10 WAYS IN-MEMORY DATABASES ARE HELPING ENTERPRISES GET AHEAD

In-memory databases and grids have entered the enterprise mainstream. New offerings from pure-play in-memory database providers as well as the large relational database management systems vendors are helping organizations that are scrambling to keep pace with the demands of an always-on, real-time economy. These in-memory databases are emerging in many forms—from extensions of relational database management systems to NoSQL databases to cloud-hosted NoSQL databases.

These new technologies couldn't come a moment too soon. With enterprises of all stripes seeking to compete on analytics, speed is everything. Enterprise systems are being overwhelmed by new workloads coming through web and mobile channels, as well as through the Internet of Things, making faster processing is a necessity. In today's ultra-competitive economy, enterprises need information and insights as soon as they are generated.

Traditional business intelligence tools—weighed down by the challenges inherent in disk-to-processor-based workloads—are not meeting this demand, and enterprises can no longer afford to be held back by these outdated and latency-ridden database environments.

In-memory databases and technologies enable decision makers to get to the information they are seeking rapidly and more readily. These fast databases also are opening up enterprise data to new questions that decision makers have never been able to ask before on their slower, more cumbersome systems. "Shifting the data storage layer from disks to main memory can lead to more than 100x theoretical improvement in terms of response time and throughput," according to a recent IEEE paper authored by Hao Zhang, Gang Chen, Beng Chin Ooi, Kian-Lee Tan, and Meihui Zhang.

With in-memory, data is stored in the random access memory (RAM) portion of servers and computers, versus transferring it off disks. There is none of the latency as data makes round trips between disks and memory. This also offloads much of the mundane work around query-building and report generation from IT departments. The ever-declining costs of memory are also making in-memory databases economically feasible.

While in-memory technology has been on the market for many years, it's only lately that the price of memory has dropped to the point where processing data in RAM compares favorably to moving it in and out of disks. There are also growing business demands for real-time analysis and insights that make in-memory a compelling option. Enterprises increasingly seek real-time capabilities, as well as the ability to manage information coming in from the emerging Internet of Things.

As a result, in-memory is being actively deployed across many enterprises today—from online ad placement to financial trading to production systems. The benefits of in-memory, of course, depend on overall application performance—there may be issues that are occurring outside of any latency resulting from data moving between disk and processor. But there is a lot of disk-induced latency to wring out of infrastructures.

The benefits in-memory databases provide to enterprises, both for IT departments and for the business at large, are many-fold, including the following:

1. Easier handling of big data. Enterprises of all types are now capturing and storing a wide variety of data—from images to documents to traditional relational data. In many cases, this data is being stored for longer amounts of time—potentially for decades. The data may be stored in older systems or storage formats, making it costly and time-consuming to

extract and manage. In-memory databases also work well with Hadoop environments. In addition, in-memory databases work well for the deep data analytics performed against large datasets for sophisticated functions such as predictive analytics, employed by data scientists and analysts.

2. Greater ability to leverage the real-time Internet of Things: Businesses are increasingly relying on real-time data streams from devices or sensors embedded into products, out in their supply chains, or among employees and contractors. Real-time responses to events require analytics performed at blazing speeds, in-memory databases enable such rapid response times.

3. Faster and more distributed decision making. The increased ability of decision makers to build their own queries and dashboards with little or no involvement from IT will shift analytics initiatives to the business side. End users will be able to conduct more sophisticated segmentation and market analysis because entire datasets will be available for immediate processing. Large volumes of data can be quickly crunched and analyzed, often reducing wait times for reports from hours to a matter of minutes, if not seconds.

4. More effective reporting. In-memory databases can support many more queries and run much faster than traditional analytic environments, so they also can more readily support highly graphical, intuitive dashboards—which usually require processing power on the back end. Visually enhanced dashboards mean more intuitive interfaces, and, therefore, reduced requirements to hire Ph.D.s in statistics to decipher rows and columns of numbers.

5. Greater business agility. Since in-memory databases have enormous capacity and are able to support large datasets to be queried, any end user at any level in the organization may be able to quickly identify and isolate any subsets of enterprise data for further investigation. Previous BI environments required that data be extracted in chunks, slowing down analysis as new data needed to be swapped between the BI platform and disk arrays. In-memory data stores can also simply be cleared from the system when complete, making room in-memory for new jobs. The increased speed and flexibility offered through in-memory processing means a transformation in the way organizations approach decision making. End users no longer need to send requests to the IT department to pull the information or reports they need. More ad hoc analysis is possible this way, since end users can access data almost instantaneously, and conduct analysis on the spot, asking any questions that come to mind.

In-memory databases work well for the the deep data analytics performed against large datasets.

6. Greater end user adoption of data analysis services and applications. Non-technical end users may be more likely to adopt analytical applications, especially with the increased access and ease of use enabled through in-memory databases. In a world in which many end users are accustomed to the instantaneous responsiveness of online sites such as Google, the slow response times of many enterprise query engines is no longer acceptable.

7. Offloading processing jobs from over-taxed transactional databases. Running analytics against a transactional system can severely tax the performance of the system, and the ability to quickly launch analytic applications from an aligned database will stave off such a severe performance penalty.

8. Faster application development. Developers will not be consumed with the requirement to program disk performance features into applications. As a result, applications will be turned around and put into production at greater speeds, while seeing increased performance.

