Database Performance Today:
THE NEED FOR SPEED & SCALE

Best Practices Series

Database Performance at the SPEED OF BUSINESS

Best Practices Series

When it comes to digital transformation or competing in the digital economy, it's all about the data. Business success hinges on how rapidly and effectively decision makers and the systems they depend upon can access essential data—often in real time. AI, machine learning, advanced analytics, and IoT are front and center. The question is: Are databases ready for the challenge?

All enterprises need to be able to respond to challenges and opportunities as they arise in real time or close to it. At the same time, many are still supporting legacy environments that may have limited capacity and get bogged down by large workloads or sudden spikes in usage by large numbers of end users or customers.

The challenges faced in today's data environments were borne out in a recent survey of 203 data managers conducted by Unisphere Research, a division of Information Today, Inc., in conjunction with Oracle. Just about all data managers, 95%, have seen increases in data volume over the last 3 years, and expect this trend to continue. At the same time, more than nine in 10 respondents are finding it difficult to keep up with this growth—and it's affecting their ability to assure peak performance. Three out of four DBAs manage more than 10 databases, with some handling hundreds of databases. Six in 10 data managers are performing multiple major database updates per year, many on a monthly or even weekly basis. Three in four DBAs feel that applications can be deployed faster with increased

database management automation ("2019 Quest IOUG Autonomous Database Adoption Survey").

"Relational databases have made tremendous improvements in performance, availability, and security over the past couple of decades," writes Lawrence Miller in *Autonomous Database for Dummies*. "They can run up to 100 times faster; can be configured for zero data loss; and have hardened security capabilities that can protect against malicious internal and external threats. These attributes have been enhanced by cloud databases and infrastructure services that deliver elastic scalability and provisioning for real-time agility and growth. Database workloads that were deemed too large or 'mission critical' to run outside corporate data centers just a

few years ago now run in public clouds. In addition, capabilities such as database resource deployment, monitoring, and management can also be automated, leading to greater operational efficiencies and cost savings. So, what's missing? The degree of manual intervention required to manage today's cloud databases and all these attributes inhibits true Database as a Service (DBaaS)—as a utility, or 'driverless' offering, if you will. As a result, enterprises are unable to realize the full operational and financial benefits of the cloud."

The following are ways to address this looming database performance crisis:

Plan for growth. It's rare these days to find a data environment that isn't growing at a clip of at least 5% per month. At this rate, the need for enhanced hardware or cloud services will keep accelerating—for both storage and memory. But meeting growth requirements means more than throwing additional hardware—or cloud capacity—at the problem. Tuning and auditing data environments must be part of an ongoing strategy.

Accelerate the move to automation in all its forms. While major vendors such as Oracle have been promoting the concept of completely autonomous databases, there are steps along the way to automate. Key functions such as backup and recovery, provisioning, alerting, and monitoring are already increasingly being turned over to machines. By migrating to autonomous data management solutions, data managers' roles will shift to tasks more integral to their businesses, such as advising management on data strategies, or helping identify and secure data and analytics resources. Many of the day-to-day aspects of data management—troubleshooting, patch management, security, and provisioning— can be quickly and efficiently handled by the machines. Automation—whether via cloud services or through internal systems—is seen as critical to ensuring that complex and sophisticated data systems are delivering what's needed to the business.

Develop a real-time backup strategy. A real-time backup strategy is perhaps one of the most critical aspects of data management to ensure that business applications and data have resiliency, so

that in the event of a problem, enterprise users and customers may not even become aware of it. A real-time backup strategy helps avoid the staggering losses that come with business interruptions—which can cost thousands of dollars per minute. This is an area that automation can address in a profound way. Manual backup and recovery processes mean hours, if not days, of potential downtime to the enterprise. When it comes to introducing automation to data operations, backup and recovery plans come first, a survey from Unisphere Research shows. As public cloud adoption grows, much of this growth is driven by backup and recovery supporting transaction environments. Data managers

Meeting growth requirements means more than throwing additional hardware— or cloud capacity—at the problem. Tuning and auditing data environments must be part of an ongoing strategy.

and professionals see the greatest value in automating backup and recovery processes, to ensure that their data assets are protected, important, followed by business continuity and disaster recovery, according to the "2019 IOUG Data Environment Expansion Survey," conducted by Unisphere Research and sponsored by VMware.

Formulate a highly focused data strategy. An enterprise seeking to promote greater data-driven decision making and analytics applications needs a clear strategy for turning this data into actionable insights. Let's face it, the volumes of data now surging through enterprises from every corner—such as IoT and edge devices—can be overwhelming, and much of it may not be of material importance to the business at the moment. A data strategy helps both

business and data managers concentrate on those data streams that matter the most. That way, processing and human resources are not wasted on managing datasets with marginal value to the enterprise.

