How to Fight the Big Data Avalanche

Real-time analytics on big data sets requires a new infrastructure approach—not a rehash of old methods

CONTENTS:

Real-time analytics is the cornerstone of a sustainable advantage for today's enterprises. Faster insights and their resulting actions increasingly make the difference between success and failure for many businesses.

But the continuing trend toward more data and faster analytics capabilities creates significant challenges for today's data workers: How do you act quickly enough to ensure the data insights are still valuable, and better than your competitors? What investment choices will help keep you ahead of this monumental data growth?

While software is essential to any big data strategy, it is not enough. Organizations must transform their infrastructure based on innovation in storage and memory to align with the increased need for faster, real-time decisions.

TechTarget Custom Media

A TECHTARGET WHITE PAPER

In short, infrastructure really matters. Why? Because organizations are losing money and missing competitive opportunities while they wait for data. It's essential that IT teams continue investing in the underlying infrastructure that underpins next-generation big data solutions. But it's not just a question of buying the right new hardware. Enterprises need to carefully optimize and tailor existing hardware systems to ensure that they have sufficient performance, capacity and throughput to handle data velocity and move to the prescriptive phase of analytics.

Of course, there are a number of hurdles that must be surmounted to make big data and sophisticated analytics workloads like Hadoop and Spark a reality—and among the biggest challenges are those involving the underlying data. Specifically, these center on the so-called "3 Vs" of data: variety, volume and velocity. While all three of these data characteristics are vitally important, velocity—the accelerating speed of data creation and decision-making—stands out.

Why is this focus on velocity and speed so essential to optimizing big data? It's due, in part, to the other two V-factors—variety and volume. Big data systems are capturing and processing far more than text; sensors, video, audio, social media and a wide range of unstructured data have the potential to create agonizing bottlenecks for in-depth analytics. Of course, the sheer volume of current—and especially future—data exacerbates that challenge. Without an intense focus on speed—and on the underlying storage and memory infrastructure—big data projects will fail to deliver on their full potential, at a significant financial and operational cost to the enterprise.

Faster Storage

Interviews with thousands of North American IT decision-makers from among TechTarget's buying audience note that the top storage features they are looking for in big data storage requirements are latency, capacity and bandwidth.[1] In particular, latency is a huge problem that must be properly addressed in big data initiatives. Latency essentially chokes the server's ability to process bigger and bigger data sets, bringing workloads to a grinding halt and making sophisticated analytics more expensive and more time-consuming to complete. The truth of the matter is that latency, particularly for big data and other advanced analytics workloads, matters far more than either capacity or IOPS.

The reality is that frame-based arrays—still the prevalent format for data center infrastructure—have reached their functional limit with spinning hard disks as the storage media. This has essentially maxed out the capabilities of many data centers in terms of real estate, power/cooling costs, CPU cycles and storage capacity. That, in turn, has forced organizations to open new data centers, usually at a seven-figure cost.

Increasingly, this is prompting a move toward flash/solid state storage. Innovative organizations are re-architecting their data centers to move storage off NAS appliances and SANs in favor of flash-rich servers. This is helping to reduce data center sprawl, decrease physical footprints, reduce power and cooling costs and drive overall operational efficiencies.

1 Interviews with IT buyers on big data storage purchasing intentions, Internal data from TechTarget, December 2015

"Most big data evangelists smile when you ask the best way to host big data and answer: 'On flash, of course!'" notes storage industry analyst Jon Toigo. "It's not surprising that database makers are reinventing their products today following the lead of Oracle and SAP, whose big data appliances feature a full suite of flash memory and dynamic RAM to host their in-memory databases."[2]

The move from Hadoop to Spark—either in collaboration with each other or with Spark replacing Hadoop—has put pressure on storage and memory infrastructure as data scientists search across blocks for insights. Spinning hard disks are simply too slow to keep up with iterative machine learning workflows necessary to extract critical insights. An investment in even hybrid arrays may save costs initially, but is likely to cause critical insight gaps as they will need to wait for hard disk drives to locate the necessary data.

Finally, organizations need to overcome latency in order to more quickly spot and act on real-time data. Obviously, workloads such as fraud detection, compliance, e-discovery and corporate governance all require the ability to find the needle in the haystack that might prevent a failed audit or a data breach. But public-facing applications such as OLTP, customer self-service, call centers, CRM, supply chain management and marketing automation also demand that combination of precision and speed. Hybrid systems are evolving that perform transaction processing and real-time analytics, for which latency of storage becomes critical.

Without the velocity provided by SSDs, big data analytics simply could not provide a sustainable competitive advantage for most applications.

Faster Compute

One of the critical steps in taking full advantage of faster storage systems is ensuring that the server's compute center can handle the higher utilization required to act on that data. Choosing the right balance of DRAM for the target application and workloads is essential to enabling a server configuration to hit performance targets without going over budget.

And now there's a new twist in the memory/storage hierarchy to consider—persistent memory. As the name implies, persistent memory sits on the DRAM bus and provides the extraordinarily low latencies of DRAM but with the assurance of data persistence. DRAM latencies are orders of magnitude better than even the fastest SSD, due to the limitations of the storage bus and the larger access patterns required for flash memory.

