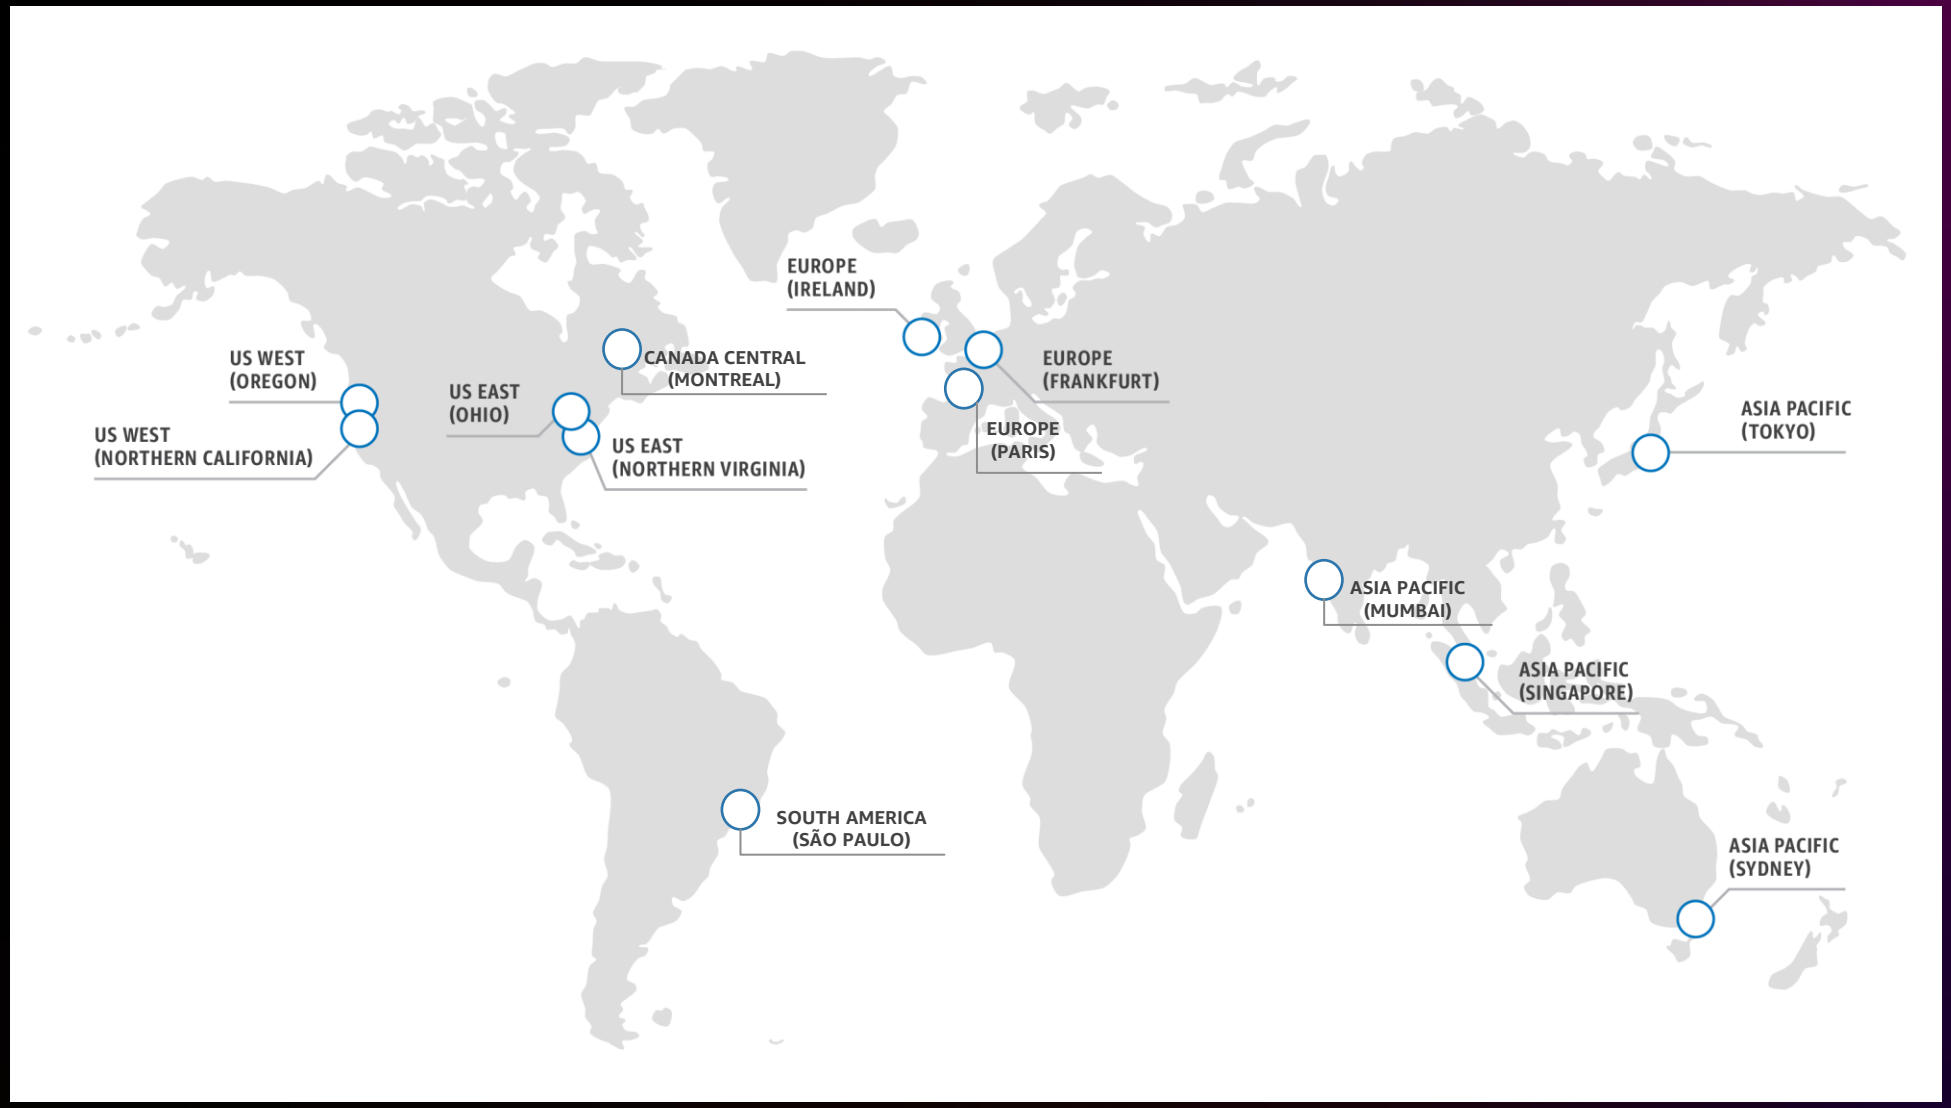
Node type	vCPU	Network bandwidth	RAM (GiB)	SSD (GiB)	Total capacity (GiB)	Node hourly price*	Price* per GiB-hour
R6gd.xlarge	4	Up to 10 Gbps	26.32	99.33	125.65	$0.781	$0.006
R6gd.2xlarge	8	Up to 10 Gbps	52.82	199.07	251.89	$1.560	$0.006
R6gd.4xlarge	16	Up to 10 Gbps	105.81	398.14	503.95	$3.120	$0.006
R6gd.8xlarge	32	12 Gbps	209.55	796.28	1005.83	$6.240	$0.006
R6gd.12xlarge	48	20 Gbps	317.77	1194.42	1512.19	$9.358	$0.006
R6gd.16xlarge	64	25 Gbps	419.10	1592.56	2011.66	$12.477	$0.006

