Enterprise Strategy Group™ by TechTarget

TECHNICAL VALIDATION

# Pure Storage FlashArray//XL

## Workload Consolidation without Compromising Performance

By Alex Arcilla, Senior Validation Analyst
Enterprise Strategy Group

March 2023

# Contents

# Introduction

This Technical Validation from TechTarget's Enterprise Strategy Group (ESG) documents our evaluation of the Pure Storage FlashArray//XL. We validated how the product supports a mixed workload without compromising overall storage performance.

## Background

According to Enterprise Strategy Group (ESG) research, 53% of respondents consider their IT environments more complex than they were two years ago. Two of the top ten reasons for this complexity can be attributed to higher data volumes (34%) and an increase in the number and types of applications used by employees (29%) (see Figure 1).[1]

**Figure 1.** Top Reasons for Complexity in IT Environment

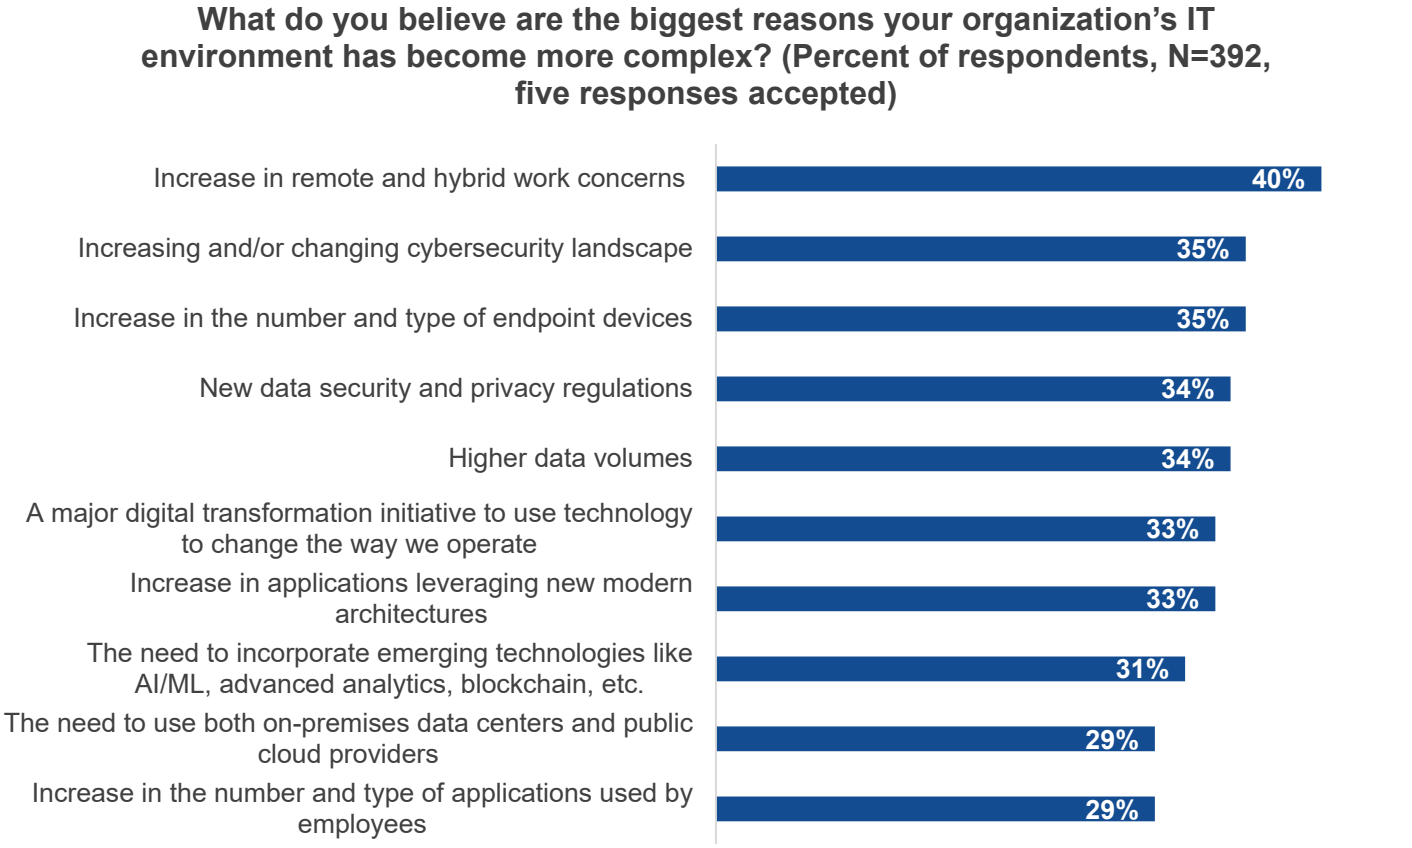


*Source: Enterprise Strategy Group, a division of TechTarget, Inc.*

It is no wonder that increases in both the amount of data used and the applications leveraging it can add unwanted complexity to the IT environment. Data that organizations must collect and manage to achieve their business objectives continues to grow. In fact, ESG research found that 52% of organizations anticipate their storage capacity for primary/active data to grow between 11 and 50% annually.[2] At the same time, more applications are

[1] Source: Enterprise Strategy Group Research Report, *2023 Technology Spending Intentions Survey*, November 2022. All Enterprise Strategy Group research references and charts in this technical validation have been taken from this research report, unless otherwise noted.
[2] Source: Enterprise Strategy Group Survey Results, *2021 Data Infrastructure Trends*, September 2021.

being deployed to help end users accomplish their daily tasks. Organizations face the possibility of purchasing more storage, leading to unwanted capital expenses. Moreover, operational expenses can also accrue, especially if organizations are managing and tuning storage to accommodate different application performance and availability requirements.

