

ScaleArc for SQL Server

Overview

Organizations around the world depend on SQL Server for their revenue-generating, customer-facing applications, running their most business-critical operations on this database. The relatively late arrival of replication to SQL Server databases have made it challenging for organizations to drive high availability with their SQL Server deployments. In this solution brief, we will discuss the challenges in designing high availability and scalable performance into SQL Server environments. We will also provide an overview of how ScaleArc for SQL Server – with support for SQL Server 2005 through 2016 – addresses these challenges in a transparent manner, without requiring changes to the apps or database. The ScaleArc software has enabled organizations to increase revenue, decrease operational costs, and deliver new apps and services faster. With ScaleArc, organizations can upgrade their apps to consumer grade with no code changes, making their apps never down, always fast, and scale anywhere.

ScaleArc for SQL Server enables organizations to deliver “consumer grade” apps – apps that are never down, always fast, and scale anywhere – with no code changes.

Database Industry Trends

The need for faster, easier to scale, highly available database environments has always been at the forefront of any IT organization. The explosion of data in the recent years coupled with the dramatic reduction in data storage costs has put an enormous strain on relational database systems. While some firms are exploring the implementation of NoSQL database solutions for some of their applications, the general consensus is that relational database systems, especially for online transaction processing (OLTP) environments, are here to stay. On the compute front, Intel’s focus for over a decade on increasing the number of cores on a physical CPU socket while keeping peak CPU speeds constant has fundamentally challenged scale-up solutions. Some RDBMS products, (e.g., MySQL) have deployment architectures that exploit distributed compute and storage systems like master/slave with scale out, read-only database servers. In contrast, SQL Server installations, especially those editions prior to SQL Server 2012, have traditionally “scaled up,” using ever-larger servers, to meet the performance demands of new applications.

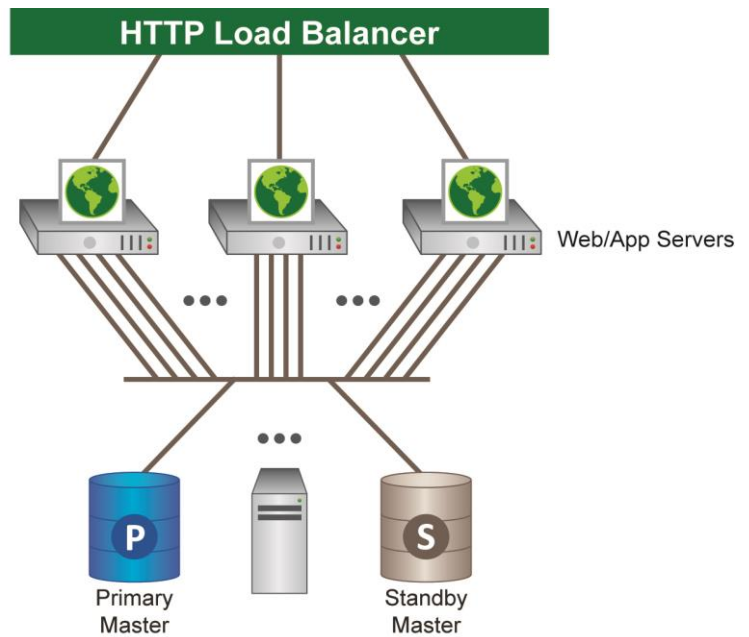
Challenges in Today's SQL Server Environments

SQL clustering and SQL mirroring, the two most popular “scale-up” SQL Server deployment architectures, provide high availability and data integrity only upon node failure. Scaling SQL Server environments usually involves investments in more powerful servers and expensive shared storage systems. Redundant hardware is usually in passive standby mode for high availability and does not add to any performance gains, resulting in wasted resources and reduced ROI. Shared storage systems, on top of being prohibitively expensive, can become single points of failure and can cripple remote Disaster Recovery (DR) architectures. While web servers have been able to scale n+1 active/active setups across data centers along with Global Server Load Balancing (GSLB) at the web tier, SQL Server databases at the data tier tend to be active/passive.