9. Scaling with almost no limits. In-memory configurations can sit on groups of processors—symmetric multiprocessing—which support large memory capacity. The capacity of such systems can already go as high as those of large disk-based databases. Potentially, all stored in a RAID array could potentially be moved right into machine memory—which could even scale into the exabyte range in some server farms.

10. Reduced infrastructure—and infrastructure maintenance—is required. Rather than having a database system linked to disk arrays, and adversely affecting database servers with analytic workloads, this all can be handled within in-memory environments. Many typical database infrastructure features may even slow down in-memory performance. "When data access becomes faster, every source of overhead that does not matter in traditional disk-based systems, may degrade the overall performance significantly," according to the IEEE researchers. "The shifting prompts a rethinking of the design of traditional systems, especially for databases, in the aspect of data layouts, indexes, parallelism, concurrency control, query processing, and fault-tolerance."

IN-MEMORY DISRUPTION

The impact of in-memory databases and technology is extending well beyond the data center—it is a potentially disruptive technology and approach to information management. Decision makers will have increased latitude to explore new areas of opportunity for the business, or to address challenges in new ways. Database operations will work much closer with the business, delivering value at the moment it's needed. ■

—Joe McKendrick

HOSTBRIDGE® TECHNOLOGY

HostBridge Redis® for z/OS®

System of Record Data at In-Memory Speed and Scale

REDIS, THE IN-MEMORY NoSQL data store, powers many of today's ultra-scalable cloud and mobile apps—Twitter, Tumblr, Weibo, Stack Overflow, and others. Twitter, to cite just one example, uses Redis to process 39 million queries per second. Now, for the first time, HostBridge Technology brings the speed, scalability, and flexibility of Redis to IBM® z Systems™—the platform that powers the largest enterprises on the planet.

HostBridge Redis for z/OS enables z Systems customers to put their system of record data—residing in DB2®, VSAM™, and other data systems—to work in new mobile, cloud, and analytics applications. Customers can thus leverage their highest-value data to meet exploding demand far more efficiently and economically, and compete successfully in the data-driven world.

WHY REDIS ON Z/OS

IBM z Systems customers tell us they need solutions addressing today's big market drivers—cloud architectures that embrace z/OS assets, data access from any device, and end-to-end analytics—all in near real time. As the following scenario explains, they need something other than traditional data systems.

For years, customers of an investment services firm have accessed CICS® and DB2-based account information via web and mobile devices—information that is updated daily, hours after markets close. The company faces two challenges. First, customers will soon demand instant information gratification, which the company will have to deliver to keep pace with competitors. Second, in a volatile world, online traffic can spike at any moment. If the Dow rises or falls by 1,000 points in a day, the company might see one hundred times their normal traffic.

Aware that traditional data systems were not designed with today's devices and transaction volumes in mind, and that rising demand can lead to skyrocketing processing costs, the company is seeking a more flexible, economical data technology.

HOSTBRIDGE REDIS FOR Z/OS

HostBridge Redis is a modernizing extension to IBM mainframe transaction processing systems and database systems. Using HostBridge Redis, customers share their system of record data with any other system on any platform, inside or outside the enterprise, at lightning speed, with unparalleled scalability. With Redis, they also add value to existing systems while reducing processing costs.

Built on the latest Redis code base, HostBridge Redis offers capabilities that deliver significant benefits:

- In-memory data caching: permits much faster request/response processing than traditional databases
- Advanced key-value store: provides more flexible data structures to transcend the limitations of relational or hierarchical models
- Data persistence: yields dependability, scalability, fault tolerance, and ACID compliance
- High availability: ensures virtually 100-percent uptime
- Asynchronous replication via master-slave and publish-subscribe models: enables ultra-scalability and automated sharing in near real time across applications and data assets running on all z Systems and non-z platforms.

HostBridge Redis drives revenue opportunities by making z Systems customers' highest-value data interoperable—at speed and scale—with new mobile, cloud, and analytics apps. And when running on the zIIP specialty engine, HostBridge Redis makes these apps even faster, easier, and less costly to deploy.

HOW HOSTBRIDGE REDIS DELIVERS

A use case illustrates how z Systems customers benefit. A claims company processes very high volumes of prescriptions. Written by thousands of health care providers, the prescriptions are fulfilled by pharmacies and dispensaries nationwide. The company must share the data they collect not only with the providers and dispensaries, but also with insurance companies, government agencies, and patients. As data requests multiply, so do CICS and DB2 transaction volumes and processing costs.

HostBridge Redis handles this processing volume more effectively. Prompted by DB2 triggers and CICS events, the Redis for z/OS in-memory cache is populated with system of record data and transactions. Requestors then access this data from Redis on z/OS or, thanks to Redis replication, from any other platform—mainframe or distributed.

The claims company gains two significant benefits. First, by sharing data more widely, they *reinforce* the value of z/OS as their system of record. Second, they lower costs by running Redis workloads on the zIIP, all while enabling new applications and uses.

FROM Z TO ANYWHERE

A guiding principle at HostBridge is that IBM z Systems customers gain the greatest return on their investment when they expose z/OS apps and data using the same high-performance, standards-based technologies used by today's cloud, analytics, and mobile innovators. HostBridge JavaScript (HB.js) and Redis—both z/OS based—adhere to this principle and are designed with one primary objective: put the power, flexibility, and economy of standards-based interoperability to work for z Systems customers. ■

HOSTBRIDGE TECHNOLOGY
www.HostBridge.com/redis