To minimize unwanted capital and operational expenses, organizations can seek out a storage solution that provides the current and future capacity required by multiple applications without sacrificing overall performance.

## Pure Storage FlashArray//XL

The Pure Storage FlashArray is a software-defined all-flash, all-NVMe, unified block and file storage array designed to offer the benefits of:

- ***Performance***: Low and predictable 150 microsecond to 1 millisecond of total latency for all supported applications. The self-optimized array can consolidate block (NVMe/NVMe-oF) and file (SMB/NFS) workloads with little to no performance impact. DirectMemory cache modules optimize read performance with Intel Optane solid-state drives (SSDs).
- ***High availability***: Designed to provide 99.9999% availability with built-in business continuity and disaster recovery to help organizations meet SLAs.
- ***Storage efficiency***: Highly efficient 5:1 data reduction and 10:1 total efficiency provide up to 5.5PB of effective capacity in a reduced footprint.
- ***Operational and management simplicity***: With the cloud-based Pure1 data management platform, organizations can use the single interface to monitor, analyze, and optimize any Pure Storage model from anywhere in the world, To reduce time and effort spent on daily operations, storage administrators can leverage REST application programming interface (API) automation to facilitate task execution. Furthermore, Pure1's support for artificial intelligence for IT operations (AIOps) and full-stack monitoring enables organizations to proactively prevent, identify, and resolve high-severity outages and other critical issues. And, as with previously released Pure Storage models, organizations can benefit from non-disruptive updates, upgrades, capacity expansions, and integrated and predictive support so that the impact of these business disruptions is minimized.