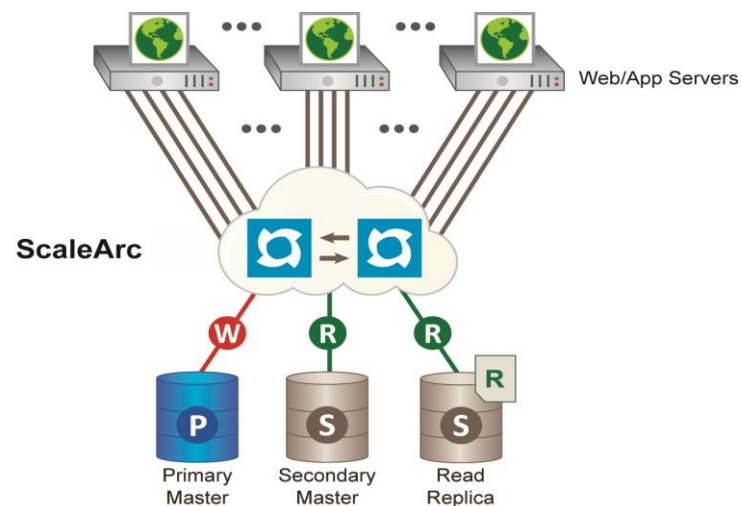
The lack of simple scale-out deployment models prior to SQL 2012 really place SQL Server at a disadvantage for growing Web 2.0 firms in eCommerce, Software as a Service (SaaS), healthcare, financial services, and social networking as well as for traditional enterprises across various industry verticals. The high numbers of SQL connections originating from the scaled-out web server farms usually overwhelm the SQL Server, especially during peak traffic hours or during traffic bursts. SQL Server does not have any built-in query cache that helps alleviate read-intensive workloads that are typical of most database workloads. Organizations opting to go with peer-to-peer replication or log shipping soon realize the shortcomings, like lack of SQL-aware load balancing, transparent failover and increased application complexity.

Finally, OLTP systems today are monitored heavily to correlate and track customer experience with end-to-end system performance data. SQL Server does not natively provide visibility on a per-query basis to aid application performance monitoring of production environments.

SQL Server's AlwaysOn architecture makes significant improvements to both scale out and high availability. With AlwaysOn, SQL Server can allow you to make your secondary servers readable for read-only queries, and it supports auto-failover / auto-promotion as well. This architecture makes it much easier to set up more SQL instances to offload your primary instance and have better reliability for your primary instance, but it still does not provide for the ability to load balance between the secondary instances or to automatically send read-only queries to secondary instances, since applications have to be modified to be able to work with read-intent connections to effectively utilize the read-only-replicas.



Without ScaleArc, organizations are at higher risk of app downtime, have no scalability, and suffer poor performance.



ScaleArc enables zero downtime for apps and 2x performance.

ScaleArc for SQL Server: Instant Scale and Availability without Any Application Changes

The ScaleArc database load balancing software solves the challenges of building consumer-grade SQL Server deployments. ScaleArc adds a transparent layer between application servers and SQL servers, enabling auto failover, instant scale up, and transparent scale out. ScaleArc lets you “upgrade” your apps to be consumer grade – never down, always fast, scale anywhere – with no code changes. Adding our database load balancing software to the data tier has enabled organizations big and small to drive substantial revenue increases and operational cost reductions. For example, we’ve enabled:

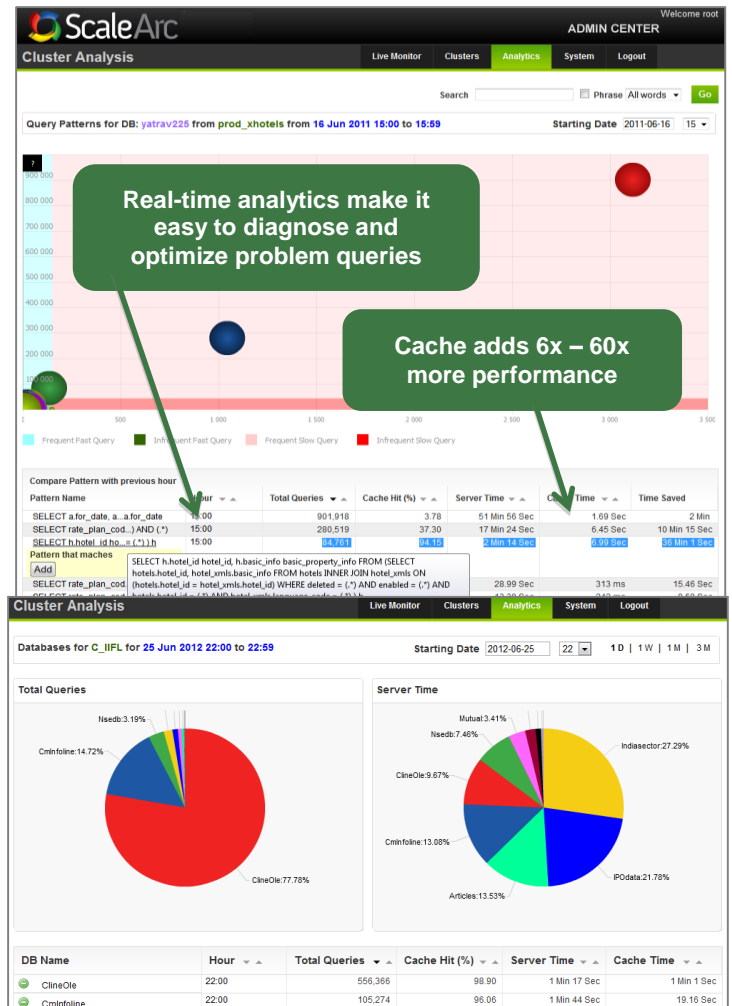
- **\$40K in savings**, in just three minutes – from avoiding application downtime during database failover
- **\$2M in increased revenue**, every year – from avoiding downtime from maintenance windows
- **\$3M in increased revenue**, every year – from doubling website performance
- **\$320K in savings** – from avoiding dev time recoding for database scaling
- **2x faster app rollout** – from dev time avoided

