

Vexata VX-100 Scalable Storage Systems

Shared storage with in-memory responsiveness for database and analytics workloads

Date: October 2017 Author: Tony Palmer, Senior ESG Lab Analyst

Abstract

This ESG Lab Review documents validation testing of the Vexata VX-100 family of scalable storage systems with a focus on production transactional and analytics performance and enterprise-readiness for online applications.

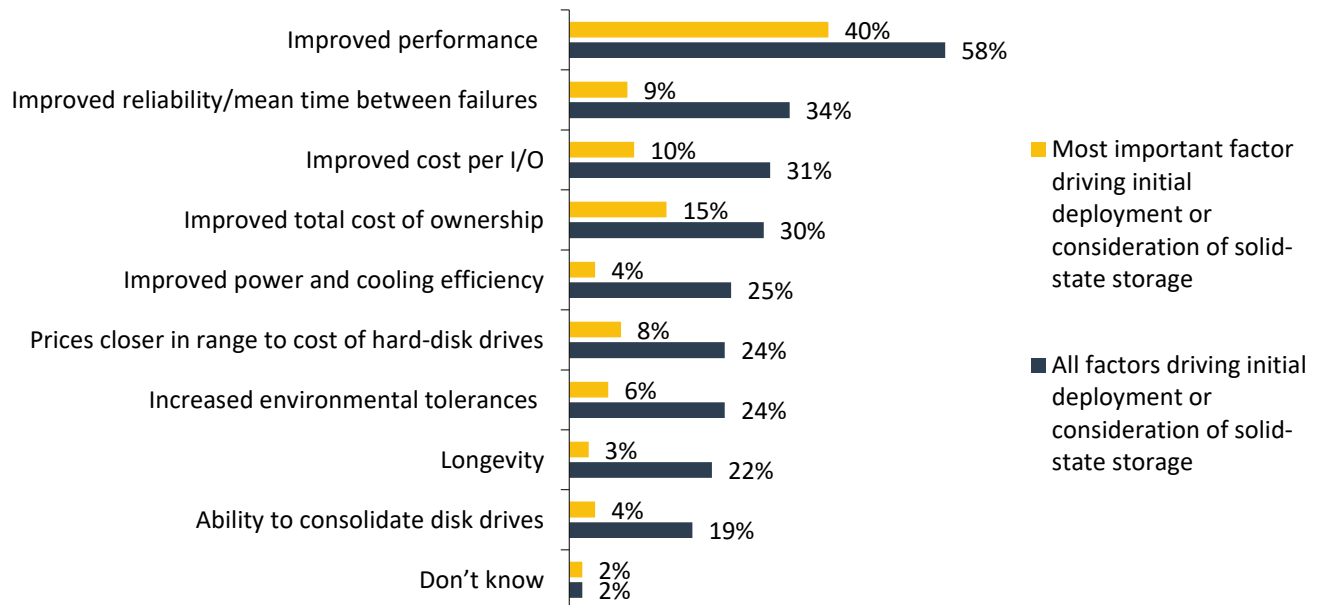
The Challenges

The move to digital business puts new performance and scale demands on database and analytics platforms. Systemic growth drives more sessions and transactions, which drives larger data sets. New user models, e.g., mobile, drive expectations of real-time performance and real-time notifications from applications. And users now expect to interact with and manipulate the increased amount of data at the core of these applications in near real time. These expectations combine to put focus on a key infrastructure characteristic, performance.

ESG research shows that improved performance is the most cited factor responsible for organizations' deployment of solid-state storage, selected by a majority of respondents—58%—when they were selecting all of the factors driving their organization's deployment.¹

Figure 1. Factors Driving Solid-state Storage Usage/Consideration

To the best of your knowledge, which of the following factors were responsible for your organization's initial deployment or consideration of solid-state storage? (Percent of respondents, N=309)



Source: Enterprise Strategy Group, 2017

¹ ESG Brief, [Flash Storage: Growth, Acceptance, and the Rise of NVMe](#), September 2017.