**Figure 2.** Pure Storage FlashArray//XL

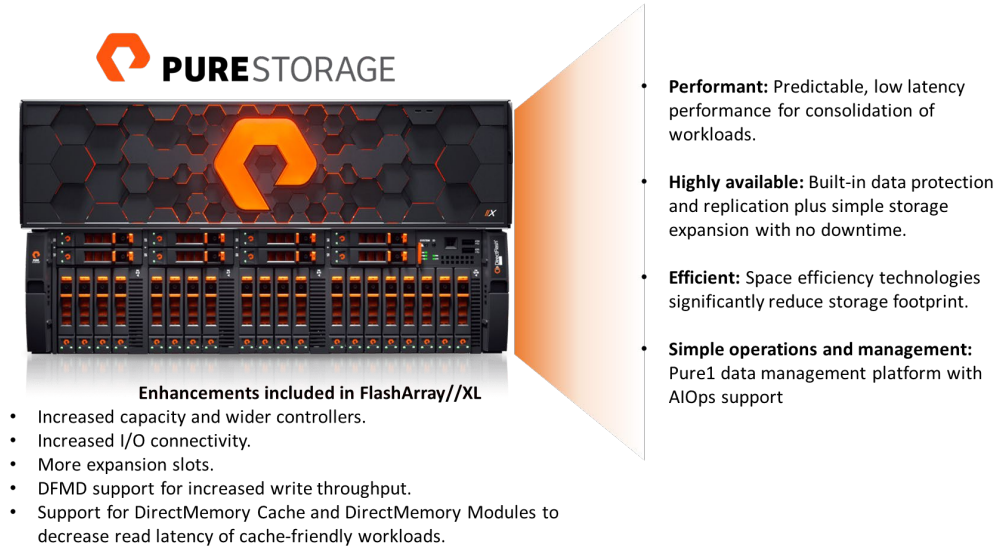


*Source: Enterprise Strategy Group, a division of TechTarget, Inc.*

The Pure Storage FlashArray//XL is designed to provide highest performance and density for business and mission-critical applications. Compared to previously released models, the FlashArray//XL offers:

- **Increased capacity and performance:** The larger five rack unit (RU) chassis accommodates higher-powered multi-core CPUs with more space for fans and airflow for improved cooling efficiency. Wider controllers enable FlashArray//XL performance to scale today, as well as into future generations of //XL.
- **Increased connectivity and improved redundancy:** The FlashArray//XL includes twice as many host I/O ports compared to FlashArray//X, up to 36 ports per controller. It also holds more expansion slots, each with twice the bandwidth, for configuration flexibility, including full bandwidth for mixed protocols. This array version also supports multiple 100GbE NVMe-over-TCP links to provide faster controller failover speed. greater reliability, and improved redundancy.
- **Support for DirectFlash Modules with distributed NVRAM (DFMD):** Increasing the number of DFMDs adds NVRAM capacity and write bandwidth as well as overall array capacity scale. Combined, these lift the limit on write throughput.

All FlashArray//XL devices ship with Direct Compress Accelerator, an FPGA-based card that offloads inline compression from the CPUs. The accelerator card has been designed to immediately react to different data types on the fly, adjusting the compression process based on the data. This is especially useful when consolidating workloads on the FlashArray//XL, as the card can help to free up processing power to boost performance, while further stretching storage capacity.

# Enterprise Strategy Group Technical Validation

Enterprise Strategy Group validated how the Pure Storage FlashArray//XL can run a mix of workloads associated with applications typically encountered in an enterprise IT environment. We observed live testing via remote demonstrations conducted in Pure lab facilities. During testing, we noted both the latency and bandwidth consumed as the Pure Storage FlashArray//XL supported individual and multiple applications simultaneously. Testing was conducted on the Pure Storage FlashArray//XL170, with a total available capacity of 968 TB.

We simulated these five workloads:

- **A virtual desktop infrastructure (VDI) workload** comprised of 12,000 desktops, generated by LoginVSI. Each desktop supported the Microsoft Office Suite (Outlook, Word, Excel, PowerPoint), along with Freemind (Java-based mind-mapping software), Doro PDF Writer, Internet Explorer, and 7-Zip (an open source file compression application).
- **A MySQL server workload** comprised of two host servers, with each host using a TPC-C[3] schema containing 5,000 warehouses. The workload simulated 102 virtual users (51 users for each host), with each user generating 1,000,000 transactions over a five-hour period.
- **An Oracle database workload** generated by the Silly Little Oracle Benchmark (SLOB) workload generator. Four host servers each contained a database with 128 schemas. Collectively, the four databases housed 2TB of data. The workload simulated 100 virtual users running transactions on each host.
- **A Microsoft Exchange workload** simulated using Jetstress 2013. A single host server contained five 165GB databases (for a total of 825 GB of data). A total of 12MB/s reads (with an I/O size of 256KB) and 970MB/s writes (with an I/O size of 10KB) was generated over 16 threads. This load simulated 110,000 end users.

