fasthosts ProActive

Digital defence

Geo-redundancy

Welcome to the 'digital defence series', where we break down essential cybersecurity strategies to protect your business. In this final installment, we're talking about geo-redundancy, an important tactic for keeping business continuity and safeguarding your digital infrastructure against regional disruptions.

What is geo-redundancy?

Geo-redundancy is a strategy that involves distributing critical systems, data, and services across multiple physical locations in different geographic regions. The idea is simple – if one location experiences an outage or disaster, another location can immediately take over, minimising downtime and maintaining continuous service availability. Geo-redundant systems are designed to detect failures and automatically switch operations to a backup site, creating resilience against both natural disasters and cyberattacks that could impact a single region.

Why geo-redundancy matters

In today's global business environment, system uptime and availability are important. Unexpected events – like regional power outages, severe weather, earthquakes, or even cyberattacks – can take down data centres, leaving your systems vulnerable to downtime. Geo-redundancy makes sure that if disaster strikes, your business remains operational with minimal disruption. It's a critical defence mechanism that not only protects against data loss, but also maintains customer trust, productivity, and revenue flow.

How geo-redundancy works

- Data replication – If copies of your data are stored in multiple locations, this ensures that if one data centre is compromised, your data is still accessible from another site. This can be done in real-time (synchronous replication) or with a slight delay (asynchronous replication), depending on your business needs. Synchronous replication would be best for businesses needing real-time data accuracy, such as financial institutions, healthcare, and e-commerce platforms. Whereas asynchronous replication would be more beneficial for businesses that can tolerate slight delays, like SMBs, media streaming, and disaster recovery-focused companies.
- Failover mechanisms – If one site goes offline, failover systems automatically reroute traffic to the backup site without any manual intervention. This ensures that end-users experience minimal service disruption, even during significant outages.
- Load balancing – Geo-redundant systems can also balance loads across different data centres, preventing any single location from being overwhelmed by traffic. This is particularly useful in scenarios like Distributed Denial of Service (DDoS) attacks, which we discussed in the first instalment of our 'digital defence series'.

The benefits of geo-redundancy

- Minimised downtime – Geo-redundancy helps make sure that your systems remain online, even during regional disasters or technical failures. With automatic failover in place, downtime is reduced to mere seconds or minutes, ensuring business continuity.
- Disaster recovery – If a major event, such as an earthquake or flood, knocks out your primary data centre, a geo-redundant site will continue running your operations. This avoids the nightmare of losing critical data or access to services.
- Enhanced performance – By distributing your services across multiple regions, you can reroute traffic to the nearest data centre, reducing latency and improving the user experience for global customers.
- Increased security – Geo-redundancy protects against localised cyberattacks. If one region is compromised by a ransomware attack, for example, your backup sites remain unaffected, preserving your data and services.
- Regulatory compliance – Some industries require data to be stored in multiple regions to comply with laws like GDPR, which mandate how and where data can be stored. Geo-redundancy helps businesses meet these regulatory requirements.

What to consider...

While geo-redundancy is an invaluable asset, there are several factors to consider when implementing it:

- Cost – Geo-redundant systems can be more expensive than single-location systems because they require maintaining infrastructure in multiple regions. However, the cost of a prolonged outage often far outweighs the investment in geo-redundancy.
- Replication type – Synchronous replication offers real-time data mirroring between sites, but it can be slower due to latency, especially if sites are far apart. Asynchronous replication is faster but may result in minor data loss if a failure occurs before the data is fully replicated.
- Latency – Routing traffic across geographically distant locations may introduce latency. It's important to optimise network routes and strategically place data centres near major business hubs to minimise this impact.
- Data sovereignty – When replicating data across borders, be mindful of data sovereignty laws. Different countries have varying regulations on data storage, and compliance is critical to avoid legal issues.

Best practices for geo-redundancy

To fully leverage geo-redundancy, it's crucial to follow some best practices:

- Strategic location choices – Place your data centres in locations that aren't vulnerable to the same risks. For example, avoid placing all your data centres on the same coastline or within the same seismic zone
- Regular testing – Frequently test your failover systems and disaster recovery plans to make sure they work as expected when an emergency occurs. Regular testing is important for identifying any gaps or weaknesses in your geo-redundancy setup.
- Data encryption – Ensure that all data replicated between sites is encrypted both in transit and at rest. This protects sensitive information during the replication process and reduces the risk of data breaches.
- 24/7 monitoring – Put real-time monitoring tools in place that keep track of the health of your data centres and services. Early detection of issues allows for faster responses and ensures smooth failovers in the event of a disruption.

In today's always-on world, downtime is not an option. Geo-redundancy isn't just a backup plan – it's an investment in resilience.

Get in touch today by calling 0333 111 2000, or book a consultation to explore how Fasthosts ProActive can help implement geo-redundancy for your business. Stay protected, stay connected!