everRun® SplitSite®

Metro-wide availability protection

Disasters, whether natural or caused by human error, can result in the total loss of a physical data center, potentially leaving your business unable to function for days, or even weeks. In regulated industries, a site-wide problem can lead to data loss that risks compliance, adding significantly to your downtime costs. That's why businesses in regulated industries like pharmaceuticals, manufacturing and financial services use everRun SplitSite protection to ensure that all their data is safely replicated and remains available at all times.

everRun with SplitSite extends the protection of your business from localized power failures and building-wide problems using physical machines located in different buildings or data centers. With everRun SplitSite, if disaster strikes in one location, applications and data are immediately available, up-to-date, and fully operational at another location without IT staff intervention. SplitSite connects two physical machines (PMs) in two geographically separated sites. SplitSite provides application availability across both sites using synchronous replication. Both HA high availability (HA) and fault tolerant (FT) protection levels can be selected. And as in a single-site configuration, everRun automatically detects disk and network failures and operates around them. For virtual machines with FT protection, SplitSite will keep VMs running with no downtime, even through a PM or site failure. When a failed site or PM is returned to service, everRun SplitSite will automatically resynchronize disk drives and VM memory.

everRun's SplitSite supports disaster-tolerant deployments that maintain hardware redundancy, as well as redundancy of physical computer rooms and the buildings containing them. By supporting geographical separation, this powerful disaster tolerant solution further safeguards your business from major downtime due to potentially catastrophic events, such as flooding and power outages. everRun SplitSite eliminates the cost and complexity associated with typical reactive recovery products. Customers often use SplitSite in larger campus or metropolitan settings, as a real-time alternative to multi-site disaster recovery.

Key benefits

- Protection against localized power failures and building-wide problems
- Protection against physical machine failures within a metro-wide area
- Automatic resynchronization when returning failed sites to service

SplitSite requirements, and licensing

There is no universal distance limitation for SplitSite, as a number of factors can come into play. Any intervening network switches add to latency, and increase the possibility of losing the connection between the nodes resulting in a split brain condition. For all SplitSite configurations, Stratus requires that you also use a quorum service.

Instead, SplitSite configurations are subject to maximum latency specifications--No more than 10ms round trip A-Link latency for HA VMs, and 2ms round trip A-Link latency for FT VMs. Separation of PMs up to 10km (using 1 Gbps fiber) is a common A-Link network topology that can meet most latency requirements. Individual performance, even within these latency specifications, can depend upon the specific application.

The only conditions for Stratus support are a license, use of quorum, and compliance with latency requirements. Otherwise, any networking equipment and topology are accommodated. In a typical, mainstream network, a safe distance between servers is 5km to 10km. However, Stratus does have some customers with very fast networks who are using SplitSite today in scenarios where the PMs are 50km or more apart from each other.

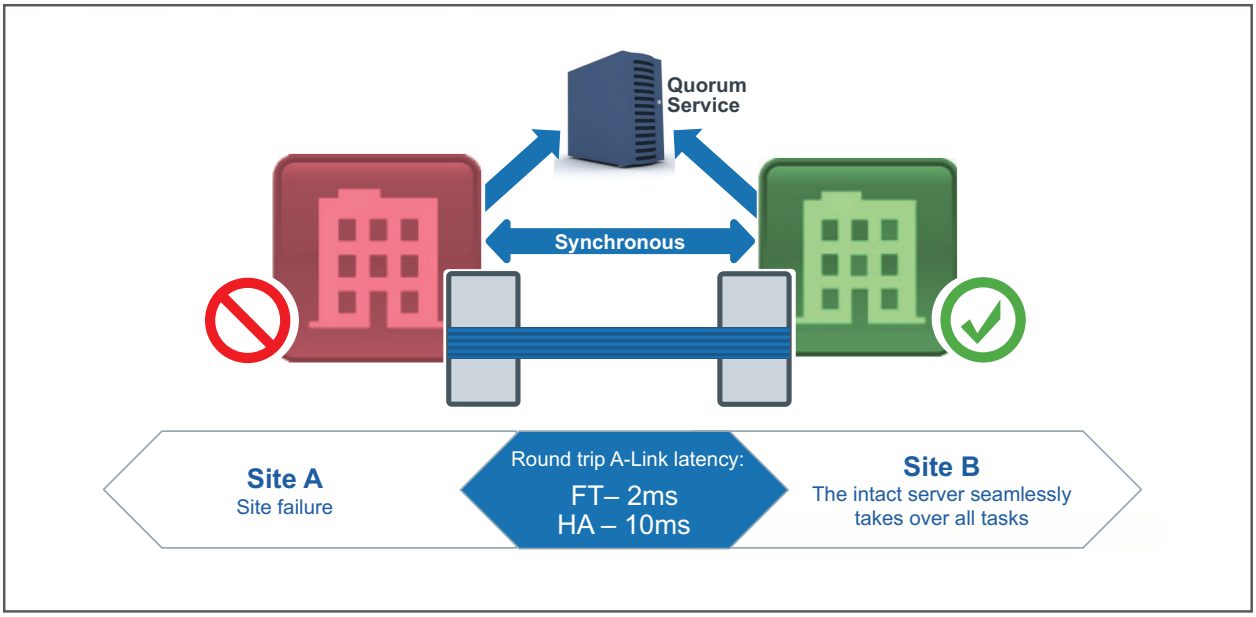

SplitSite and Quorum servers

Use of quorum is required for SplitSite configurations to protect against data loss (due to split-brain) and to safely enable VMs to start up automatically if a second everRun PM or site has failed. In a SplitSite configuration, you will need at least one, and optimally two, quorum servers. These servers are used to protect against network failures which might cause the two everRun nodes to lose communication with each other and operate in a split brain scenario. Quorum availability is improved, and mandatory VM shutdown scenarios are minimized, if quorum servers are placed at a third location and an appropriate quorum networking design is implemented. Quorum servers don't require dedicated hardware or have any specific network latency requirements. They can run as a Windows service and be installed on almost any Windows workstation or server that's used for other purposes, as long as the computer is left running 24 hours a day.

Quorum server considerations

- Quorum service software can be installed on any general-purpose computer or laptop running Windows Server 2012, Windows Server 2008, Windows Server 2003, Windows Vista, Windows 7, or Windows 8; always powered on and with 100MB minimum disk space and a network interface card with connectivity to the everRun configuration via the management network.

SplitSite network requirements

- NICs must be at least 1 Gb and fully-duplexed; use 10 Gb, if possible.
- Switches and/or fiber-to-copper converters connected to the private network must be non-routed, non-blocking and support IPv6.
- For systems running FT-protected VMs, A-Links require:
  - A minimum bandwidth of 1 Gbps per VM
  - A maximum inter-site latency* of 2ms, round-trip time
- For systems running only HA-protected VMs, A-Links require:
  - A minimum bandwidth of 155 Mbps per VM
  - A maximum inter-site latency* of 10ms, round-trip time
- Do not use a common card (multiport NIC) for multiple A-Links.
- A-Links can be dedicated point-to-point fiber connections or on a VLAN. VLANs used to connect the A-Link ports must not filter any communications between the two everRun nodes.

** Calculate latency at 1ms for each 100 miles of fiber, plus any latency added by non-routed, non-blocking switches or fiber converter*

Stratus

www.stratus.com