[3] The Transaction Processing Performance Council - Benchmark C (TPC-C) is an industry-standard online transaction processing (OLTP) benchmark.

- **A Microsoft SQL Server workload** with 256 virtual users generated using the HammerDB, an open source database benchmarking application. The simulated database adheres to a TPROC-C[4] benchmark with users accessing 5,000 warehouses (approximately equivalent to a 1.1TB database size).

Performance was observed using the array's dashboard. Results were displayed in increments of five minutes. The X-axis in each figure within this report notes times in 15 second intervals of a displayed five minute period.

## Mixed Workload Performance on FlashArray//XL

As organizations use more applications and collect more data, the amount of storage will increase. However, supporting disparate storage platforms across multiple applications only adds to unnecessary complexity and capital and operational expenses.

With the Pure Storage FlashArray/XL, organizations can consolidate multiple workloads onto this single storage platform, simplifying both their IT environment and eliminating unwanted costs. More importantly, organizations can maintain high application performance, as the Pure Storage FlashArray//XL delivers millisecond storage latency across all workloads.

## Enterprise Strategy Group Testing

Enterprise Strategy Group (ESG) validated the performance of the FlashArray//XL by observing performance as individual workloads were loaded. We first saw both internal read and write latency (latency internal to the array) and bandwidth consumed by running the VDI workload.

During our observation, "read" bandwidth did not exceed 6 GB/sec (see top of Figure 3, depicted by the blue line). This was expected as the VDI workload was designed to simulate the use of multiple applications associated with knowledge workers. (Bandwidth associated with "writes" remained fairly low, never exceeding 2 GB/sec, as shown by the orange line).

However, both read (blue) and write (orange) latency rarely exceeded 0.30 milliseconds (see bottom of Figure 3). In fact, once this specific workload reached steady state, read and write latency, separately, hovered around 0.20 milliseconds.

[4] The Transaction Processing Benchmark derived from the TPC "C" specification (TPROC-C) is the OLTP workload implemented in HammerDB. Derived from the TPC-C specification, TPROC-C has been modified to make running HammerDB straightforward and cost-effective on any of the supported database environments.

**Figure 3.** Performance on FlashArray//XL – VDI only


*Source: Enterprise Strategy Group, a division of TechTarget, Inc.*

We then loaded the MySQL workload, on top of the VDI workload, onto the FlashArray//XL. As the dashboard displayed how the added workload affected the storage array's performance, we observed latency and bandwidth consumed. Again, "read" bandwidth did not exceed 6 GB/sec (see top of Figure 4), while bandwidth associated with "writes" remained fairly low, never exceeding 2 GB/sec. We should note that spikes in "read" bandwidth could be attributed to the workload simulating a ramp-up in end users sending queries.

At the same time, we saw that read latency began to go above 0.20 milliseconds, yet it did not exceed 0.30 milliseconds. We expected this as the FlashArray//XL was processing read requests from both workloads simultaneously. Write latency hovered around 0.20 milliseconds.

**Figure 4.** Performance on FlashArray//XL – VDI and MySQL


*Source: Enterprise Strategy Group, a division of TechTarget, Inc.*

ESG then observed the FlashArray//XL's performance as the Oracle workload was loaded. As shown in Figure 5, while bandwidth consumed by "reads" peaked beyond 6 GB/sec at times, it reached a steady state not exceeding 4

GB/sec. We also noted that except for a couple of spikes, read latency remained steady around 0.30 milliseconds, while write latency remained steady slightly above 0.20 milliseconds.