*Displayed pricing for on-demand nodes in the US East (N. Virginia) region

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

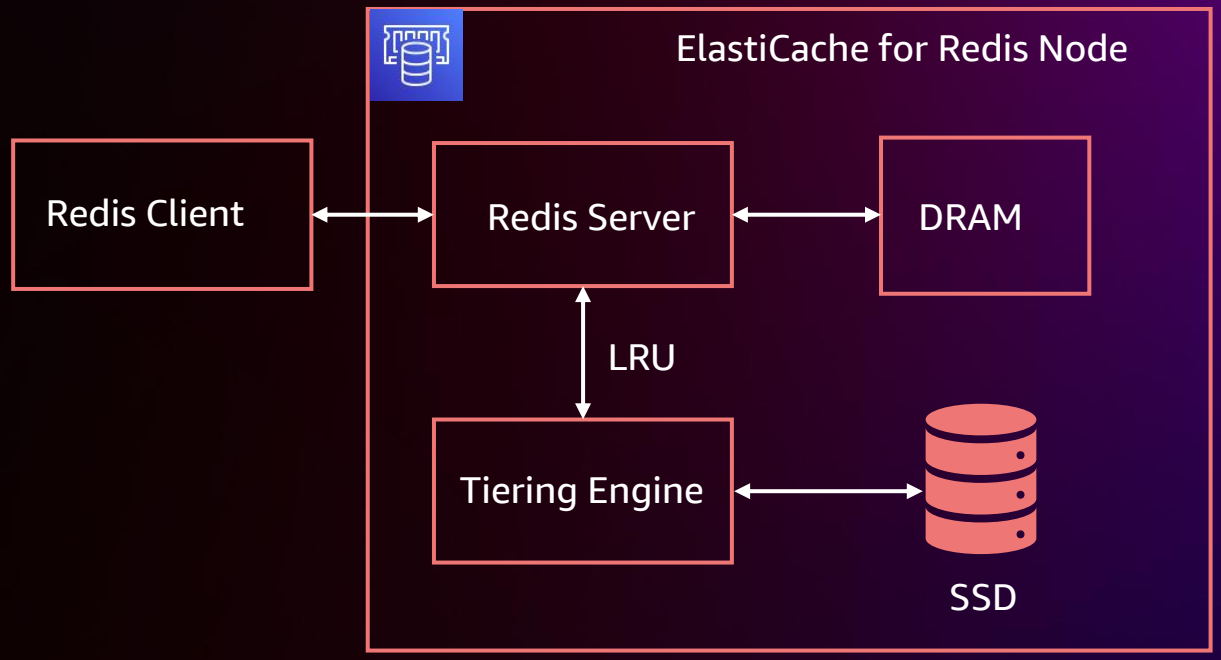
ElastiCache for Redis R6gd Region availability



