Workload Consolidation for Today and Tomorrow:
The Role of AI and Machine Learning

PROTECTING AND EXTENDING STORAGE ASSETS AS WORKLOADS CONSOLIDATE

As the mix of traditional enterprise applications and real-time analytics-driven modern application workloads rapidly evolves, pressure has intensified on storage infrastructure to evolve with it. This paper looks at how organizations should rethink their environments to maximize their IT investments and future-proof their storage infrastructure as workload compositions change.

Mainstream enterprise workloads have always demanded a lot from storage infrastructure. Capacity, performance, resiliency and data protection requirements have continued to expand as mission-critical applications like Oracle, SAP, Microsoft SQL and others extend their footprint and stretch the limits of storage systems.

Now, add to that a growing number of real-time analytics-powered workloads introduced into enterprise application portfolios, such as geospatial advertisements in retail marketing, risk modeling in financial services, supply chain error remediation in manufacturing and augmented, automated decision-making in healthcare.

Today, the consolidation of enterprise and modern applications is dominated by traditional applications, making up an estimated 80% of today's workloads. But that's changing fast, and in a big way. Although workload consolidation has been around for a number of years, the nature and mix of those workloads are rapidly evolving.

In the next several years, analytics-driven workloads will both increase in number and become increasingly important in the overall enterprise application portfolio. As the need for greater insights and more granular analysis of more and more data increases, real time analytics-powered modern workloads will become mission-critical, requiring the same level of capacity, performance, resiliency, availability and security traditionally demanded by enterprise applications.

DELL EMC

TechTarget | Custom Media

As a result, real-time analytics-powered modern workloads are expected to make up the majority of workloads over the next few years. And, of course, traditional enterprise workloads aren't going away, either. Applications like databases, ERP, CRM and decision support have long been essential to enterprises' core business operations and must continue to be supported with best-in-class infrastructure.

This is already placing a lot of pressure on storage infrastructure to evolve with this changing workload mix. Storage infrastructure must deliver the high IOPS, low latency and expandable capacity that go with demanding analytics-powered workloads. And this must be done in a way that protects CapEx investments over longer periods of time, while also ensuring a future-proof architecture that intelligently and automatically grows, evolves and adapts with changing workload requirements.

In short, storage infrastructure must become smarter, more adaptable and better attuned to the rapidly changing needs of analytics-powered workloads. Storage administrators and IT decision-makers now need a storage platform that anticipates and realigns itself as workload patterns change over time.

This paper looks at a new class of storage solution that delivers the investment protection and future-proof architecture required as more analytics-powered workloads dominate the enterprise application portfolio.

How real-time analytics-powered modern applications change demands on storage infrastructure

Digital transformation is more than an industry buzzword: It's the new paradigm based on leveraging the incredible depth, breadth and diversity of data to improve everything from how global supply chains are shortened to how to analyze population health.

Traditional applications continue to command substantial resources from IT organizations that need to ensure availability, performance and security of those core workloads. But next-generation applications are growing fast, and they demand an even higher level of performance from storage systems.

Not only do enterprises need to dramatically increase capacities to account for richer data formats and bigger data volumes, but modern applications also need higher IOPS and bandwidth, lower latencies and greater resilience. Storage infrastructure must accommodate all of those requirements, and must adapt to fast-changing workload characteristics and demands as portfolios become increasingly analytics-centered.

Another challenge is the need to deal with a dramatic uptick in the number of copies of data being created and needing to be managed across analytics workloads. This means that data compression and deduplication are now vital storage requirements, especially with the preponderance of copies created with modern analytics-powered applications.

Although many data centers have moved toward flash storage arrays to deal with this huge data influx, many all-flash arrays still fall short on the ultra-low latency and ability to learn autonomously from analytics-powered workloads.

At the same time, service-level agreements are becoming more stringent than ever in the face of higher and higher demands for resiliency, availability and hyperscale performance. This is also happening as enterprises continue to consolidate workloads on a single storage array, which can increase fault domain risks with traditional storage infrastructure designs.

As workloads continue to shift more rapidly toward analytics-driven applications, how will your storage environment change to keep pace now and in the future?