The Solution: Vexata VX-100 Scalable Storage Systems

The Vexata VX-100 Scalable Storage Systems are enterprise-class all solid-state block storage systems designed to meet the I/O, throughput, and response time requirements of database and analytics applications. The Vexata VX-100F system uses NVMe SSDs for storage and the Vexata VX-100M system uses Intel Optane 3D XPoint modules for storage. The Vexata architecture is designed to enable users to take advantage of the performance characteristics of the back-end storage media with both systems. Vexata’s stated performance for both systems is up to seven million IOPS at 220 microsecond latency for the flash-based VX-100F and 40 microsecond latency for the 3D XPoint-based VX-100M. This is accomplished without drivers or changes to the host or network stack and allows Vexata to deliver line rate read and write traffic across all 32Gbps Fibre Channel Ports.

The systems are 6RU appliances that deliver high availability and scalability (see Figure 2). There is no single point of failure in the system—the dual controllers run in an active-active mode and there are dual redundant power supplies. Both systems can be configured with from four to 16 hot-swappable Enterprise Storage Modules (ESMs). Each ESM contains 4 NVMe SSDs: Flash SSDs for the VX-100F and Intel Optane 3D XPoint SSDs for the VX-100M. ESMs can be added non-disruptively to scale capacity and throughput. RAID 5 or RAID 6 protection delivers data availability for the back-end capacity. The VX-100F configurations start at 20 TB capacity and scale to 180 TB useable while the VX-100M configurations scale between 2.25 TB and 34 TB usable capacity. Host connections are via up to 16 ports 32 Gb/s Fibre Channel—which was the basis of ESG Lab’s testing in this report—or 16 ports of 40GbE NVMe over Fabric, which was not tested by ESG Lab.

The VX-100 Systems are based on VX-OS software that combines concepts from advanced networking systems and distributed storage. VX-OS software is a lockless, user-space design that runs on three separate processing planes: Control, Router, and Data. VX-OS Control performs I/O command processing and data services, including high availability failover, thin provisioning, and space-efficient snaps and clones. VX-OS Router is embedded into FPGA firmware to handle cut-through I/O distribution to the Storage Modules, RAID 5 or RAID 6 data protection processing, data-at-rest encryption, and system-level garbage collection. VX-OS Data runs on each Storage Module processor and handles SSD I/O scheduling and metadata management. VX-OS management interfaces for the system include a GUI, a CLI, and a REST API. The Vexata systems can call home for support and provide detailed I/O performance and usage analytics.

Figure 2. Vexata VX-100 Scalable Storage Systems



The Vexata VX-100 systems are designed for enterprise use cases including high transaction databases, business intelligence, big data, and real-time analytics. The VX-100F’s NVMe SSDs offer higher capacity while the Optane 3D XPoint

modules allow the VX-100M to approach in-memory performance and latency while providing the traditional benefits of SAN-connected shared storage.

ESG Lab Tested

ESG Lab recently tested the Vexata VX-100M 3D XPoint Optane system in Vexata’s San Jose, California facility, evaluating the performance and enterprise readiness of the platform. Performance testing was designed to validate the claims that the VX-100M provides sufficiently high IOPS and low latency to virtually eliminate I/O waits, while delivering massive bandwidth to support real-time analytics with no I/O impact to other workloads, all while using industry-standard storage connectivity and drivers. ESG Lab also audited previously conducted testing of Vexata’s VX-100F NVMe flash system for comparison to the VX-100M.


Performance

The performance test environment included four dual socket x86 servers, each with dual 28-core Intel Xeon Scalable processors and 512 GB of memory, running Oracle Real Application Clusters (Oracle RAC) 12c. Each server was connected to a single VX-100M system using two dual port Emulex 32G FC HBAs per server via dual Brocade G620 32G FC switches. Testing used a 5TB database configured with several different types of tables and numerous record and population sizes. Workloads tested included: a complex OLTP workload with several concurrent transactions of differing types, an analytics workload with a combination of high data ingest and query processing, and a hybrid transaction processing and analytics (HTAP) workload that combines the two. Performance was tested using HammerDB, an open source database load testing and benchmarking tool in addition to the SLOB (Silly Little Oracle Benchmark) test tool.