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

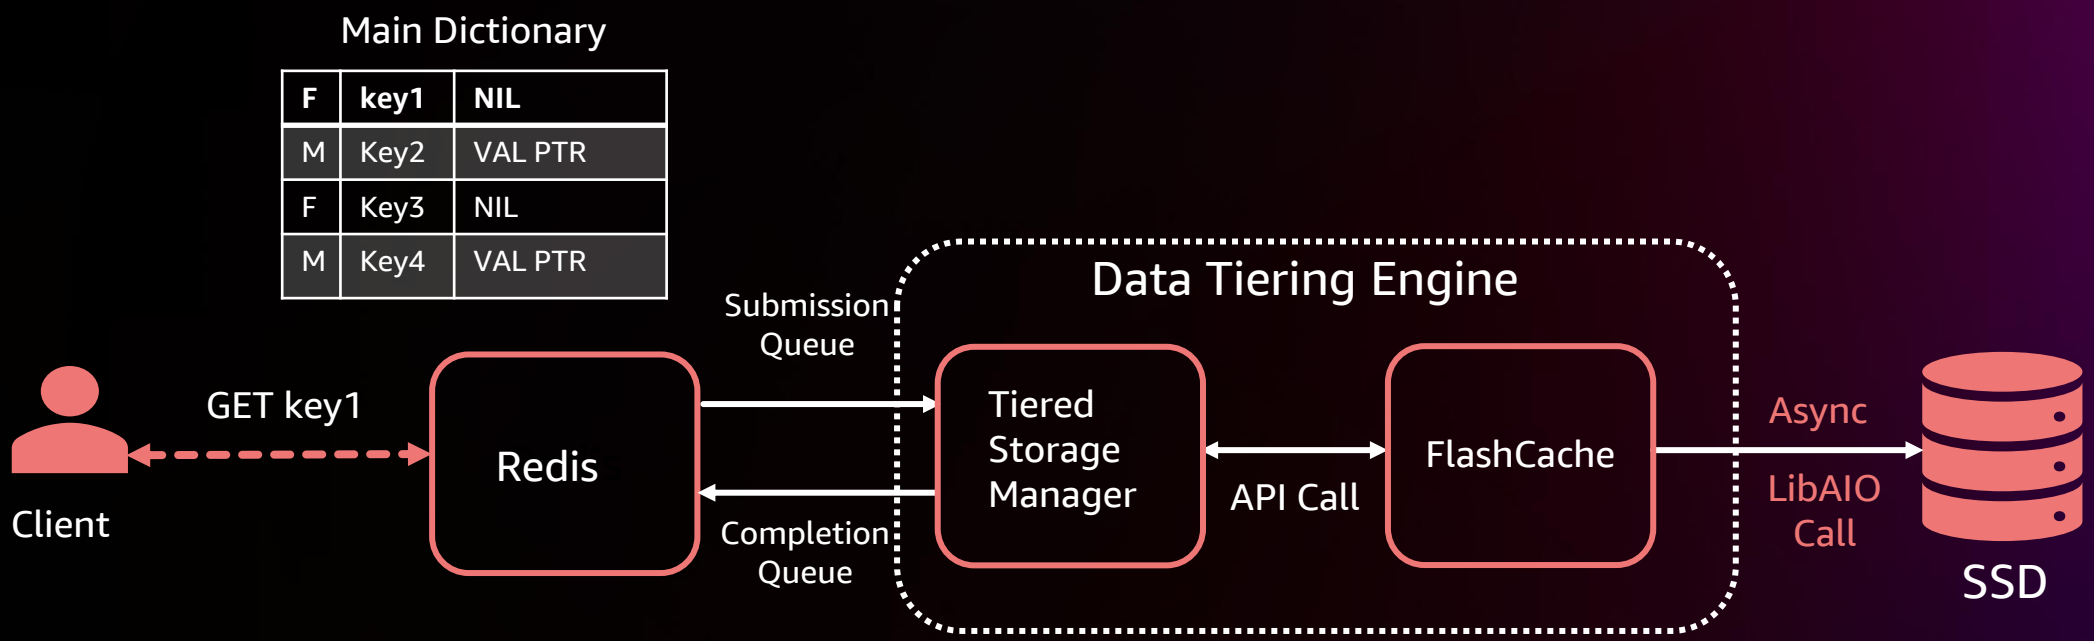
How data tiering works

- Fixed, 16-byte overhead per key
- Operates atomically at item level
- Keys are always stored in memory
