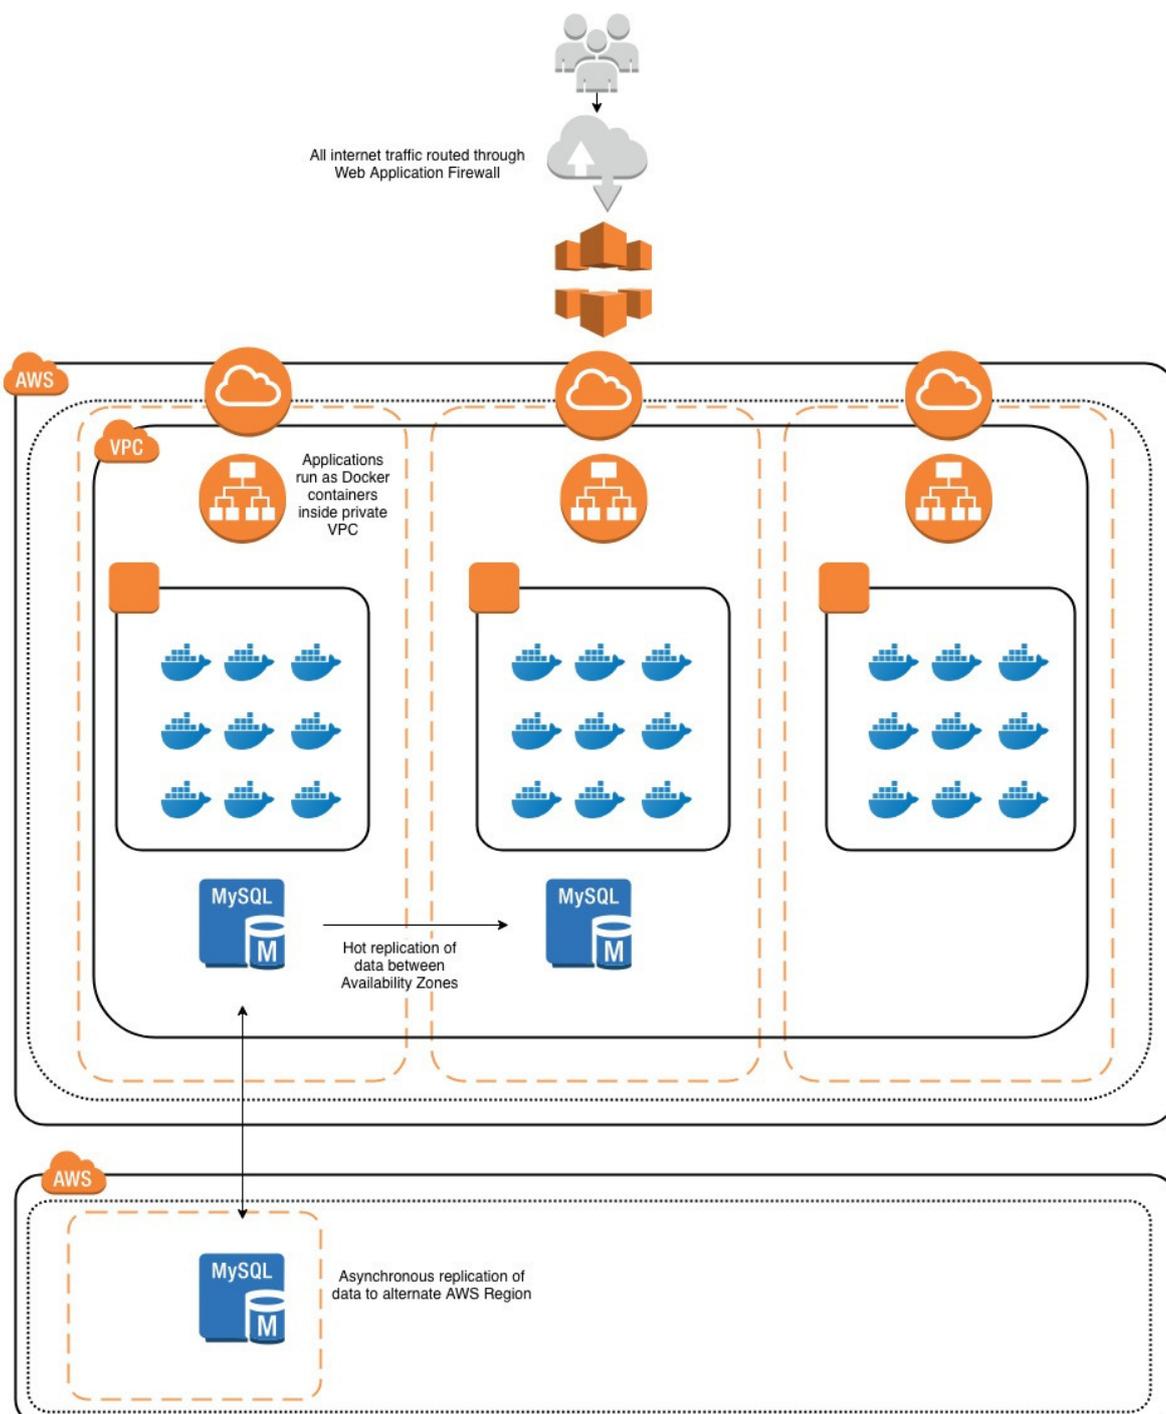


Technical due diligence

Goji's platform technology is built with security, resilience and scalability as primary concerns. Our platform customers use the same technology as our own investment customers so you can be sure that we treat your customers' data and funds as carefully as we do our own.

Infrastructure and architecture

Our platform runs in Amazon's cloud infrastructure (AWS). This allows us to create a best of breed architecture that is secure, resilient and secure. All applications are deployed as Docker containers and are managed by Hashicorp's Nomad container scheduler. This ensures deploys are automated and repeatable and failed containers are automatically replaced.



Security

All Goji services run in a Virtual Private Cloud (VPC) that cannot be directly accessed from the internet. Only authorised Goji staff can access running servers. All internet traffic to the platform is routed through an AWS firewall which actively monitors traffic for malicious activity and can block access to potential threat actors.

Scalability

The Goji platform is designed to scale with customer volume and demand. Our microservice oriented architecture means we can easily horizontally scale to match load by increasing the number of running instances.

Data protection

The Goji platform is hosted in the Ireland AWS region with a backup in Germany. All incoming connections from the internet are encrypted over HTTPS. Personal data stored in the database is encrypted (AES-256). Access to the database is not possible directly from the internet. Access to internal Goji systems is restricted to authorised Goji personnel.

Resilience and disaster recovery

In the event that a system failure occurs, whether an internal failure or the loss of an external dependency, there are a number of levels of redundancy and failover.

The VPC spans three AWS Availability Zones (AZ) within an AWS Region. Each AZ is a physically isolated data centre. Applications run a minimum of two instances and they run in separate AZs. If an AZ fails, traffic can still be served by the alternate running instance.

The internal load balancers monitor the health of all instances and only route traffic to available and healthy containers. If a container fails, the job scheduler automatically takes it out of service and deploys a new one.

All the EC2 instances run within AWS auto-scaling groups (ASG). These ASGs monitor the health of the servers and ensure that a minimum number of health servers are running at all times.

The MySQL database runs in a hot-hot configuration between two AZs. If one instance fails, the alternate is automatically promoted to primary.

If the entire AWS region fails (ie all three AZs become unavailable) then the platform can manually failover to the Germany AWS region. This requires manual intervention and the full system can be recovered in under one hour.

Monitoring and alerting

We monitor application logs and infrastructure health 24x7x365. We utilise automated alerting technology to alert us if there are issues.

Audit

In addition to ongoing internal compliance monitoring, we are subject to annual external security penetration by our partners at <https://www.onsecurity.co.uk/>