Tighten your infrastructure. For data to flow and be processed at real-time speeds, the data infrastructure needs to be ready, with the right tools and pipelines. Big data has been on the scene for some time now, and the need for tools and platforms that can efficiently handle these workloads has never been greater. Such solutions can handle transactional data, analytic queries, real-time streaming, and relationship graphs. While major migrations may be too expensive and disruptive, there are a number of tools and platforms that support building new solutions that extend systems—such as a range of database testing tools, or real-time streaming engines such as Apache Spark and Kafka.

Move forward with a cloud strategy. With the large volumes of data now moving through enterprises, as well as sophisticated applications such as AI now being applied, most on-premise systems will likely be overwhelmed, and become difficult to upgrade and maintain. Today's generation of cloud services offer greater capacity, as well as automation and resiliency. Organizations can choose to swap out current systems for cloud-based ones, or gradually migrate pieces of their applications to cloud providers. With a hybrid cloud architecture, enterprises may enjoy the best of both worlds—secure local processing with existing on-premise systems, supplemented by cloud services that provide backup or handle spikes in demand.

DATABASE PERFORMANCE AT THE SPEED OF BUSINESS

Organizations are under pressure to deliver superior customer experiences and solutions almost instantly. This calls for a robust backend database structure. With greater automation and cloud capabilities formulated through a forward-looking data strategy, databases can operate at the speed of business. ∎

—Joe McKendrick

Geo-Scale MySQL for Continuous Availability and Fast Response Times for Global Operations

Speedy database performance and scalability are critical to the health and growth of any enterprise, but neither has any value unless the database layer is properly protected from service interruptions … in other words, downtime.

Continuent, the MySQL Availability Company, has been providing solutions for continuous operations enabling business-critical MySQL applications to run on a global scale with the ability to perform zero-downtime maintenance since 2004.

Our customers are leading SaaS, e-commerce, financial services, gaming and telco companies who rely upon MySQL plus our solutions to cost-effectively safeguard billions of dollars in annual revenue.

Our flagship solution in particular—Tungsten Clustering—allows enterprises running business-critical MySQL and MariaDB database applications to cost-effectively achieve continuously fast operations with commercial-grade high availability (HA), geographically redundant disaster recovery (DR) and global scaling.

The best practices for MySQL database management that we have built into our solutions allow us to measure MySQL availability uptime in months and years. And what better way to share some of these best practices than by letting our customers showcase them for us?

USE CASE I—GEO-SCALE MYSQL FOR CONTINUOUS GLOBAL OPERATIONS WITH VERY FAST RESPONSE TIMES

This customer is a global gaming company with several hundred million world-wide player accounts. The challenge here was servicing a single application for a geographically-distributed audience; in other words, a pretty unique challenge.

It requires fast local response times for read traffic, a limited number of updates, and a single consolidated view of the data across the world, which is very typical for gaming applications, and for all account/subscription management systems in general.

The solution we implemented for this customer is comprised of four (4) geo-distributed Composite Tungsten Clusters, with one active cluster in USA West accepting writes and updates and handling local read traffic, and three passive Tungsten Clusters in USA East, EMEA and APAC providing very fast local reads to access the player accounts.

The benefits that this solution provides are clear: geo-scale, availability and disaster recovery.
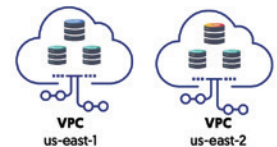
More specifically, it includes low-latency, geo-distributed data access providing fast response times for read traffic as well as local, rapid-failover automated high availability. This combined with simple administration, system visibility and stability also helps create high return on investment.

USE CASE II—MYSQL HIGH AVAILABILITY, ZERO-DOWNTIME MAINTENANCE AND DISASTER RECOVERY FOR SAAS

This customer is a SaaS provider with sensitive medical data, which offers EMA, the Electronic Media Assistant solution (that "speaks the language of doctors").

The challenge they were facing in fact came from using AWS, which allowed them to rapidly provision database and application servers. However, the instances, underlying storage, and management interface were not highly available. By using Tungsten Clustering, they could and can quickly deploy clusters in AWS and recover from multi-zone AWS outages.

The deployment that we implemented for this customer's growing needs includes 17 Tungsten MySQL clusters in AWS (totaling 68 MySQL servers) using a Pod Architecture.

The Pod Architecture provides practically infinite scalability for the SaaS providers as they can just keep adding new Pods when they sign up new customers. Each Pod includes a 3-node Tungsten Cluster deployed in multi-AZ AWS, and a 3-node DR cluster deployed in another AWS Region with all replication traffic encrypted in-flight.

The benefits this SaaS provider customer is able to reap from our solution include continuous operations, high availability, scalability and better data protection.

REAPING THE CONTINUENT TUNGSTEN FOR MYSQL BENEFITS

These are just two examples among many others as to how our solutions and best practices can help customers who manage business-critical MySQL environments achieve the availability, scalability and safe operations with fast response times they are looking for.

If you're looking for speed and scale when planning MySQL database performance, do keep in mind that this requires a solid foundation … and that downtime is not an option.

Consider a complete solution such as Tungsten Clustering which provides the peace of mind that comes from the knowledge that downtime won't occur. This in turn makes speed and scale a real possibility.