

Adding a Database Layer

DATABASE SELECTION · AMAZON RDS & AURORA · DYNAMODB · PURPOSE-BUILT DATABASES · AWS DMS · WELL-ARCHITECTED

STUDY & REVIEW

01 MUST KNOW

if you remember five things, remember these

- 01 Match the database to the workload:** the one-size-fits-all relational model no longer fits every application. Before choosing, weigh **scalability**, **storage** (GB / TB / PB), **data characteristics** (data model + access patterns + latency), **durability/availability**, and **cost** — and work backward from the business need.
- 02 Managed means far less to run:** **Amazon RDS** and **Aurora** automate provisioning, patching, backup, recovery, failover, and repair. As you move on-premises → database on EC2 → managed service, your responsibility shrinks until you are responsible only for **optimizing your application and queries**.
- 03 Multi-AZ ≠ read replica:** a **Multi-AZ deployment** *synchronously* replicates to a standby for **high availability** with **automatic failover** — the standby serves no reads. **Read replicas** replicate *asynchronously* for **read scalability** and are **manually promoted** to become standalone.
- 04 Aurora is the cloud-built RDBMS:** MySQL- and PostgreSQL-compatible, managed by Amazon RDS, roughly **one-tenth the cost** of commercial databases. Storage self-heals and auto-scales to **64 TB** across **three AZs**, with up to **15 low-latency read replicas**; **Aurora Serverless** auto-scales capacity on demand.
- 05 DynamoDB is serverless NoSQL:** fully managed **key-value + document**, **single-digit-millisecond** performance, flexible schema keyed by a **partition key** (+ optional **sort key**). A **GSI** uses new keys and is eventually consistent (add anytime); an **LSI** shares the partition key, allows strong reads, but must be created **with the table**.

02 KEY TERMS

TERM	DEFINITION
Amazon RDS	Managed relational database service that deploys, operates, and scales one of seven engines; uses Amazon EBS volumes for database and log storage
Amazon Aurora	MySQL- and PostgreSQL-compatible RDBMS built for the cloud, managed by RDS; self-healing storage auto-scales to 64 TB across three Availability Zones
Aurora Serverless	On-demand, auto-scaling Aurora configuration that starts, stops, and scales capacity automatically; suits variable, new, or dev/test workloads
Multi-AZ deployment	Synchronous replication to a standby in another AZ providing high availability and automatic failover (the standby does not serve reads)
Read replica	Asynchronous, read-only copy that scales out read-heavy workloads; can be promoted manually to a standalone database or created cross-Region
Amazon RDS Proxy	Fully managed, highly available proxy that pools/shares DB connections, cuts failover time up to 66%, and enforces IAM auth with AWS Secrets Manager
Amazon DynamoDB	Fully managed, serverless NoSQL database (key-value + document) with single-digit-ms performance and automatic table scaling
Primary key (DynamoDB)	A partition (hash) key alone (simple), or a partition key + sort key (composite), that uniquely identifies each item in a table
Global secondary index (GSI)	Alternate schema with new partition and sort keys; eventually consistent, add/remove anytime, separate capacity, max 20 per table
Local secondary index (LSI)	Alternate sort key on the same partition key; strong or eventual reads, created only with the table, shares base capacity, max 5 per table

TERM	DEFINITION
Global tables	Multi-Region, multi-active DynamoDB replication for fast local reads/writes and cross-Region availability
AWS DMS	Managed migration and replication service moving data between endpoints (one must be on AWS); supports homogeneous and heterogeneous migrations

03 RELATIONAL VS. NON-RELATIONAL

pick the model that fits the data

ASPECT	RELATIONAL (SQL)	NON-RELATIONAL (NOSQL)
Structure	Tabular rows and columns	Key-value, document, graph, in-memory, wide-column
Schema	Strict, well-defined rules	Flexible; each item can differ
Strengths	SQL, ACID integrity, low effort	Flexibility, scalability, high throughput
Best for	OLTP ; migrating relational workloads	Caching, JSON, single-digit-ms retrieval
Scaling	Vertical (plus read replicas)	Horizontal
AWS	RDS, Aurora, Redshift	DynamoDB, DocumentDB, Neptune, ElastiCache