It’s important to note that testing was designed to demonstrate the ability of the Vexata VX-100M system to support realistic transactional and analytics workloads at extremely low latencies, and not to validate the theoretical upper limits of VX-100 systems’ performance.

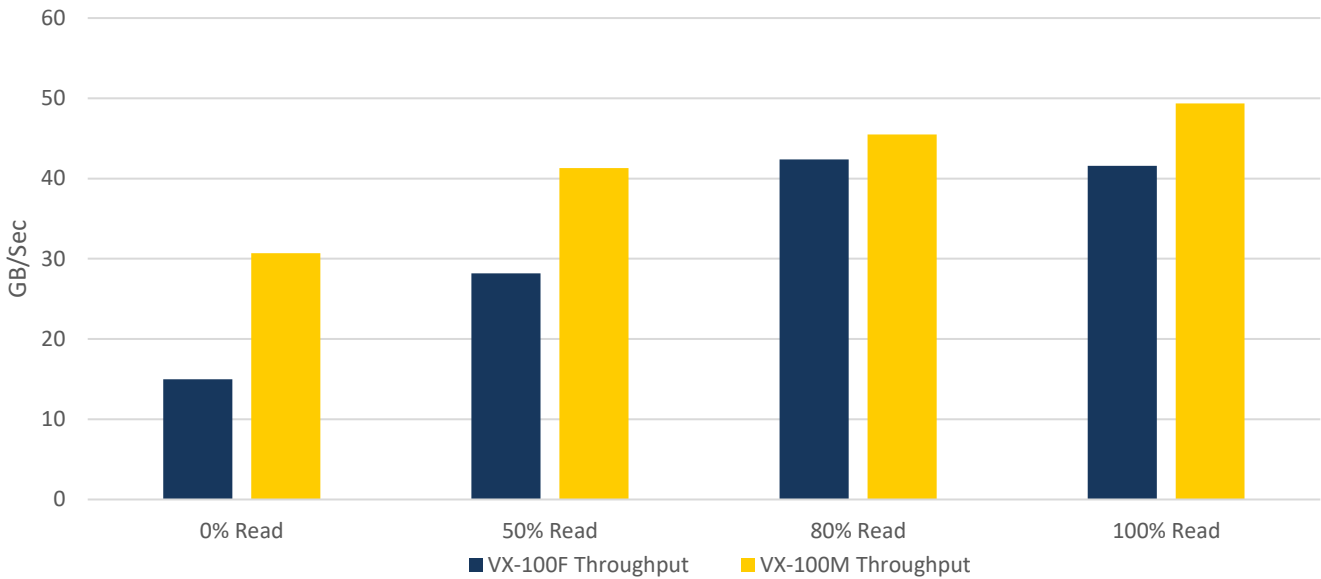
Figure 3 shows the test bed utilized by Vexata and ESG Lab. It’s also important to note that no custom drivers or interface cards were required or used in any of the testing and all testing referenced in this report for both the VX-100F and the VX-100M was executed with just four servers. All test results were cross-checked between the Vexata UI, Oracle AWR reports, and the output from the benchmark utilities used to ensure accuracy and consistency.

