IRI
Page 14
SPEED & SECURITY FOR
EVERY DATABASE

database
TRENDS AND APPLICATIONS

# DATABASE PERFORMANCE STRATEGIES: ENSURING FAST, SCALABLE, AND AGILE OPERATIONS

database
TRENDS AND APPLICATIONS

Best Practices Series

# DATABASE PERFORMANCE STRATEGIES: ENSURING FAST, SCALABLE, AND AGILE OPERATIONS



Best Practices Series

In an era of surging AI, database performance has never been more critical. AI and automation in all its forms demand well-vetted, real-time or close to real-time data to deliver business-ready results. Still, while technology providers have stepped up their platforms and tools to help databases perform in this high-charged environment, and consultants have pushed solutions, there are still too many missing pieces that need to be in place.

The pressure is on data leaders and professionals to ensure the right data is delivered at the right time to the right parts of their enterprises. This calls for a new role of data transformation managers, who would build on transformative technologies such as AI and high-level analytics and high-level automation to help organizations keep up with the demands of today's digital economy.

There's a tangible cost to the business when databases don't deliver in an effective way. For almost 2 out of 3 data sites (62%), performance is a regular, monthly issue, a survey of 218 data managers by Unisphere Research shows. Database teams are required to spend the bulk of their time doing SQL performance tuning, as indicated by many respondents. Notably, only 15% report they were "rarely" impacted by database performance issues. It's no surprise that addressing performance issues ranks as the top time-consuming task faced by database teams (DBTA Oracle Database Strategies & Support Survey—"The Diversification and Decentralization Revolution," March 2026).

*The pressure is on **data leaders** and **professionals** to ensure the **right data** is **delivered** at the **right time** to the **right parts** of their **enterprises**.*

Similarly, database teams are facing dueling realities, a survey of 500 IT executives by Everpure finds. For example, 70% say any downtime is unacceptable. With the advent of AI needs, 86% say retrieval augmented generation (RAG) will reshape their databases. Another 60% can't accurately predict infrastructure costs. A majority, 73%, agree that hybrid management is harder today than it was 3 years ago. As a result, organizations are delaying making changes, even when they can't afford to.

Superior database performance is critical to every aspect of the

*Lapses or glitches in database performance will impinge employee productivity, customer satisfaction, and, ultimately, the ability for businesses to meet the needs of a digital economy.*

business—not only for emerging AI applications, but also digital commerce, employee productivity tools, ERP systems, and supply chain management networks. Lapses or glitches in database performance will impinge employee productivity, customer satisfaction, and, ultimately, the ability for businesses to meet the needs of a digital economy.

Assuring superior performance has grown more challenging in recent times, amid rapidly expanding data volumes and workloads, regardless of data environment—be they built on cloud, on-prem, hybrid, or distributed infrastructures. This calls for more attention and resources to accelerate data processing without the administrative overhead previously associated with such efforts.

The following are best practices to help ensure the delivery of business insights and knowledge as AI and high-end automation becomes mainstream across today's businesses:

**Look to AI as an enabler.** AI exacerbates this challenge, but at the same time, it serves as a powerful performance tool. AI applications are built around pattern recognition across very large datasets, taking out much of the manual or cobbled-together integrations and troubleshooting efforts. AIOps (AI for IT operations) provides for machine learning across database performance processes through anomaly detection that finds and remedies query latency, memory consumption, and CPU load ahead of issues that may arise. In addition, through AI-enabled root-cause analysis, AI can rapidly sift through query logs and network traffic in order to prevent or deal with system bottlenecks.

Then there's the advent of the self-driving database, something that has been pitched for the past decade, but is now nearing reality, thanks to AI. These products can provide autonomous, proactive security, as well as updates and patches.

**Employ AI early and often.** AI also should be activated at the earliest stages of database application development. It can be used to optimize schema design and architecture. AI, put in place early, will also enable more effective capacity management and the prediction of scaling needs.

**Think hybrid.** Database performance will hinge on the interplay between next-generation applications—such as those linked to large or small language models—and existing applications such as ERP. It's a hybrid world now, both in terms of cloud and on-prem systems and data, but also as it applies to pre-AI and post-AI applications. The key is standardization to provide consistency between these two poles of applications.

*A high-performing database environment can carry weight between highly satisfied customers versus those leaving for competitors.*

**Keep business requirements front and center**. From a business perspective, performance makes all the difference. A high-performing database environment can carry weight between highly satisfied customers versus those leaving for competitors. In addition, a high-performing database will keep the information flowing that is needed to help the business change direction, move to new opportunities, or resolve damaging issues. It's important that business leaders are part of the process that determines how and where technology investments are directed.

**Keep data teams up to date and supported.** With the rise of greater database automation and AI, a transformation of data professional roles is underway. The roles of data managers are being elevated to ones that are more consultative, helping businesses pinpoint the information needed for competitiveness. AI-driven or highly automated performance tuning can take on much of the manual or jury-rigged processes that have slowed down database management, but humans need to remain in the loop. Data compliance mandates, for example, still require double-checking the work AI produces.