- Asynchronous communication between Redis engine and tiering engine
- No Redis client changes required



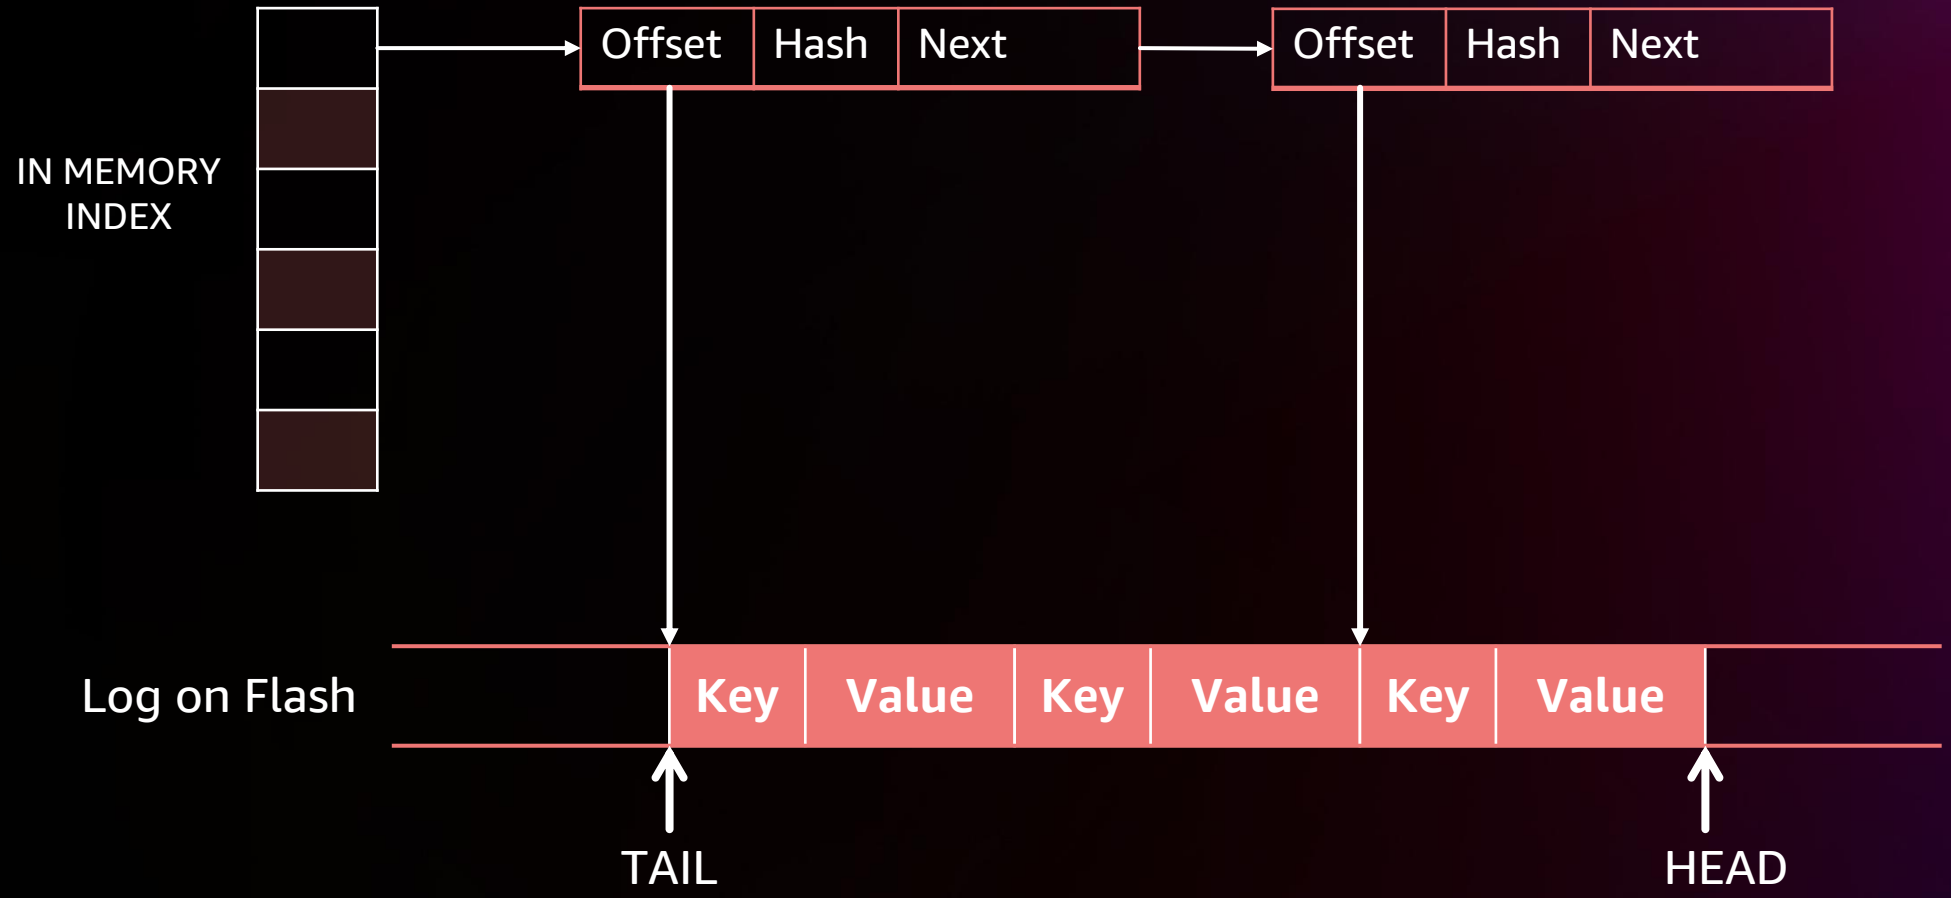
aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data tiering – High-level flow