**Figure 5.** Performance on FlashArray//XL – Combination of VDI, MySQL, and Oracle

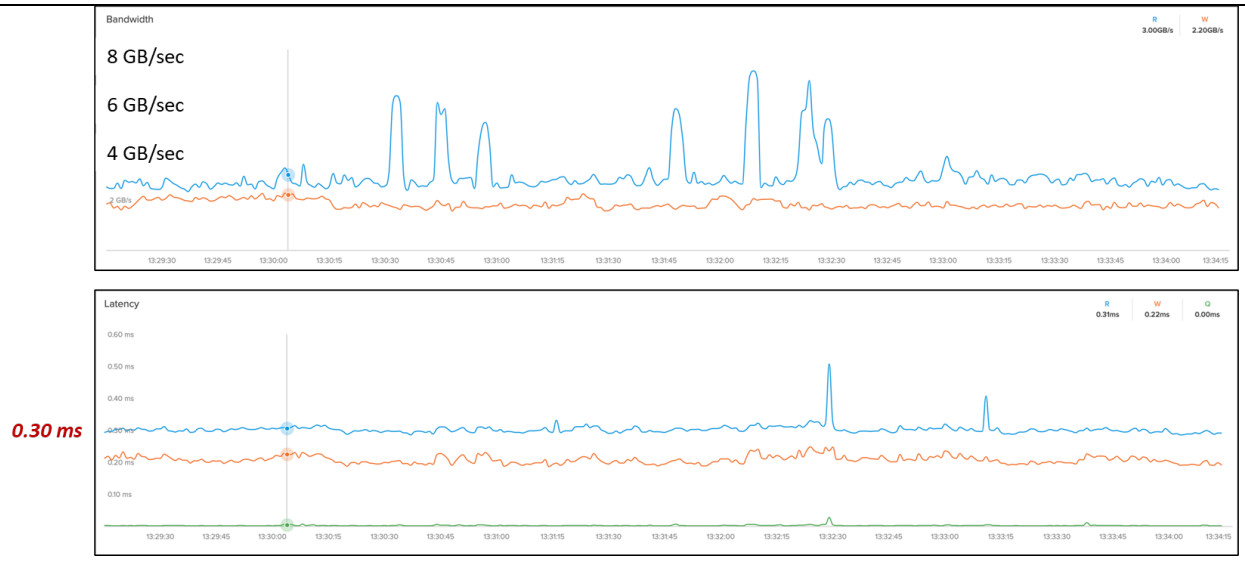


*Source: Enterprise Strategy Group, a division of TechTarget, Inc.*

Finally, the remaining workloads, Exchange and SQL Server, were loaded onto the FlashArray//XL. ESG observed the effects of the mixed workloads on storage performance, as displayed in Figure 6. While we saw that the combined read bandwidth for all applications spiked closer to10GB/sec, read and write latency increased. However, both latencies did not exceed 0.50 milliseconds (versus the earlier observations of 0.30 milliseconds).

