

# AWS

## Relational Databases

### what is it?

---

Amazon RDS makes it easy to provision a managed database instance in the cloud. At the time of writing the following database engines were available.

- Amazon Aurora - MySQL and PostgreSQL
- MySQL
- PostgreSQL
- MariaDB
- Oracle
- MS SQL Server

### disaster recovery

---

Disaster recovery relates to the backups, logs and replication instances that are maintained while everything is working fine.

These are switched on, switched over, and analysed when something does go wrong, like a hardware failure, natural disaster or even human error.

- **Failover** - multiple clusters are set up so if one fails the other can take over
- **Mirroring** - maintaining two copies of the same database at different locations. One in offline mode so we know where things are at when we need to use it
- **Replication** - the secondary database is online and can be queried. This is not only good for disaster recovery but can be useful if you utilise one instance for reporting and one for live queries

# AWS

## NoSQL Databases

### what is it?

---

You would use a non-relational database if:

- You have a lot of data with little structure
- Your data structure may change over time
- You want to make frequent changes to how the data is structured

### Dynamodb

---

A key-value database which stores data with

- **a key** - as a unique identifier
- **a value** - anything from an integer to a JSON structure.

### Amazon Neptune

---

A graph database set up to show relationships between people, places, objects and entities. If you've seen LinkedIn's 'Recommended Connections' you've encountered a Graph database.

### AWS Timestream

---

A time-series database which aims to collect data points over time.

### AWS Quantum Ledger

---

A ledger database designed to record a history of economic and financial activity.