

Improve DevOps Efficiencies by Building an Agile Data Tier

EXECUTIVE SUMMARY

Database scalability is one of the biggest challenges DevOps teams face today. Business executives expect the ability to change course quickly in the competitive race to market, but without a way to scale quickly and accurately, you'll lose your edge waiting for the infrastructure to catch up with your marketing plans. In an agile software development environment, if you must consider scalability implications every time you roll out new code, the adaptive planning and fast, flexible process quickly becomes bogged down.

ScaleArc's software brings solutions to the most common problems associated with scalability. Using an agile data tier allows teams to scale more effectively, gain control and visibility into the process, solve problems faster and win back that advantage that's rightfully yours.

The scalability challenge hits technical teams on three levels:

- At the *application* level, groups are understandably apprehensive to change applications. It's risky and time consuming. But changes are needed to accommodate different database roles and to make sure data from secondary locations is current.
- At the *database* level, IT teams are expected to provision new hardware for testing and development groups to isolate these activities from production. Whether physical or virtual, this takes hours and may even require downtime interrupting the workflow of the engineers who needed the help in the first place.
- And finally, on a *troubleshooting* level, without visibility to see exactly how and why database slowdowns occur, dev teams are hamstrung and left to try and fix errant behaviors through inefficient trial and error methods.

ScaleArc offers solutions for every challenging level.

Applications

Unfortunately applications don't automatically evolve as our business grows. Without lots of intervention, they can't take advantage of multiple database servers or know to move to the right one for high availability. They don't know the difference between read and write queries or which ones go to which servers. Database load balancing software handles that and lets applications be applications, removing the expectation that they'll continually perform well despite intricate modifications to accommodate growth. The software manages that complexity providing replication-aware load balancing, read/write split, visibility into potential caching opportunities to save resources and even automated failover.

Database as a Service (DBaaS)

Provisioning separate locations for testing and development groups has never been easier than with ScaleArc. With database load balancing software, within seconds and with no downtime, you can divide one database cluster into individual pieces creating an internal DBaaS that allows each member room to work independently. This architecture saves time, enhances the development experience and reduces costs since the organization is maintaining just one shared database cluster rather than multiple ones.

Visibility and control

In the past, how applications would behave after changes were rolled out was a guessing game. Any steps taken to proactively affect change meant teams risked over provisioning and paying for unnecessary resources or under provisioning and watching performance suffer as a result. Instead, teams need a performance-centric view to know how applications and databases will react to each other and how traffic might change as a result. ScaleArc's software offers an unparalleled level playing field to measure performance at different times before and after implementing changes. Now teams can easily compare the new analytics against the old and even rely on automated alerts that highlight significant differences. Using these actionable analytics, teams can identify security violations and quickly block threats like SQL injections that can cause significant performance disruptions.

[Learn more about building an agile data tier with ScaleArc in the full white paper. »](#)