Today's persistent memory solutions are built with both DRAM and flash technology and are known as non-volatile DIMMs (NVDIMMs). The advantage to this new technology is the ability to finally move critical data—data that you cannot afford to lose in a power failure—as close to the processor as possible. The active work is done in the DRAM[3]—like a standard DIMM—but data can be quickly written to flash in the event of a power failure (capacitor or battery systems maintain power to the module long enough to complete the backup task). It's a compelling choice for applications that have frequently accessed or modified metadata or simply need the fastest performance possible on a given data set.

2 "Big Data Storage Architecture Combines Speed and Capacity," SearchStorage.com, August 2014

3 For the sake of simplicity we are describing just type-N NVDIMMs here, largely because they offer the best path to DRAM-like performance and are the best candidate for big-data applications. Other types use the NAND and DRAM differently.

In a recent article, a leading server OEM said that Microsoft's ubiquitous SQL Server database could run two to four times faster by utilizing NVDIMMs[4]—an enormous performance gain. But what's perhaps most striking about this claim is the fact that server systems are just in the early stages of NVDIMM adoption; performance will increase significantly again when operating systems and software applications have been fully optimized to take advantage of persistent data on the memory bus. The performance advantages NVDIMM offers, coupled with the competitiveness of the server market, are driving those changes quickly, though. Any IT administrator who wants to be on the leading edge of data analytics should be considering how to move the most time-critical tasks to NVDIMM-based persistent memory systems today.

Evaluating and Selecting a Full Portfolio of Storage and Memory for Big Data

Modernizing and refitting data center infrastructure for big data and other advanced analytics is essential today, and will become increasingly important—and challenging—in the years to come. But that does not necessarily mean a full-scale rip-and-replace overhaul. Instead, smart IT decision-makers are looking for new storage and memory options that can fit seamlessly into legacy environments, as well as become the basis for new infrastructure-optimized for analytics.

Micron—a world leader in silicon-based storage and memory solutions—has a proven track record. Its broad line of SSDs, DRAM and NVDIMMs makes it an ideal partner for infrastructure administrators who need fast, reliable solutions for big data. In particular, Micron offers a full line of products designed to turn big data into "big and fast data," which has now become the basis for analytics-centric infrastructure.

Take SSDs, for instance. Micron has long been on the leading edge in developing and delivering a wide range of SSDs for new storage environments. One of the advantages in working with Micron's flash solutions is the company's intimate knowledge of the underlying flash storage media, because Micron makes and designs the media.

SSD densities have been climbing rapidly and are eclipsing those of HDD systems. With this capacity crossover, it makes more sense to buy flash storage than to stick with spinning hard disks. Additionally, the much lower power consumption of SATA-based SSDs compared with SATA-based HDDs is a big motivator for cost-conscious buyers.

Micron continues to push the envelope on flash storage with its latest 3D NAND solutions. This new technology triples the capacity previously available. With more capacity, faster performance and low power requirements, 3D NAND is purpose-built for the most demanding analytics workloads.

DRAM was Micron's first product more than three decades ago, and the company continues to be one of the world's largest suppliers. It offers a comprehensive line of technologies and form factors in DRAM, and works with every major server OEM. Micron also offers a full family of non-volatile DIMMs that combine the speed of DRAM and the persistent storage of NAND flash, as described above.

When choosing a supplier, it's also important to understand their long-range development and research skills. Micron is working on several storage and memory solutions that offer compelling advantages for future big data workloads. For instance, Micron's new 3D XPoint™ technology is an innovative non-volatile memory solution that delivers as much as 1,000 times lower latency and vastly higher endurance than traditional NAND technology, making it ideal for real-time data mining and insights. It has

4 "Hewlett Packard Enterprise Explores Persistent Memory for Chips," *Wall Street Journal*, March 29, 2016

potential to offer dramatic new features as either high-end enterprise storage or high-capacity persistent memory.

Micron is also finding innovative new ways to apply its technology to big data challenges. For massively parallel problems like bioinformatics, high-frequency trading, and fraud detection, sometimes a completely new computing architecture is warranted. Micron is developing in-memory data processing solutions—like the Automata Processor—that can solve these pattern matching problems much more efficiently than a computational processor, providing the promise of application acceleration that wasn't available previously. Both of these solutions are in early initial enablement, but show the company's commitment to delivering new solutions to fuel the next stages of big data computing.

Conclusion

As big data workloads become both more commonplace and more strategic inside enterprises, IT infrastructure has to evolve in dramatic and fundamental ways. Big data brings big opportunity for invaluable new insights, but only if systems can quickly and reliably gain access to individual morsels of data inside massive and rapidly expanding data stores.

Spinning hard disks, the staple of storage infrastructure for decades, are no longer sufficient to capture, index, parse, report and act on mountains of unstructured data. Instead, organizations are turning to SSDs and DRAM for in-memory processing in order to overcome performance-killing latency.

The best way to get the most out of any sophisticated analytics workload is to utilize infrastructure that is purpose-built for big data. Flash storage and DRAM solutions from Micron have been optimized to deliver the combination of low latency and high IOPS needed for this new wave of analytics solutions. Micron not only provides the performance necessary to overcome latency and processing bottlenecks, but it also enables organizations to deploy sophisticated new storage and memory infrastructure more cost efficiently than ever.

Combined with its history of silicon solutions leadership and its reputation for R&D excellence, Micron helps organizations achieve the goals of real-time data analytics, speed, reliability and performance in a scalable solution set.

For more information, go to
www.micron.com/solutions/enterprise-ssd-storage/big-data-and-analytics.