


An overview of all the different databases in AWS

<https://fibonalabs.com/>



AN OVERVIEW OF ALL THE DIFFERENT DATABASES IN AWS



Amazon Web Services (AWS) is an on-demand cloud computing platform which is trusted and used by organizations all over the world. AWS offers great services like computing, storage, database, analytics, development and management tools, security, networking, IoT and enterprise applications. This blog gives you a glimpse of all the database services available in [Amazon Web Services](#).

AMAZON AURORA

It is a MySQL and PostgreSQL – compatible relational database. Amazon Aurora is five times faster than standard MySQL databases and three times faster than PostgreSQL databases. Amazon Aurora is managed by Amazon Relational Database Service (RDS). This helps in automating hardware provisioning, database setup, patching and backups. It is designed to offer 99.9% availability by replicating six copies of data across three availability zones. It also offers continuous data backup to Amazon S3.

It facilitates easy migration from standard MySQL or PostgreSQL databases with minimal changes. Amazon Aurora offers multiple levels of security for the database such as network isolation using Amazon VPC, encryption at rest using keys which can be created and controlled by AWS Key Management Service. Even the data in the underlying storage is encrypted to ensure security.

USE CASES:

- Enterprise applications – Includes customer relationship management and enterprise resource planning solutions.
- SaaS offerings – Architectures which require significant storage and scalability.
- Web and mobile gaming apps – Include apps which require high throughput, availability and massive storage.

AMAZON REDSHIFT

Amazon RedShift is an analytical database tool where we can query and combine exabytes of structured or semi-structured data using standard SQL. The results can be saved back to S3 Data Lake using open formats, so that they can be used in future for other analytical services like Amazon EMR, Amazon Athena and Amazon Sagemaker.

Amazon Redshift uses the AWS design hardware and Machine Learning to analyse queries. It uses the AWS Nitro System to accelerate data compression and encryption. Graphic algorithms are used to automate organization and storage of data for better query results. AQUA (Accelerated Query Accelerator) is the hardware accelerated cache used in RedShift which ensures fast scanning, filtering and aggregation operations. Amazon Redshift has upto 3x better price performance and improved scalability in comparison with other cloud data warehouses.

USE CASES:

- Business Intelligence – Helps build powerful reports and dashboards using QuickSight or other BI tools.
- Operational analytics on events – Integrates structured and semi-structured data from the S3 Data Lake to get real time operational insights.
- Data as a Service – Enables secure sharing of data both inside and outside an organization and monitors collaboration on live data.
- Predictive analysis – Redshift Machine Learning helps in creating, training and deploying Amazon Sagemaker models in data in data warehouses.

AMAZON DYNAMODB

Amazon DynanoDB is a key-value and document database, which is fully managed, multi-regional, multi-active and durable. The data is stored as collection of key-value pairs where a key serves as a unique identifier. It offers built-in security backup and restoration. In-memory caching is used for internet-scale apps. It can support peaks of 20 million requests per second.

DynamoDB is a serverless database with built-in availability and fault tolerance. It offers both provisional and on-demand modes for cost optimization. It also supports ACID transactions, enabling smooth building of business critical applications at scale.

USE CASES:

- Serverless webapps – Helps to build web applications with automated scaling and availability.
- Mobile Backends – Can be used along with AWS AppSync to build [interactive mobile and web applications](#), with real-time updates, offline data access and data sync with built-in conflict resolution systems.
- Microservices – Flexible and reusable microservices can be built using DynamoDB as a serverless data store.

AMAZON ELASTICACHE

Amazon ElastiCache is a fully managed in-memory data store service which is supported on both Redis and Memcached. It is a type of database that relies primarily on memory for data storage, contrary to databases that store data on disks or SSDs. ElastiCache provides automated set up, hardware provisioning, configuration, monitoring, backup and recovery. Both write and memory processes can be scaled through sharing and data replications.