Upgrading your SQL Server Database with ScaleArc

The ScaleArc software provides the greatest value when paired with modern databases offering scale out and failover, but if you’re upgrading to get that functionality, you are better off upgrading with ScaleArc already in place. First, ScaleArc can provide immediate value, offloading your server ahead of the upgrade. ScaleArc then improves the upgrade process, letting you upgrade without interruption to the application. A zero-downtime upgrade isn’t possible otherwise.

Zero Downtime and Fast Performance

The ScaleArc database load balancing software abstracts application servers from SQL Server database servers to break the forced 1:1 dependency between the app and database tiers. ScaleArc enables zero downtime for apps – it provides auto failover, instant scale up, and transparent scale out. Real-time analytics enable instant troubleshooting, one-click query caching for SQL Server offload and up to 24x faster responses, wire-speed security enforcement, and much more.



The ScaleArc query analytics display provides insight into all SQL traffic, reducing troubleshooting time and helping with capacity planning.

ScaleArc can perform read/write split for queries and stored procedures, and it performs SQL-aware load balancing. The read/write split with dynamic load balancing allows for seamless deployments with peer-to-peer replication or log shipping replication for SQL Server. Now with ScaleArc, SQL Server customers have the luxury of multi-master architectures with active/active databases allowing for better availability, higher performance, and increased ROI. SQL Server customers not yet on AlwaysOn and utilizing only SQL clustering or mirroring benefit from single-server improvements such as connection management and transparent caching.

ScaleArc provides a transparent caching solution for SQL Server databases that require no changes to existing applications or database servers – with no development effort needed to realize the caching benefits. ScaleArc's SQL caching is a query-level agentless approach that uses a NoSQL (i.e., purpose-built hash map) database to store read query responses, thus enabling blazingly fast responses for subsequent queries and allowing for massive scalability of existing SQL Servers.

ScaleArc's caching works at the query pattern level; these SQL query patterns are generated at wire speed from de-duplication of ScaleArc's centralized logging of all SQL queries flowing through ScaleArc.

ScaleArc also supports caching of stored procedures. Caching increases app performance by serving content faster, resulting in a better user experience. Database CPU and I/O load is reduced significantly even in the case of SQL clustering and mirroring, freeing critical system resources, accelerating reads and writes, and extending the life of the server. In high-availability mode, the cache on a secondary (passive) ScaleArc instance is synchronized in memory in real time with the primary ScaleArc instance, resulting in no traffic surge if the primary ScaleArc instance fails.

ScaleArc also provides for unparalleled real-time visibility of all SQL traffic traversing through production database servers. The SQL analytics is derived from de-duplicating granular log data. This approach does not require any sampling (all queries are logged and all data is utilized), nor does it add any performance overhead on the application or the database server. The comprehensive query/stored procedure performance data is charted in a human-readable graph and highlights all frequent-but-slow queries for instant troubleshooting. These frequent-but-slow queries can be immediately added to the ScaleArc cache with a single click, to instantly scale and accelerate overall database and application performance. With ScaleArc's real-time SQL instrumentation, application developers and DBAs can now get a non-intrusive, performance-centric view of the SQL query load, all with a simple click of the mouse. ScaleArc's SQL analytics can also be used for auditing the SQL traffic and analyzing performance bottlenecks.

Leveraging AlwaysOn

ScaleArc makes it easy to take advantage of the failover and scalability features of AlwaysOn in SQL Server. With ScaleArc, organizations don't need to program into the app support for read-intent strings, for example. ScaleArc automatically translates read-only calls and routes them to read replicas, without any client side intervention required. In addition to that, ScaleArc also provides features like SQL query caching, load balancing, a query firewall, and deep SQL analytics to identify and fix common SQL performance and security issues.