**Figure 6.** Performance on FlashArray//XL with All Five Workloads

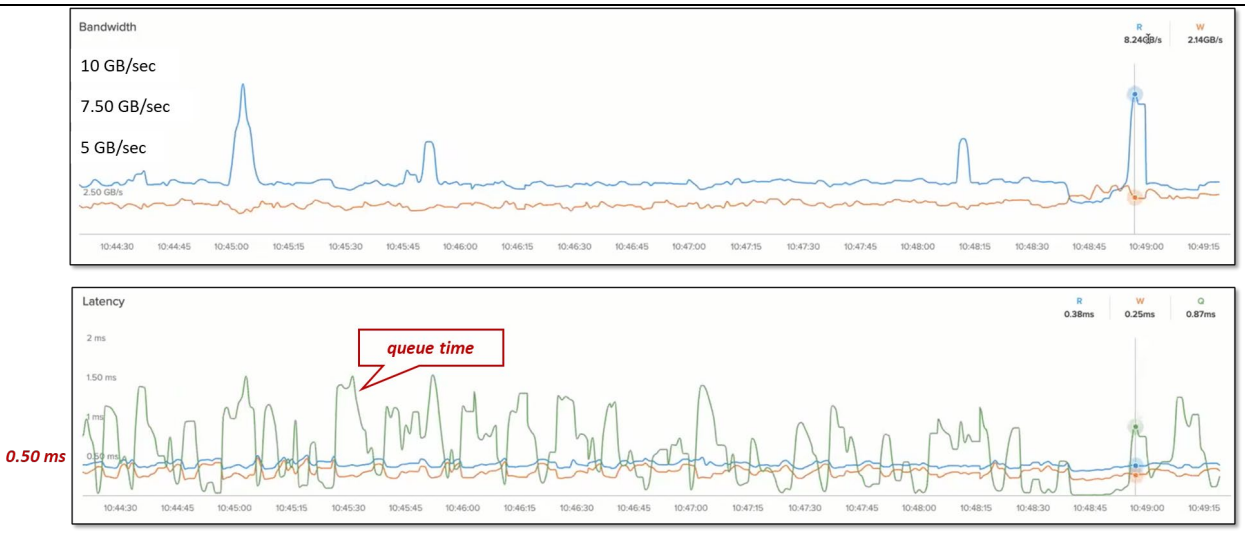


*Source: Enterprise Strategy Group, a division of TechTarget, Inc.*

ESG should mention that the queueing time for writes (shown by the green line in the latency graph within Figure 6) fluctuated and reached peaks above 150 milliseconds. The line charted the time a write request spent in the queue before being processed. This queue time accrued outside of the FlashArray//XL (e.g., request was waiting in the

storage area network connected to the FlashArray//XL). The write latency remained low and fairly constant, leading ESG to conclude that long queue times did not negatively impact the FlashArray//XL's performance.

### Why This Matters

As more organizations deploy applications, more data is being generated that requires more storage capacity. Yet, adding storage capacity to individual applications incurs additional capital expenses and operational complexity, especially when relying on disparate storage platforms. In parallel, organizations must ensure high performance consistently across all business and mission-critical applications. If using separate platforms, the constant management and tuning lead to unwanted operational costs.

Enterprise Strategy Group validated that the Pure Storage FlashArray//XL can support multiple workloads on a single array, without sacrificing individual application performance. We observed how Pure Storage FlashArray//XL delivered low read and write latencies—consistently around 0.50 milliseconds, with minimal deviation—as we loaded a total of five simulated workloads to the FlashArray//XL. Specifically, as each workload was loaded onto the FlashArray//XL, it maintained the high performance it was designed to deliver.

## Maintaining High Performance During Controller Failover

Failovers in storage arrays—due to scheduled maintenance or unscheduled outages—must be completed without sacrificing overall application performance. Otherwise, organizations risk damaging the end-user experience and preventing business tasks from being completed promptly. To address this issue, the Pure Storage FlashArray/XL has been designed to fail over to either of the redundant controllers without sacrificing overall performance.

## Enterprise Strategy Group Testing

ESG validated that performance recovered within seconds by performing a live software update on the secondary controller. All five workloads continued to run during this typical "scheduled maintenance" activity.

To kick off this software update, the secondary controller was rebooted via the Pure Storage FlashArray//XL CLI. ESG immediately observed that consumed bandwidth for both reads and writes decreased significantly (close to zero) during shutdown of the secondary controller (see top of Figure 7).

Within one 15-second interval, bandwidth consumption began to increase, signaling that read and write requests from all applications were being tracked and processed. During this time, the FlashArray//XL directed all requests to the running (primary) controller. At the same time, both read and write latency remained steady under 0.50 milliseconds. Latency continued to be recorded as the FlashArray/XL processed requests remaining in the queue after the reboot.

While ESG noted that some spikes in read latency occurred (indicated by the dotted line), read latency times eventually reached a steady state of around 0.50 milliseconds. Spikes could be attributed to a ramp up in read requests from a specific application.

