

Global Infrastructure



Description:

Resilience, Redundancy and Global Availability: These are buzzwords that get thrown around often by cloud providers. And indeed, resilience is one of the key reasons to move to a cloud-based solution.

A fundamental philosophy of architecting in AWS is that no single instance, or even whole availability zone, should be relied on. If a resource fails, the application should be designed to continue using resources in other zones.

Our solutions architects can review existing infrastructure to ensure that it has a fault tolerance which matches the customer's uptime and risk requirements.

We also get involved when designing new application infrastructure, helping to ensure that fault tolerance objectives are met, not just on a server level but across all resources.

AWS Tech Overview:

Currently AWS has:

- 16 Regions worldwide
- 2-6 Availability Zones in each region
- Over 100 Edge locations worldwide

Load Balancing and Autoscaling give the potential for fault tolerance, but the underlying architecture (and application) needs to be designed to take advantage of this. We can help with workshopping different failure modes, to ensure that the application infrastructure is designed to be resilient to multiple types of failure.

For customers serving applications and web endpoints to a global audience, CloudFront allows data to be cached at edge locations globally. Newer technologies such as Lambda@Edge, enable more complex application logic to be pushed closer to end users, decreasing latency.

At the database level, we can use technologies such as Aurora to implement clusters of database servers across availability zones, and then to roll out read replica databases across multiple regions, if required.

Cost:

Pricing available upon request.