Figure 3. The Vexata VX-100M Test Bed



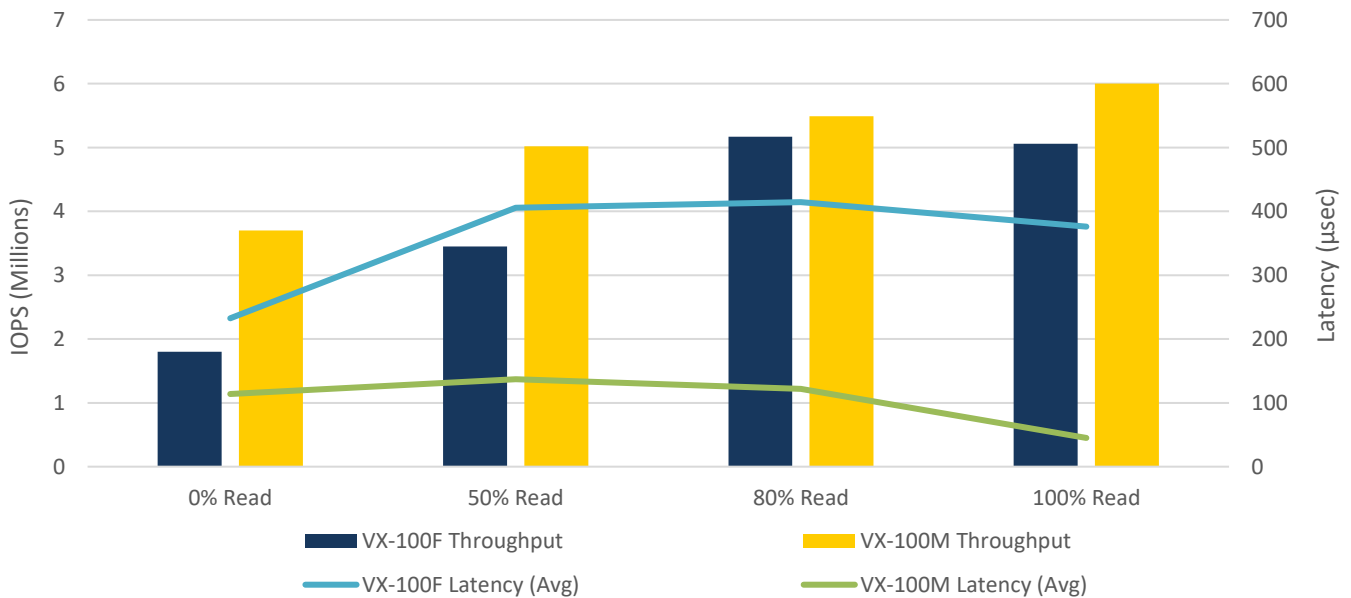
First, ESG Lab compared the throughput of the VX-100M with the VX-100F running a synthetic OLTP workload using Vdbench. Figure 4 shows the difference in throughput achieved using varying levels of read/write ratios. The VX-100M achieves nearly double the throughput of the VX-100F when servicing 100% writes, but throughput begins to converge as the percentage of reads increases. It’s important to note that at an 80% read/20% write ratio both systems are providing more than 40GB/sec of throughput.

Figure 4. Vexata VX-100 Systems—Synthetic Workload Throughput Comparison



Next, ESG Lab looked at IOPS and latency. This is where things began to get interesting. As seen in Figure 5, the delta between the number of IOPS serviced by the systems looks very much like the delta in throughput, but when we examine the response time at the system, we see that the VX-100M is servicing I/O at much lower latencies. To be fair, the VX-100F is servicing up to 5.17 million IOPS with an average latency of 414.5 μ sec, an outstanding result for any all-flash array, but the VX-100M is operating on a completely different level, servicing 6 million IOPS with an average response time of just 45 μ sec.

Figure 5. Vexata VX-100 Systems—Synthetic Workload IOPS Comparison



The implication here is that while the VX-100F will provide more than enough performance for most organizations' I/O requirements, the VX-100M is an entirely different beast, and can approach performance in the realm formerly occupied solely by in-memory databases.