04 AWS DATABASES BY USE CASE

know the one-line job of each – exam gold

SERVICE	CATEGORY / DATA MODEL	REACH FOR IT WHEN...
RDS / Aurora	Relational (OLTP)	Structured, transactional data — ERP, CRM, ecommerce; migrating relational workloads
Amazon Redshift	Relational warehouse (OLAP)	Petabyte-scale analytics and reporting; columnar storage, massively parallel
Amazon DynamoDB	Key-value + document (NoSQL)	Single-digit-ms at any scale — serverless apps, gaming, media metadata
Amazon DocumentDB	Document (MongoDB-compatible)	JSON documents, flexible schema — content management, customer profiles
Amazon Keyspaces	Wide-column (Cassandra)	High-volume, write-heavy, ms latency — trade monitoring, route optimization
Amazon MemoryDB	In-memory (Redis), durable	Ultra-low latency as a primary store — caching, leaderboards, transactions
Amazon Neptune	Graph	Highly connected data — recommendations, fraud detection, knowledge graphs
Amazon Timestream	Timeseries	Trends over time — IoT sensor data, operational and web-traffic metrics
Amazon QLDB	Ledger	Immutable, cryptographically verifiable history — financial records, claims

05 MULTI-AZ VS. READ REPLICA

availability is not scalability

ASPECT	MULTI-AZ DEPLOYMENT	READ REPLICA
Purpose	High availability	Read scalability
Replication	Synchronous to a standby	Asynchronous
Failover	Automatic to the standby	Manual promotion to standalone
Reads	Standby serves none	Offloads read-heavy traffic
Placement	Standby in another AZ	Can be created cross-Region

06 DYNAMODB SECONDARY INDEXES: GSI VS. LSI

the details are the exam

DIMENSION	GLOBAL SECONDARY INDEX (GSI)	LOCAL SECONDARY INDEX (LSI)
Keys	New partition key + optional new sort key	Same partition key, different sort key
Consistency	Eventually consistent reads only	Strong or eventual consistency (you choose)
When created	Any time — add or remove freely	With the table only — cannot add/remove later
Max per table	20	5
Capacity / size	Separate provisioned capacity; no size limit	Consumes base-table read capacity; size limits apply

07 EXAM TRAPS

where the knowledge check gets you

- × “Multi-AZ gives you read scaling.” — No. The Multi-AZ standby is a **synchronous failover target for high availability** and serves no reads; use **read replicas** (asynchronous) to scale reads.
- × “A read replica automatically takes over on failure.” — No. Read replicas are **asynchronous** and for scaling; you must **manually promote** one to standalone. Automatic failover is a **Multi-AZ** feature.
- × “DynamoDB has no schema at all.” — It has a **flexible** schema, but every item still needs a **primary key** (partition key, optionally + sort key); only the non-key attributes vary per item.
- × “You can add an LSI later.” — **LSIs must be created with the table** and cannot be added or removed afterward (max 5). **GSIs** can be created or removed anytime (max 20).
- × “GSI reads are strongly consistent.” — GSI reads are **eventually consistent only** (data is copied asynchronously). **LSIs** let you choose strong or eventual consistency.
- × “Aurora is a separate service from RDS.” — Aurora is a **database engine managed by Amazon RDS**, chosen as one of RDS’s seven engines — not a standalone service.
- × “Amazon Redshift can serve read replicas.” — Redshift is a **relational OLAP data warehouse** and **does not support read replicas**. For a highly available relational database with read replicas, choose **Aurora**.
- × “AWS DMS can migrate on-premises to on-premises.” — No. **One endpoint must be on an AWS service**; DMS cannot migrate from one on-premises database to another.

08 SELF-QUIZ

answers are upside-down below

- Q1** Which two capabilities does a Multi-AZ deployment provide that a read replica does not?

Q2 What replication type do read replicas use, and what must you do to make one a standalone database?

Q3 How many read replicas can Aurora support, and to what size does its storage auto-scale?

Q4 Name the required and the optional components of a DynamoDB primary key.

Q5 Which DynamoDB index must be created with the table and cannot be changed afterward?
- Q6** Which DynamoDB feature provides a multi-Region, multi-active database with fast local reads and writes?

Q7 Which AWS database is a MongoDB-compatible JSON document database?

Q8 Which purpose-built database stores nodes and edges for highly connected data such as fraud detection?

Q9 What does the AWS Schema Conversion Tool (AWS SCT) do during a heterogeneous migration?

Q10 An app needs a highly available relational database with eight read replicas and daily-growing storage — which AWS database fits?

ANSWER KEY (rotate the page)

Q1 Synchronous replication to a standby and automatic failover (high availability) **Q2** Asynchronous replication; manually promote the replica
Up to 15 read replicas; storage auto-scales to 64 TB **Q4** Partition (hash) key is required; sort key is optional — together they form a composite
primary key **Q5** A local secondary index (LSI) **Q6** Global tables **Q7** Amazon DocumentDB **Q8** Amazon Neptune (graph database) **Q9** Converts
the source schema and SQL code into an equivalent target schema and code **Q10** Amazon Aurora