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

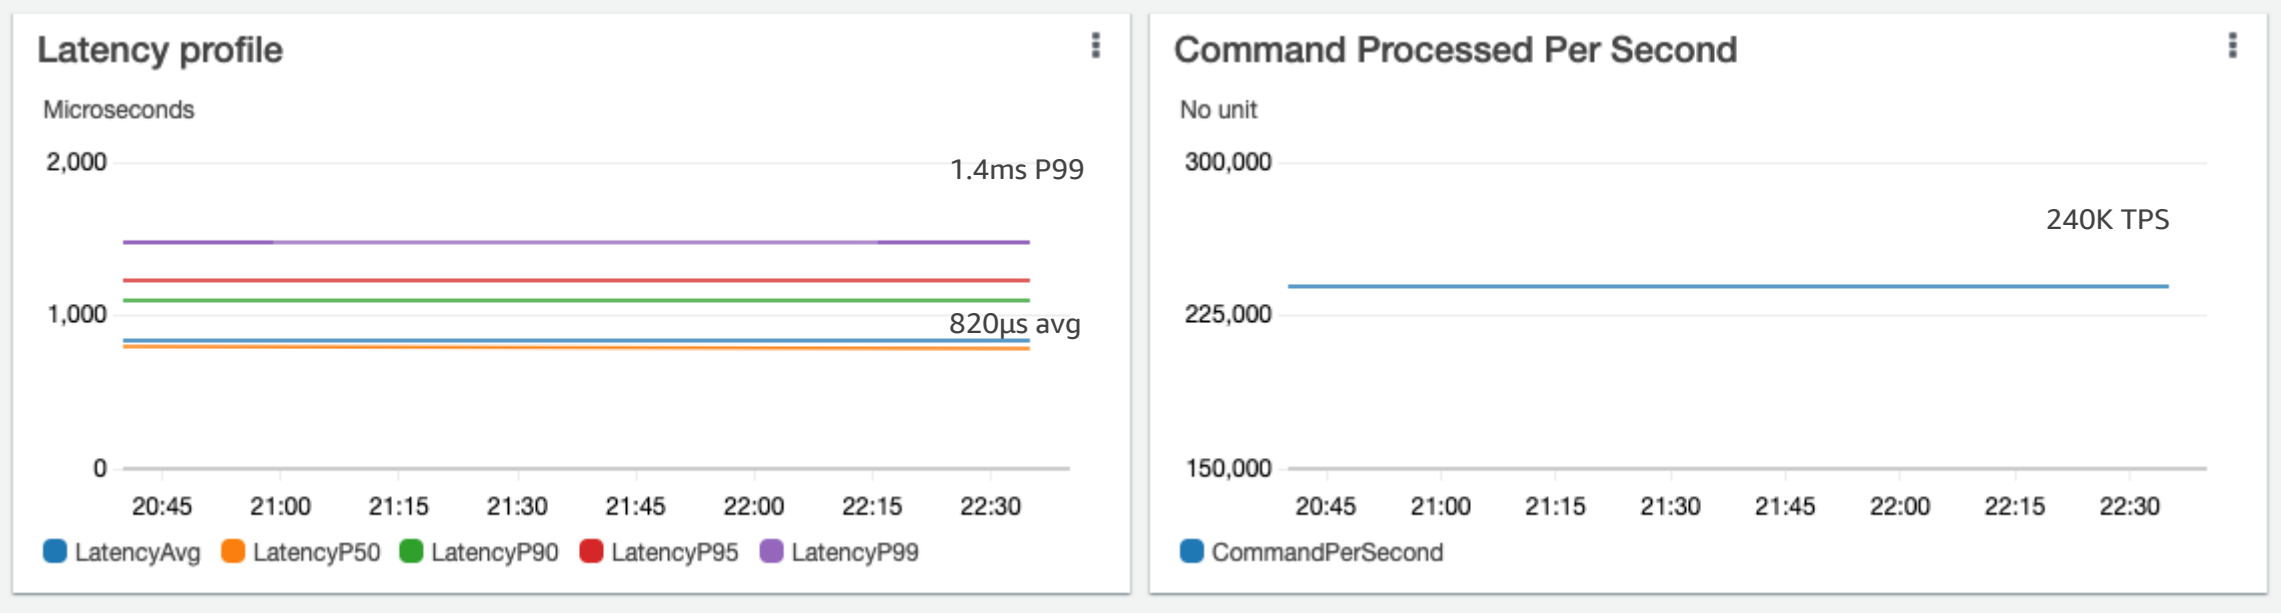
FlashCache storage engine



- FlashCache tracks the hash value of the on-flash keys in its Index Table
- Append-only log Stores the actual Key name along with the value

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Performance review



- R6xd.2xlarge
- 396M unique keys
- 16 byte keys, 500 byte string values
- 200 client connections
- 4:1 get:set ratio
- 10% of workload hits SSD tier

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Monitoring with Amazon CloudWatch

4 new metrics

- NumItemsWrittenToDisk
- NumItemsReadFromDisk
- BytesWrittenToDisk
- BytesReadFromDisk

New metric dimensions

- CurrItems
  - Tier=Memory
  - Tier=SSD
- BytesUsedForCache
  - Tier=Memory
  - Tier=SSD

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Monitoring with Amazon CloudWatch



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data tiering in a nutshell

- Use R6gd with data tiering to cost-effectively scale large-capacity Redis workloads
- Ideal for applications that access up to 20% of data regularly
- Minimal performance impact; transparent to applications
- Available today in 13 AWS Regions

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS in-memory database services



Amazon MemoryDB for Redis

(Now supports data tiering)