Database managers and professionals are also being called on to play a key role in this transformative environment, fully understanding the possibilities and limits of AI. This is where ongoing training is essential—data managers need to understand how the AI tools and platforms now at their disposal can be leveraged to deliver data more effectively to the business. In the Everpure survey, most database administrators (53%) expect to upskill in AI and machine learning, along with data integration, within the next 3 years.

Along with data transformation managers, stakeholders from the business side will also need to play an active role in emerging high-performance databases. Database activity will also require engagement from data architects, data engineers, and line-of-business managers.

Boosting database performance has become more than an exercise for more efficient data environments, although this is a benefit as well. High-performing databases will also help separate high-performing enterprises from less agile competitors. ■

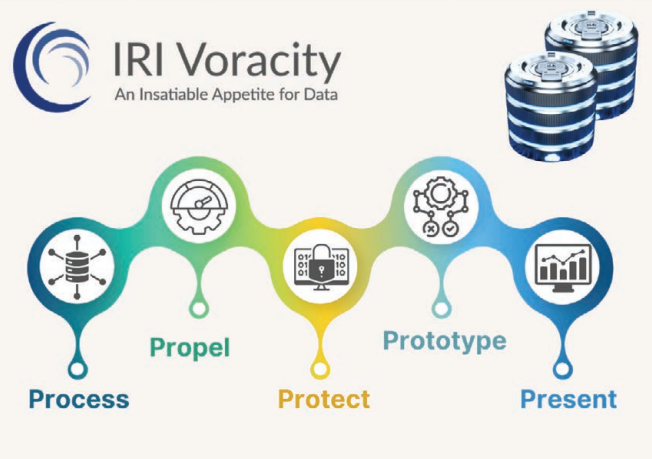
—Joe McKendrick

# Speed & Security for Every Database

IRI Voracity
An Insatiable Appetite for Data

Because of the volume of operational data now available in on-premise and cloud databases, and the risks associated with PII, the speed and safety of database operations are more critical than ever. In use cases involving database migration, ETL, data cleansing, testing, and privacy law compliance, IRI Voracity is a proven solution platform for DBAs, GRC teams, and developers.

Powered by the CoSort data transformation engine and featuring mask/test tools like FieldShield and RowGen, Voracity streamlines database DataOps and DevOps. Voracity consolidates data discovery, integration, migration, governance, and analytics. Its combinable tasking reduces job steps, runtimes, and the costs of multiple tools.




## MULTIPLE SOLUTIONS

***Process***

Voracity rapidly performs high-volume structured data processing operations, including: searching, profiling, and classifying data, change data capture/replication, query integration and acceleration, bulk table extraction and pre-sorted loads, offline reorgs, database migrations, external transforms, and cross-platform SQL operations.

***Propel***

Voracity optimizes the performance of several of the above tasks through task consolidation, state-of-the-art algorithms, and the multi-threading of extraction, classification (PII search), and sorting. Voracity sort/join/aggregation outperforms SQL transforms, ELT appliances, other ETL tools, and even mainframe utilities.

***Protect***

Voracity governs data through built-in data quality and data masking functions. Validating, cleansing, enriching, and standardizing data improves its utility and reliability. Encryption, pseudonymization, redaction, scrambling, hashing, and anonymization functions de-identify PII in production and test environments for breach nullification and legal compliance.

***Prototype***

Voracity can also subset and synthesize smart, safe test data for prototyping, DevOps, demos, and benchmarking. Its "RowGen" component parses DDL to rapidly generate and populate—or otherwise provision—pre-sorted, structurally and referentially correct test data for an entire schema.

***Present***

Voracity can report on data in databases and files as it searches, transforms, cleanses, and/or masks that data. Its embedded BI includes custom detail and summary report formatting with cross-calculation and statistical functions. Alternatively, Voracity users routinely transform and wrangle data to rapidly prepare and feed data to their own analytic, visualization, or AI environments.

## USE CASES

***Optimizing Database Performance***

Organizations in multiple industries have leveraged Voracity to optimize their database performance. For example, a financial services company used Voracity to accelerate their ETL processes, reducing data processing times by 50%. This improvement enabled them to deliver real-time insights to their clients, enhancing customer satisfaction and competitive advantage.

***Ensuring Data Privacy and Compliance***

Healthcare and BFSI companies leverage the data classification, masking, re-ID risk scoring, and robust auditing features in Voracity components like FieldShield or DarkShield to comply with HIPAA, PCI-DSS, etc. By masking sensitive information and providing robust audit logs, IRI customers are protecting patient and customer privacy, and avoiding costly data breaches.

***Streamlining Data Migration***

A global retail company faced the challenge of migrating their data to a new cloud-based database. With Voracity, they were able to discover, map, and replicate their data efficiently, minimizing downtime and ensuring a seamless transition.

## CONCLUSION

The ability to accelerate and govern database-related operations is essential. IRI Voracity is a popular and affordable data management platform that combines high-speed data integration and cleansing with multi-platform test data masking and generation. By leveraging the power of Voracity, you can accelerate your database performance, ensure data quality and compliance, and drive business success.

Learn more at https://iri.com/solutions/database-acceleration and https://iri.com/voracity. ■