Zero Downtime
in SQL Server Environments

EXECUTIVE SUMMARY

Downtime is a dirty word. Every year organizations worldwide blame downtime for lost time, money, productivity and consumer confidence. Until now, outages seemed inevitable. Planned and unplanned downtime happens. Feeding the need for round-the-clock, insatiable uptime has been impossible because nonstop was not an option. Until now.

ScaleArc has a solution for planned and unplanned downtime dilemmas. With automated, fast failover, unplanned outages can be mitigated across various data centers. Planned downtime such as applying updates, doing backups or just running diagnostics on a SQL Server no longer require servers to be offline. Instead, you'll get maintenance done with an innovative approach and keep applications up and running simultaneously.

How it Works

In most environments today, applications are connected directly to database servers. If that connection drops, your customers get errors. Not with ScaleArc. Using an abstraction layer between the application and the server, applications are protected from failure. If a server goes offline, your traffic is automatically directed to another source assuring seamless, reliable user experiences. ScaleArc offers your team new methods for uptime all the time during planned and unplanned outages.

Unplanned

Today's leading high availability (HA) solutions – Windows Failover Clustering and SQL mirroring – limit failover to one datacenter. Multiple locations can be used, but not without heroic efforts such as modifying applications or making changes at the DNS or TCP level. Virtualization and cloud computing have changed the game. ScaleArc eliminates the problem and allows for multiple datacenter failover. The software knows the difference between read and write queries and automatically routes each one to the proper server. In SQL Server 2012 and 2014, AlwaysOn technology is key to HA. ScaleArc integrates seamlessly with AlwaysOn Availability Group Listener or Virtual Network Name (AGL/VNN). Now the care and feeding of server farms are easier because teams can see each server's current role, which ones are part of an availability group and what kind of replication is being employed so servers are constantly ready.

Planned

Even the best laid plans have critics. Customers don't care about your SQL Server upkeep schedule or your need for diagnostics. With ScaleArc, your IT team can perform planned (or emergency) maintenance like upgrades or patching behind the scenes with no downtime to users. ScaleArc offers various ways to reduce server load depending on how gradually or suddenly organizations need to move and whether the server in question is still online or has been taken offline.

ScaleArc offers two methods for conducting maintenance on a server while it's still *online*: load balancing bias and reduced server connection. Load balancing bias gradually reduces traffic while the server's spare capacity performs backups and diagnostics. The reduced server connection method is more dramatic. It instantly lowers the number of queries to a server and allows companies to see an immediate and significant decrease in server traffic. It's best used for storage maintenance or any other kind of work that doesn't require the server to be offline but does require a sudden traffic reduction.

Offline server maintenance is needed for making changes to a server including hardware replacements, software updates and RAM upgrades. For organizations that need to bring a server offline, ScaleArc offers a unique queuing system integrated with its software. During failover, traffic from the server that's being taken offline is queued up then diverted to a secondary server transparently ensuring minimal interruption.

Learn more about zero downtime in a SQL server environment. »