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon MemoryDB for Redis

REDIS COMPATIBILITY WITH MULTI-AZ DURABILITY

Ultra-fast performance

Lowest latency database offered by AWS

Redis compatibility

Flexible and friendly Redis APIs and data structures

Fully managed

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

High scalability

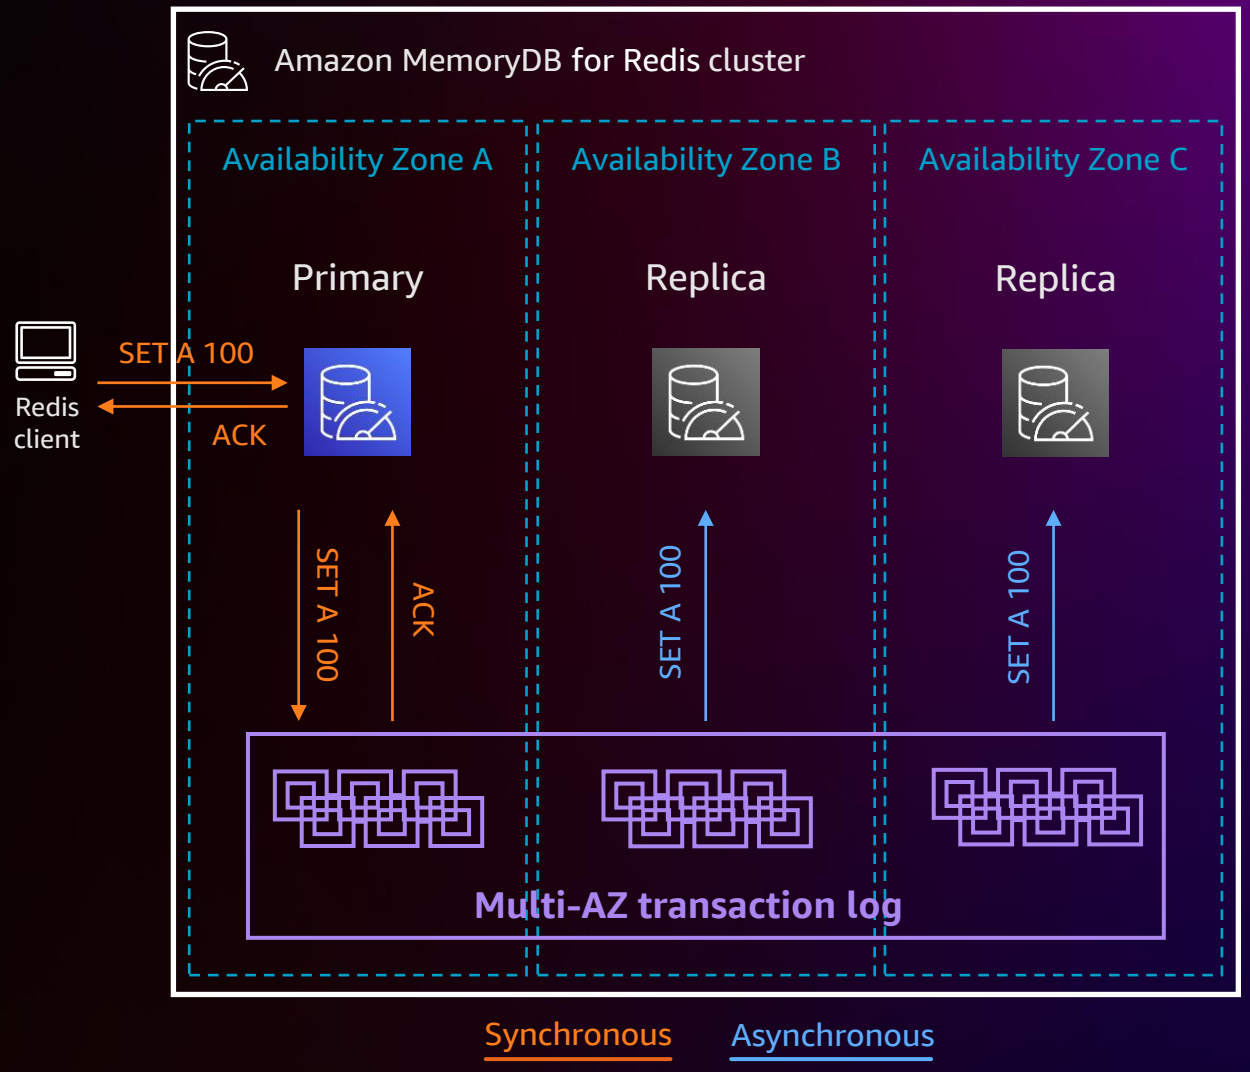
Up to 500 nodes and 128 TiB of in-memory storage per cluster (with 1 replica per shard)

Durability and high availability

Multi-AZ transactional log for durability and replicas for high availability

Security

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



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon MemoryDB for Redis Use cases



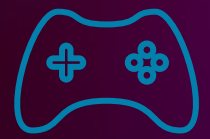
Web and mobile

User content data stores
Session management
Geospatial indexing
Chat and message queues