**Figure 7.** Performance During Software Update of Secondary Controller

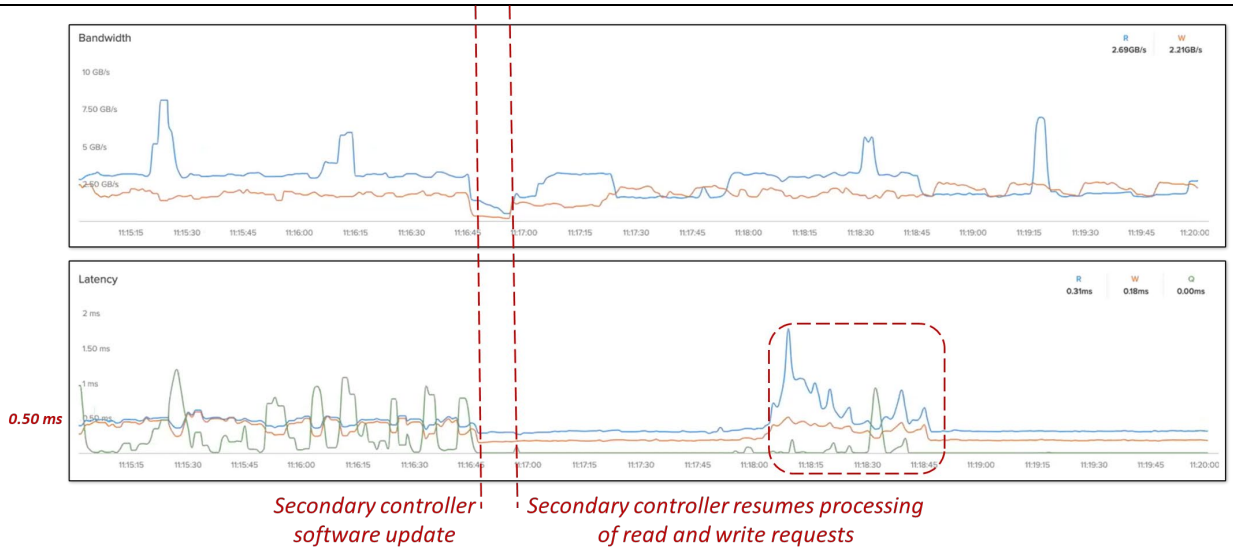


*Source: Enterprise Strategy Group, a division of TechTarget, Inc.*

ESG then observed FlashArray//XL performance during a forced failover (i.e., the primary controller was shut down) to simulate a hardware outage. Before this controller was shut down, bandwidth consumption was steady around 2.5 GB/sec (see top of Figure 8).

