

# The Top 5 Ways to Maximize SQL Server Availability

EXECUTIVE SUMMARY

Maximizing SQL Server availability is a double-edged sword. In a perfect world, every user in every organization would have every system available 100% of the time. But as companies face the costs and logistics to make that happen, the opportunity arises for important dialogue between IT professionals and executives to determine how best to increase high availability in a cost effective way without adding technical complexity. Consider these strategies:

## 1. Service level agreements

SLAs help turn executives into IT partners. With business objectives in mind, clearly show management what “100% up time” costs the organization in hardware, software, and resources. You’ll gain understanding, empathy, and a training ground to work together toward a solution that’s realistic and achievable.

## 2. The right tool for the right job

Getting a job done right is easier with the right tool. Thanks to clear SLAs, you can address the needs of each database independently knowing the business’ expectations. Plan accordingly, mixing and matching SQL Server’s four most popular HA (high availability) options along with their pros and cons.

1. AlwaysOn Availability Groups bring three benefits to the table. They span multiple data centers, can address DR concerns and handle scale-out operations. But they aren’t supported on SQL Server Standard Edition.
2. Starting in SQL Server 2012, AlwaysOn failover clusters span multiple subnets making them great solutions for HA and DR concerns. But there’s no redundant data here posing an architectural problem with only one single point of failure.
3. Log shipping is easy to automate and lends itself well to DR situations. But extensive management and additional resources are needed to make it applicable in HA and scale-out situations.
4. Replicating is great for scale-out – after all, its supported on SQL Server Standard Edition – but it’s not a solid solution for HA or DR due to data constraints that are typically built in.

## 3. Prep and practice for failover

IT professionals should understand each step associated with failover and the time required to complete each one. Be comfortable and conversant on what happens to core components when failover occurs as well as what happens to operations or transactions that are in motion when the failover happens. Prepare talking points on recovery time and performance expectations to share with users once the recovery process is complete.

## 4. Consider the applications

Rather than just looking at SQL Server uptime, organizations could focus also on how applications are connecting and responding during failover. Applications must be modified, however, and making complex changes is risky.

## 5. Regularly test and validate solutions

Without regular testing and validation, not only is it possible for slight and subtle changes to the environment to break fault-tolerance or failover capabilities, but the IT professionals who manage these solutions are also letting their operational and disaster management skills with these systems atrophy.

---

Considering all the strategies listed above, the cleanest method for maximizing SQL Server and application availability is by deploying a third-party abstraction layer between the application server and the SQL Server back end. It’s more simplistic than customizing applications. It’s cheaper than picking up additional server licenses. It’s more flexible for cloud migrations and datacenters in multiple locations. And applications connect to a proxy reducing failover errors.

Database load balancing software gives IT professionals a new weapon to go along with that double-edged sword. It can reduce the complexity and time associated with the failover process and brings organizations closer to that perfect balance between high availability and manageable costs. [Click here to read the full white paper.](#) »