Microsoft Customer Case Study

Business Challenge

The answers.microsoft.com site is amongst Microsoft's most popular community platforms, supporting 2 million users a day. As part of Microsoft's "Get Current/Stay Current" initiative, its IT organization planned to upgrade from SQL Server 2008 to SQL Server 2014. At the same time, the organization's server hardware was approaching end of life, so the team decided the timing was right to move to a virtualized environment.

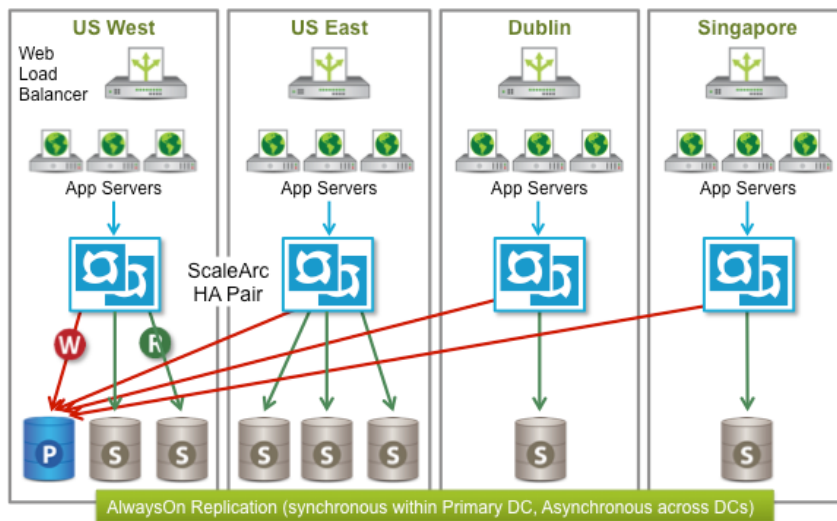
Senior services engineer Michael Schaeffer set his main objective as ensuring 24x7 availability across the site's four geographically distributed data centers. The previous architecture relied on Windows Failover Clustering to ensure uptime, but the new virtual environment no longer supported that mode, so Microsoft could not rely on clustering for the Availability Groups Listener failover. In addition, while SQL Server 2014 supports the number of secondary servers Microsoft IT needed across the four data centers, it does not support any form of load balancing.

Schaeffer and his team did not like any of the options for overcoming these challenges – deploying a TCP load balancer or rewriting the app to load balance the reads. Speaking to other infrastructure groups within Microsoft, Schaeffer learned about ScaleArc's database load balancing software, and the team chose ScaleArc to ensure a zero downtime environment for its website.

“Because our website functions across four data centers with over 2M hits a day, site availability is our number one issue. ScaleArc provides us with a seamless failover process, providing us with greater levels of uptime than we’ve ever seen before.”

-- Michael Schaeffer

Senior Services Engineer, Microsoft



The answers.microsoft.com site provides auto failover across four data centers, avoiding app errors during failover.

Solution

The Microsoft team had initially estimated it would take 9 months just to update the web application to leverage AlwaysOn. However, with ScaleArc's database load balancing software, the team was able to reduce the timeline to just four months, without re-architecting the app, servers, or network. ScaleArc for SQL Server augments the failover capabilities in AlwaysOn, enabling cross-data center failover that is transparent to the users of the site.

Results:

- Enabled the team to “go live” in 4 months vs. 9 months with no app changes
- Provides cross-data center failover, without app errors
- Enables zero downtime maintenance and patching for continuous availability
- Enables read/write split with no app changes
- Delivers geo-aware load balancing
- Supports 8x traffic capacity – all servers used

Summary

SQL Server drives the critical business of organizations across the globe. Building an HA, scalable environment has proven challenging.

ScaleArc database load balancing software lets organizations tap into the power of modern databases without writing a single line of code. As a result, ScaleArc customers can instantly upgrade their apps to be consumer grade – never down, always fast, and scale anywhere.

For more information about ScaleArc for SQL Server, visit www.scalearc.com or call us at 408.414-7191.



2901 Tasman Drive, Suite 205
Santa Clara, CA 95054
Phone: 1-408-780-2040

info@scalearc.com
www.scalearc.com



ScaleArc enables consumer-grade apps for today's digital business – apps that are never down, are always fast, and scale anywhere. ScaleArc's database load balancing software helps organizations of all sizes eliminate application downtime from database outages or maintenance, improve application performance, and scale database capacity – all without writing a single line of code. As a result, ScaleArc customers increase revenue, reduce operational costs, and accelerate time to market. Learn more about ScaleArc and its customers and partners at www.ScaleArc.com.

© 2016 ScaleArc. All Rights Reserved. ScaleArc and the ScaleArc logo are trademarks or registered trademarks of ScaleArc in the United States and other countries. All brand names, product names, or trademarks belong to their respective holders.

10/07/2016