BEST PRACTICES FOR REDUNDANCY AND AVAILABILITY WITH REDCOM® SIGMA®

Communication Service Providers face considerable challenges when managing the performance of their network. Multiple layers of redundancy are necessary to ensure maximum service uptime. These recommendations for availability represent varying levels of redundancy that service providers should consider when deploying REDCOM's Sigma software.

Server best practices

REDCOM strongly recommends implementing as many of these best practices as possible, as they are a relatively low-cost means of avoiding hardware failures.

- Dual power supplies with independent power feeds
- UPS battery backup
- RAID arrays to protect against data loss
- Use of industrial-grade and/or ruggedized servers and gateways
- Maintain proper operating environment as directed by hardware manufacturer (air flow, temperature, humidity)

Options to maximize uptime

In addition to the best practices mentioned above, choosing one of these options for server redundancy will further improve system uptime.

Option 1: Use a secondary standby server

REDCOM Sigma supports hot standby with minimal manual intervention. This is a low cost option suitable for bare metal
or non-virtualized installations.

- Two servers, each running an independent copy of REDCOM Sigma software
- Servers can be co-located or geographically separated
- REDCOM Sigma configurations are automatically synchronized between servers
- Upon failure, services are restored when a manual switchover instruction is given

Option 2: Use fault tolerance provided by a hypervisor

Certain hypervisors offer different ways to maximize availability, including High Availability Mode, Clustering, and Fault Tolerance.

- Two or more servers, each running a hypervisor
- Used to meet Five 9s Availability
- Highly automated with automatic failover
- No dropped calls
- Suitable for larger customers with many users or mission-critical communications

Redundancy option 1: A secondary standby server

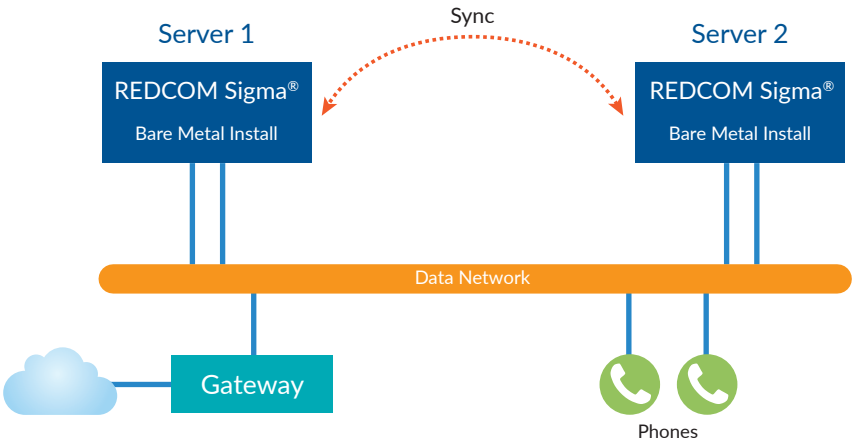

Redundancy option 2: Fault tolerance on a hypervisor