**Figure 8.** Performance Before and After Primary Controller Shutdown

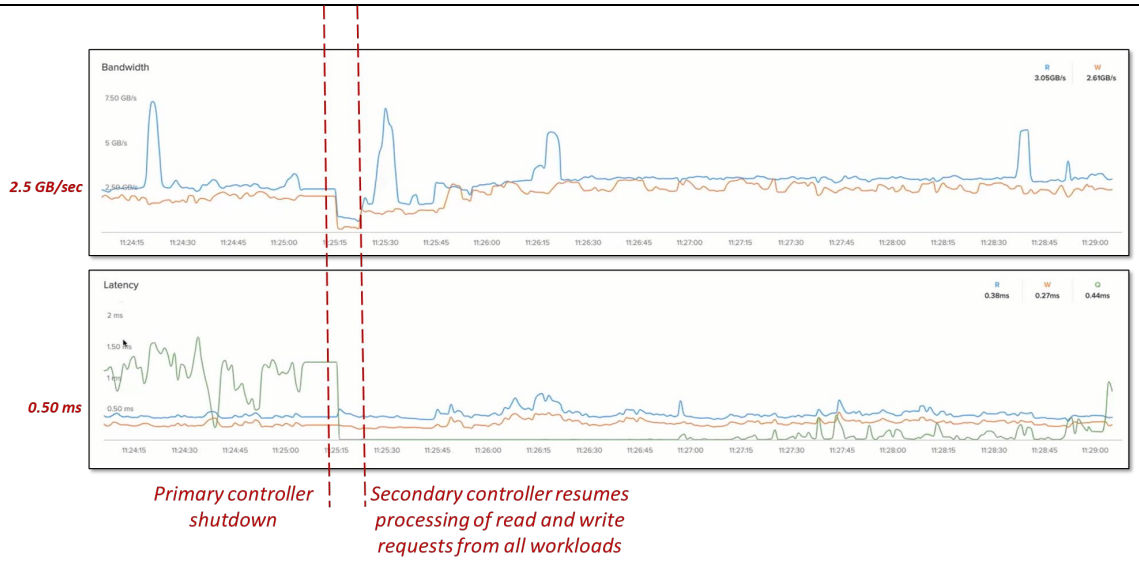


*Source: Enterprise Strategy Group, a division of TechTarget, Inc.*

Once the shutdown occurred, all read and write requests were directed to the primary controller. Less than 15 seconds elapsed before the read and write requests began to increase, eventually consuming up to 2.5 GB/sec. More importantly, ESG saw that read latencies remained steady around 0.50 milliseconds, while write latencies were consistently lower. We should note that the green line went down to zero after shutting down the primary controller, indicating that all initiators stopped sending requests, which allowed the array to process any requests remaining in the queue.

### Why This Matters

Whether scheduled or unscheduled outages occur (e.g., software upgrades, hardware failure), end users in an organization still expect to access the data they need to complete their work. Any extended downtime is unacceptable.

Enterprise Strategy Group (ESG) validated that the Pure Storage FlashArray//XL can recover from outages in noticeably short periods of time, without sacrificing performance. With all five workloads loaded onto the FlashArray//XL, ESG validated that the storage array maintained mixed workload performance around 50 milliseconds, with little deviation when observing scenarios—updating the software on the secondary controller and forcing a failover to the secondary controller.

# Conclusion

The amount of data that organizations believe they must collect and store increases daily. Couple this with the growing number of applications that organizations use to extract business value from the data, and the resulting IT environment can become complicated to operate and manage. While workload consolidation may be an option to decrease such complexity as storage arrays support higher capacities, organizations must ensure that individual workload performance does not degrade.

The Pure Storage FlashArray//XL is the latest version released as part of the FlashArray product line. As with previous FlashArray releases, the FlashArray//XL has been designed to offer benefits of high performance, availability, and efficiency. With the FlashArray//XL, organizations can consolidate workloads without adversely affecting the performance of any individual workload, as the array can deliver high and predicable storage performance. With higher controller port counts and increased slot bandwidth, the FlashArray//XL can support higher bandwidth consumption, especially in mixed workload scenarios. Enhanced redundancy capabilities help to maintain overall high availability. Higher data reduction capabilities enable organizations to extract up to 5.5PB of effective capacity within a reduced footprint. To simplify daily operations and management, Pure Storage offers the cloud-based Pure1, now equipped with AIOps capabilities so that organizations can proactively prevent, identify, and resolve high-severity outages and other critical issues.

Throughout our evaluation, ESG validated that the Pure Storage FlashArray//XL can deliver high storage performance when supporting multiple workloads simultaneously. We specifically observed how the FlashArray//XL delivered low and consistent read and write latencies when:

- Individually loading the FlashArray//XL with five simulated workloads representing typical enterprise workloads with multiple users—database, VDI, Exchange—on top of each other.
- Simulating both a controller software update and forced failover to the secondary controller while all five workloads continued to operate.

Except for some random spikes in latency, ESG observed no degradation in storage performance. When the FlashArray//XL supported all five workloads, read latencies remained around 50 milliseconds, with little deviation, while write latencies remained consistently lower.

ESG has had the opportunity to validate the performance of previous FlashArray releases. We have observed how the array has evolved to help organizations simplify their storage environments without sacrificing performance and this latest validation is no different. If you seek to simplify your storage environment in light of the continued explosion of data and growing number of application required to extract value from that data, we urge you to place the FlashArray//XL on your short list of evaluations.

**About Enterprise Strategy Group**
Enterprise Strategy Group is an integrated technology analysis, research, and strategy firm that provides market intelligence, actionable insight, and go-to-market content services to the global IT community. © TechTarget 2023.

contact@esg-global.com
www.esg-global.com