Retail

Customer profiles and accounts
Inventory tracking and fulfillment



Gaming

Player data stores
Session history
Leaderboards



Banking and finance

User transactions
Fraud detection



Media and entertainment

User data stores
Real-time streaming

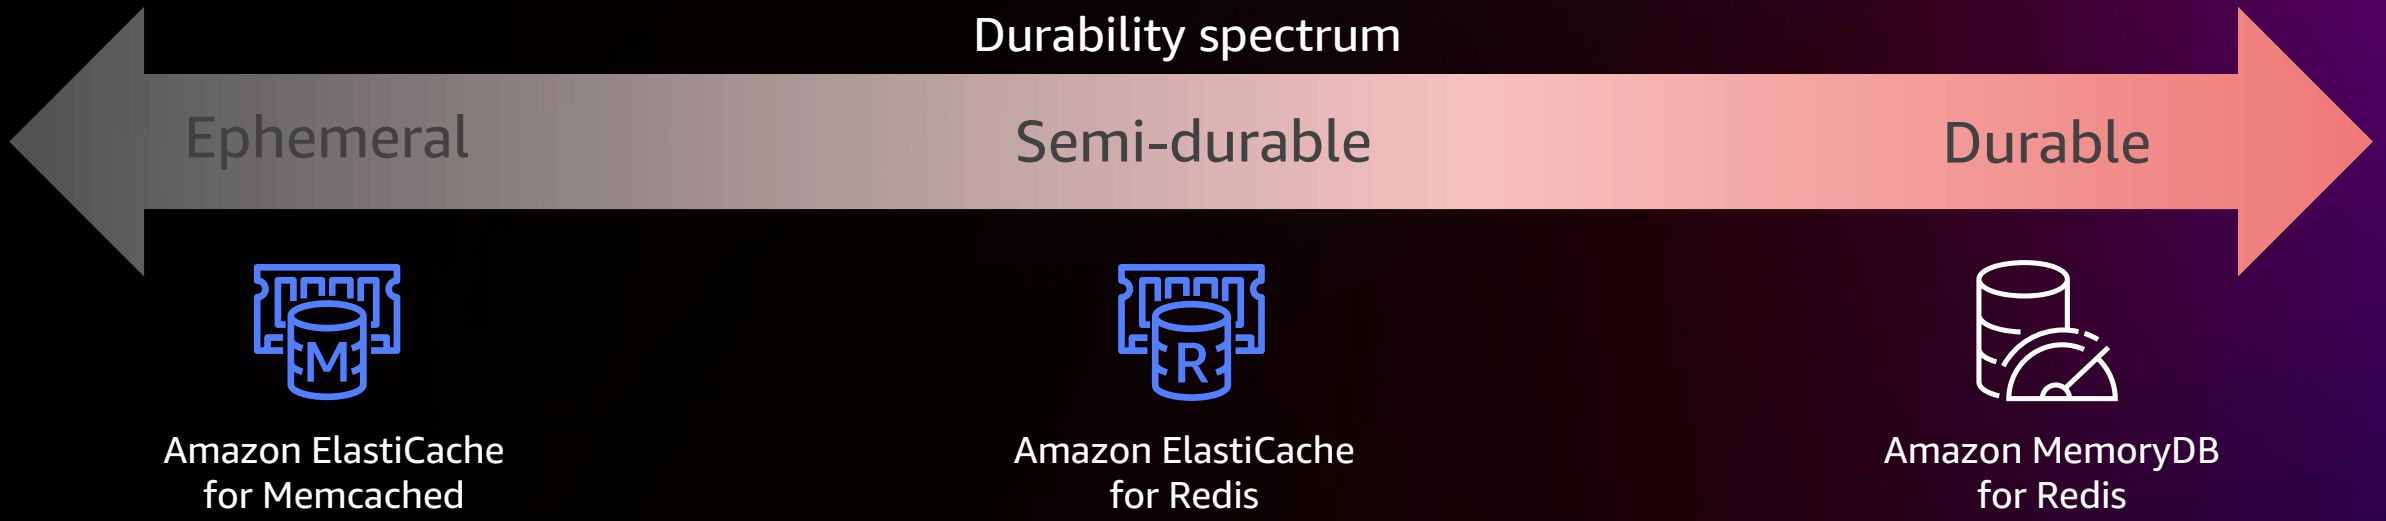


IoT

Streaming device data
Operational insights

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS in-memory managed solutions



Amazon ElastiCache for Memcached

Amazon ElastiCache for Redis

Amazon MemoryDB for Redis

	ElastiCache for Memcached	ElastiCache for Redis	MemoryDB for Redis
Maximum RAM:	20 nodes/13 TiB	500 nodes/340 TiB	500 nodes/128 TiB
Replication:	No	Yes	Yes
Failover:	Not available	Automatic	Automatic
Data types:	1	Many	Many
Commands:	9	100+	100+

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Amazon ElastiCache – New features

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

ElastiCache new features added in 2022



New and improved management console



Memcached v 1.6.12



Memcached TLS



Memcached FedRAMP and HIPAA eligible



IPv6 for both Redis and Memcached

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

ElastiCache new features added in 2022



Redis log delivery through Kinesis Data Firehose and CloudWatch Logs



Native JSON support



AWS PrivateLink



Redis 7 (ACL v2, functions, sharded pub/sub)



Redis IAM authentication

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

ElastiCache 2022 availability



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

How to get started, access technical resources



Amazon ElastiCache

Access webinars and videos, technical blogs, whitepapers, and presentations.

Deepen your skills with digital learning on demand.

Access ElastiCache **Resources**

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Thank you!

Itay Maoz
itay@amazon.com

Siva Karuturi
maravisk@amazon.com

Maulik Shah
maulik@yahooinc.com



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

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.