How artificial intelligence and machine learning make a new storage platform possible

Storage infrastructure has evolved in modest, predictable patterns over the years, initially driven largely by the need for massive boosts in capacity and later for the demand for increased IOPS. But those incremental improvements no longer are sufficient to accommodate the dramatic shift toward real-time analytics-driven workloads making up more and more of enterprise application portfolios.

This quantum leap in storage advances is made possible by artificial intelligence (AI) and machine learning (ML) integrated into next-generation storage infrastructure. A new storage

platform is emerging, one that ensures investment protection and future-proofing even while delivering impressive performance improvements and adapting to changing workload requirements.

This new class of storage infrastructure has adopted ML algorithms that ensure:

- Automatic, lightning-fast and intuitive tiering across different storage media.
- Optimized data reduction to cut down on unnecessary copies and improved storage efficiency.
- Efficient data placement across NAND-based SSDs and emerging media like Storage Class Memory (SCM).

The result is a modernized, analytics-driven storage platform that allows the automatic, intelligent scaling of performance across multiple servers, controllers and both physical and virtual storage and compute infrastructure. It also enables the capture and analysis of IO statistics for performance tuning without the need for manual monitoring and intervention by storage administrators.

High storage capacity is also delivered, but without the impact on performance at scale that often results with legacy storage infrastructure designs.

Finally, this new generation of storage platform uses real-time, rapidly evolving intelligence to drive better decisions on the optimal approaches for storing, managing and analyzing data.

The Dell EMC PowerMax Platform

Dell EMC industry-leading storage portfolio includes a platform optimized for this changing mix of workloads, especially those powered by analytics. Dell EMC PowerMax is a next-generation all-flash array, with a high-performance, low-latency, end-to-end NVMe architecture. It supports lightning-fast performance (up to 10 million IOPS), ultra-low latency (less than three hundred microseconds) and extremely high throughput (up to 150 gigabytes per second).

Most importantly, PowerMax is specifically designed for this changing application landscape shaped by analytics-powered

applications, because it incorporates a range of artificial intelligence features that enable the storage platform to "learn" from workload I/O patterns.

A built-in ML engine automatically places data on the correct media type—either flash storage or SCM[1]—leveraging IO pattern recognition. The ML engine gains insights from data patterns and responds to changes in real time, rebalancing the delivery of appropriate resources as necessary. It leverages predictive analytics and pattern recognition to analyze and forecast over 40 million data sets in real-time. PowerMax is also NVMe-oF-ready, which will support the extreme high performance required for both traditional enterprise and modern, analytics-powered applications.

An important feature in PowerMax's ability to ensure quality of service and high availability is its granular control of resource usage. Customers have the option of setting latency requirements by application, allowing vital server applications to operate at necessary performance levels and preventing a single application from usurping unneeded performance. Service levels can be set on a per-workload basis that reflect the business priority and according to the needs of the particular workload.

PowerMax is an ideal solution for the rapidly changing workload mixes of today and tomorrow, due in part to its efficiency enabled by inline deduplication, compression, space-efficient snapshots and thin provisioning.

Finally, PowerMax is optimized for any workload that demands mission-critical availability, self-service data protection, highly secure storage meeting data governance and compliance requirements, and simplified, automated management that reduces the impact on already burdened IT staff and storage administrators.

1 SCM coming in early 2019.

Conclusion

Enterprise workload mixes are changing fast due to the growing adoption of analytics-driven applications. But at the same time, that doesn't mean that traditional enterprise applications are being retired or replaced abnormally fast.

Organizations must acknowledge and account for this changing workload mix and put in place a storage architecture that accommodates both traditional and modern applications, and does so in a way that preserves capital investments over time and ensures a future-proof storage infrastructure.

This modernized infrastructure needs to embody the intelligence, automation and adaptability required with analytics-driven workloads, while continuing to meet the ever-changing demands for traditional enterprise workloads as well.

The Dell EMC PowerMax platform delivers the essential high performance and extremely low latency in an end-to-end NVMe storage design required for rapidly evolving mixed-workload environments increasingly shaped by analytics-powered applications. PowerMax's integrated machine learning engine and intelligent data placement are among several smart software capabilities that help organizations ensure a future-proof platform that delivers robust data protection, rock-solid security, high availability and scalable performance that keeps up with evolving workload requirements.

To learn more about the artificial intelligence-powered Dell EMC PowerMax platform and the future-proof storage loyalty program, please visit www.dellemc.com/powermax.