USE CASES:

- Session stores – Provides responsive session stores for web applications and sites.
- Gaming applications – Provides leaderboards and chat services.
- Real-time mapping and analytics – Includes real-time mapping, Internet of Things (IOT) sensor processing and AI applications.

AMAZON DOCUMENT DB

Amazon Document DB is designed to store and query data as JSON-like docs using the same format in the application code. It enables flexible indexing, powerful ad hoc queries and analytics over collections of documents. It is a MongoDB compatible database service which is purpose-built for JSON data management at scale. It is fully managed and integrated with AWS. It offers high durability and scalability for working with critical MongoDB workloads. It supports upto millions of requests per second without application downtime, regardless of the size of the data. DocumentDB replicates six copies of data across three AWS Availability Zones (AVS).

USE CASES:

- Content and catalog management – Used to store and query content for shopping sites, online publications, digital archives, point-of-sale terminals and self-service kiosks.
- Profile management – Helps to manage millions of user profiles to facilitate online transactions, preference management and user authentication.
- Mobile and web apps – Helps to build [high-performance mobile and web applications](#) with minimal operational burden and development time.

AMAZON KEYSACES

Amazon Keyspaces can be interpreted as a two-dimensional key-value store. It is an elastic, scalable and highly available wide column database service.

With Amazon Keyspaces, we can run Cassandra workloads on AWS using the same Cassandra application code and development tools. It is serverless and hence adopts a pay-as-you-go approach, with virtually unlimited throughput and storage. Amazon Keyspaces offer continuous data backups using point-in-time recovery.

USE CASES:

- Build low latency applications – Helps to build apps for trade monitoring, fleet monitoring, route optimization etc.
- Build apps using open source technologies – Can be used to build apps using open source Cassandra APIs and drivers for JAVA, Python, Ruby, .Net, NodeJS, PHP, C++, and Perl.
- Data storage – Stores information about devices for IoT and gaming applications, entries in log files, message history in chat applications.

AMAZON NEPTUNE

It is a fully managed graph database service enabling applications with highly connected datasets. It supports storage of massive relationship data sets with low-latency access. Neptune supports various graph nodes and languages, eg. RDF, SPARQL and Gremlin. It features point-in-time recovery, replica reading and content backup.

USE CASES:

- Social networking – Includes user profiles and prioritizes content.
- Recommendation engines – Storage of user contacts, purchase histories and preferences.
- Fraud detection – Stores details related to overlapping email addresses, IP addresses and credit/debit card numbers.

- Life sciences – Helps in gene patterning, disease model development and much more.
- Networking and IT operations – Creates network visibility, monitoring and forensic analysis.

AMAZON TIMESTREAM

Amazon Timestream is an automated, time series database service for IOT and operation applications, it eases storage and analysis of trillions of events per day up to 1000 times faster and is cost effective in comparison with other applications. It is a purpose-built query engine which grants access to and analyses recent and historical data. We need not specify the tier location of the data. It helps to identify trends and patterns in in near real time. The data is always encrypted whether at rest or in transit, offering high security.

USE CASES:

- DevOps – Supports performance monitoring and management, network optimization and server monitoring. It also supports clickstream analysis and monitoring.
- IoT applications – Supports IoT analytics for smart devices implementation.
- Industrial telemetry – Helps in monitoring industrial equipment, fleet management, trade monitoring and route optimization.

AMAZON QUANTUM LEDGER DATABASE

Amazon Quantum Ledger Database is a NoSQL Database which uses an immutable, transparent and cryptographically verifiable transaction log. This transaction log is owned by a central authority. It is used to track application data changes with a verifiable history.

Amazon QLDB avoids the need to build a custom ledger application and its related verification tools, which saves time and resources. It is an automated system where data can be queried using a SQL-like API.

USE CASES:

- Financial services – Helps in tracking credit and debit card transactions, claim transactions and auditing.
- HR and payroll – Helps in maintaining employee benefits, performance histories and certifications.
- Retail and supply chain management – Helps in batch tracking, product recall processes and shipping details maintenance.

THANK YOU