

A Guide to Architecting the Active/Active Data Center

EXECUTIVE SUMMARY

All things being equal, every company would choose an active/active architecture. It means nonstop availability of applications, higher capacity and better performance. But high availability usually comes with high costs and complications. Organizations resist implementing an active/active data center due to the need to modify applications, databases and networking configurations. Applications must be altered to understand the difference in read vs. write procedures and know how to bridge various data centers. A typical compromise is to run a less expensive active/standby system despite their lengthy recovery time after a failure and the secondary site is sitting idle until needed.

Today there's a new option for achieving easy, affordable active/active systems that works for 8 out of 10 organizations. The game has changed. IT associates don't need to build support into applications. That's too complex and labor intensive. They don't need a third-party cluster management system for aligning capabilities. That can make for inconsistent applications and reduce capacity. And they don't need to build or change anything within the SQL server. You can accomplish data consistency without it. Even in the cloud.

Imagine architecting an ideal active/active system that doesn't need continual maintenance. With an abstraction layer that separates the applications from the databases, you can. You'll have a neutral zone that provides all the necessary traffic management and still understands the SQL state and the database architecture. It supports replication and has performance tools baked in.

Your organization can easily, quickly implement ScaleArc's database load balancing software and leverage the following four critical advantages immediately.

Read/write split

No need to modify applications with read-intent strings. ScaleArc's software brings top performance and reliability by knowing the difference between reads and writes and distributing the traffic appropriately.

Geo-aware balancing act

With a built-in understanding of response times, the software sends queries to the closest servers. Servers with better performance automatically get more loads.

Smarter replication

Your organization dictates what kind of replication delay is tolerable. ScaleArc's monitoring tools sniff out lags and immediately stop sending traffic to a node that has fallen behind.

Stronger failover processes

ScaleArc's software handles failovers effectively in one data center, across multiple locations or even in the cloud. It reserves queries in a queue once a failover has begun, saving users from application errors. Once the secondary server is in place as the primary, the queries are automatically forwarded. If necessary, other read queries can be routed to other secondary servers for higher availability and continuity of service.

This improved method for architecting active/active data centers even comes at a reduced cost. Active/standby operations cost 2.5 - 3 times more than a single data center. Active/active operations cost only 1.25 - 1.8 times more than a single data center. Each site operates less equipment and resources are more efficient in the active/active model allowing organizations to save money while ensuring system reliability.

For more information on how ScaleArc's software can help your company, read the full white paper. »