451 Research | Business Impact Brief

Automating Storage Tiers Can Drive Faster, Deeper Analytical Insight

The 451 Take

It is no secret that data volumes will continue to grow at an accelerated rate, particularly at the enterprise level. There is also a broad array of data types that enterprises collect, depending on the vertical in which they operate. These include data from internet searches, data uploads and downloads; social media sentiment data from photos, videos, tweets and messages; financial data from transactions and trades; communication-related data from emails, texts, profiles and messaging applications; and IoT data from embedded applications and connected devices. And the list goes on.

Enterprises generally want to capture as much data as possible because the data likely feeds analytical processing systems that these enterprises maintain. Systems may consist of business intelligence platforms for reporting and visualization and data warehousing environments, but enterprises are increasingly moving into AI-based analytics tools that leverage machine learning (ML) to carry out predictive and prescriptive analysis, and which run on specialized hardware such as GPUs.

For some enterprises, depending on the analytical workload, gigabytes of data may be sufficient. But with AI/ML-based analytical workloads, data requirements often move beyond gigabytes into petabyte territory and beyond. And therein lies the challenge for many enterprises. How do enterprises manage and aggregate petabytes of data while efficiently feeding GPUs that serve as the analytical backbone?

The answer lies in managing the storage tiers. Because if enterprises can optimize their storage, AI/ML-based workloads run more efficiently. Turns out that enterprises not only worry about the growing amount of data but also about how to effectively manage that data. In fact, 451 Research's recent Voice of the Enterprise: Storage survey (see figure below) reveals that 63% of respondents want to improve their storage efficiency, particularly as a way to reduce labor expenses, and are looking to AI as means to address this. Further, a more streamlined approach to storage indicates fewer incidents (tickets), improved remediation, and the ability to better meet RTO and RPO requirements.

Streamlining Data Management and Support with AI

Source: 451 Research's Voice of the Enterprise: Storage, Budgets and Outlook 2019

Q. How would your organization measure the value of machine learning/AI enhancements? (Select all that apply.)

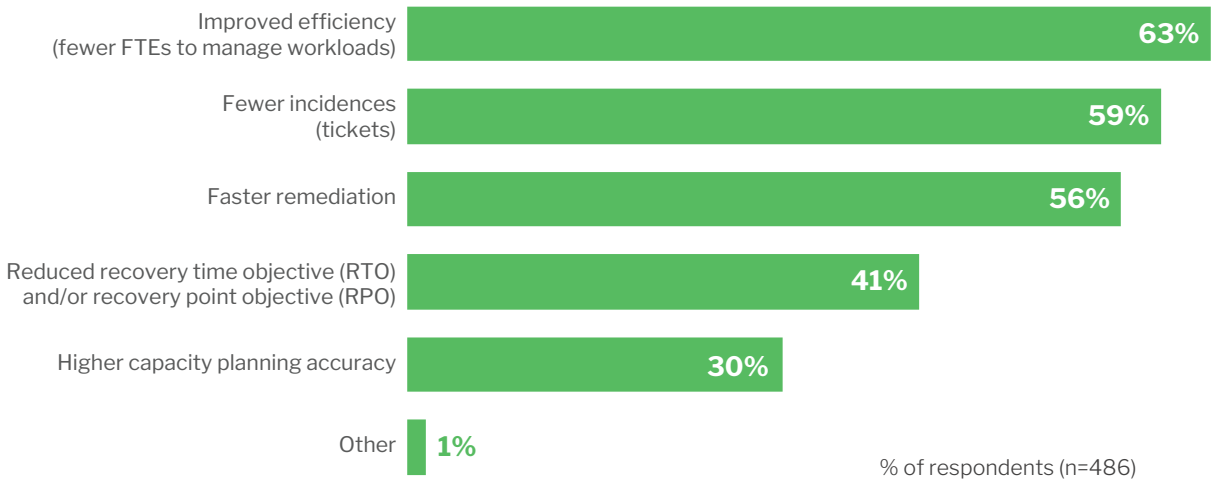

451 Research® | Business Impact Brief

Business Impact

REDUCING COSTS TO MANAGE THE STORAGE TIERS. In an ideal world, enterprises might prefer to load all of their data in memory or into pricey flash arrays for analytic processing, but that approach can be cost-prohibitive. A more cost-effective approach is to ensure that the data, regardless of where it resides—memory, flash or spinning disk—is optimized for the storage media on which it resides, thereby reducing an enterprise's overall storage costs.

KEEPING ANALYTICAL ENGINES FULLY UTILIZED. GPU hardware can be a significant investment for enterprises. Not having fully optimized or saturated GPUs essentially means that an enterprise's analytical investment is being wasted. But more importantly, the impact of underutilized GPUs means potentially longer analytical results, which can then delay business decision-making.

REDUCING PRE-PROCESSING OF DATA FOR ANALYTICAL PROCESSING. Storage-optimized systems not only need to feed GPUs at a rate that will keep them fully saturated and utilized, but in order to do this, the data coming off of storage needs to be prepared to be processed by the GPUs. Because GPUs process data in parallel, this approach ensures that the overall environment can deliver high throughput and low latency for AI/ML-based workloads.

CUTTING DOWN ON MANUAL EFFORTS FOR DATA MIGRATION. For many enterprises, data migration is an inevitable task because the data may be coming from other systems and needs to be moved to analytical systems for processing. Ideally, enterprises need an environment in which the storage tiering is managed automatically, effectively eliminating or greatly reducing the need for manual data-migration efforts.

Looking Ahead

Enterprises have become amazingly optimistic and open to adopting AI/ML initiatives, according to 451 Research survey data. And for good reason. Implemented correctly, AI/ML can provide a number of benefits for enterprises. For instance, the building of ML-based models can be leveraged for analytic processing such as predictive analysis, leading to improved decision-making capabilities. AI/ML can also be applied internally for optimizing storage environments, as illustrated in the figure above, with benefits that include lowering labor costs. Further, data volume is expected to continue growing at significant rates. And while this data growth may be a good thing for data-hungry AI/ML processing, there's still a challenge for enterprises in terms of effectively managing that data for analytic processing.

As enterprises grapple with some of these challenges, one place they should look to is the storage layer. Traditionally, the storage layer has simply stored data. But with the increased focus on analytics and the adoption of AI/ML workloads, the storage layer is expected to play a crucial part in the modern IT environment going forward.

While not as widely adopted, optimized storage environments provide a number of benefits for enterprises, the least of which is having a system to streamline workflows from end to end. Enterprises don't want bottlenecks within the environment. Optimized storage systems address those bottlenecks by providing an environment that operates in harmony by addressing data migration, scalability and availability, and also providing a seamless path for AI/ML analytic processing.