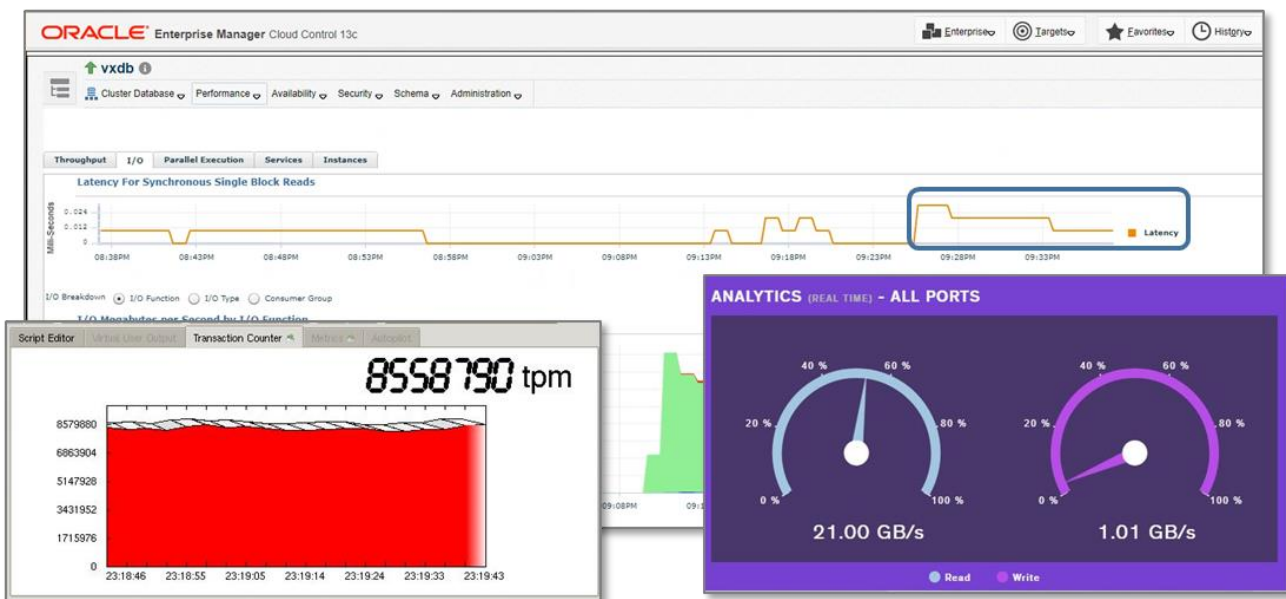
Next, ESG Lab looked at the performance of the VX-100M with an HTAP workload. We used Hammer DB to drive an OLAP analytics workload and SLOB (Silly Little Oracle Benchmark) to drive an OLTP workload against a 5TB Oracle database.

Figure 6. Vexata VX-100M HTAP Workload—OLTP, Data Ingest, and Query Processing



As Figure 6 shows, the VX-100M serviced an average of 54.26 GB/sec of I/O, with an average latency of approximately 40 usec at the array. It's important to note that this test simulates multiple workloads running simultaneously against the same data set on the VX-100M system. A complex OLTP workload with several concurrent transactions of differing types is combined with an analytics workload with a combination of concurrent data updates and queries run against the same database tables simultaneously. Figure 7 shows the HTAP workload test from three different points of view. In this instance, the HammerDB benchmark utility was used to generate the two workloads, and the performance was validated both at the array, and from the applications.

Figure 7. Vexata VX-100M HTAP Workload



As seen in Figure 7, the VX-100M serviced more than 8.5 million transactions per minute, while an OLAP query workload consumed more than 21 GB/sec of throughput against the same database. The response time at the host—reported by Oracle Enterprise Manager—ranged from just **12 to 30 µsec**.



Why This Matters

High performance OLTP and data analytics are not esoteric requirements for a niche market any longer. With 69% of respondents to recent ESG research having already deployed or planning to deploy the technology,² solid-state is decidedly mainstream. While many of its applications are exotic—genomics, high performance computing, and weather modeling—data analytics are needed for tasks such as keeping up with customers' online transactions, handling seasonal workload increases, and simply enabling organizations to outmaneuver their competition.

ESG Lab validated the extreme performance of the Vexata VX-100 series running OLTP and analytic workloads against a 5TB data set in an Oracle RAC environment and found that while the VX-100F provided outstanding performance for an all-flash system, the VX-100M took performance to the next level, providing memory-class IOPS and response times while servicing OLTP and analytics workloads against the same database at the same time.

The Vexata VX-100 family eliminates the need for proprietary storage networks and connectivity and enables IT to spend more time and money on productive business efforts. Companies get faster time to information—the actual answers to their questions—with a vastly simpler application environment and a storage solution that packs a huge punch in a small footprint.

Enterprise-ready

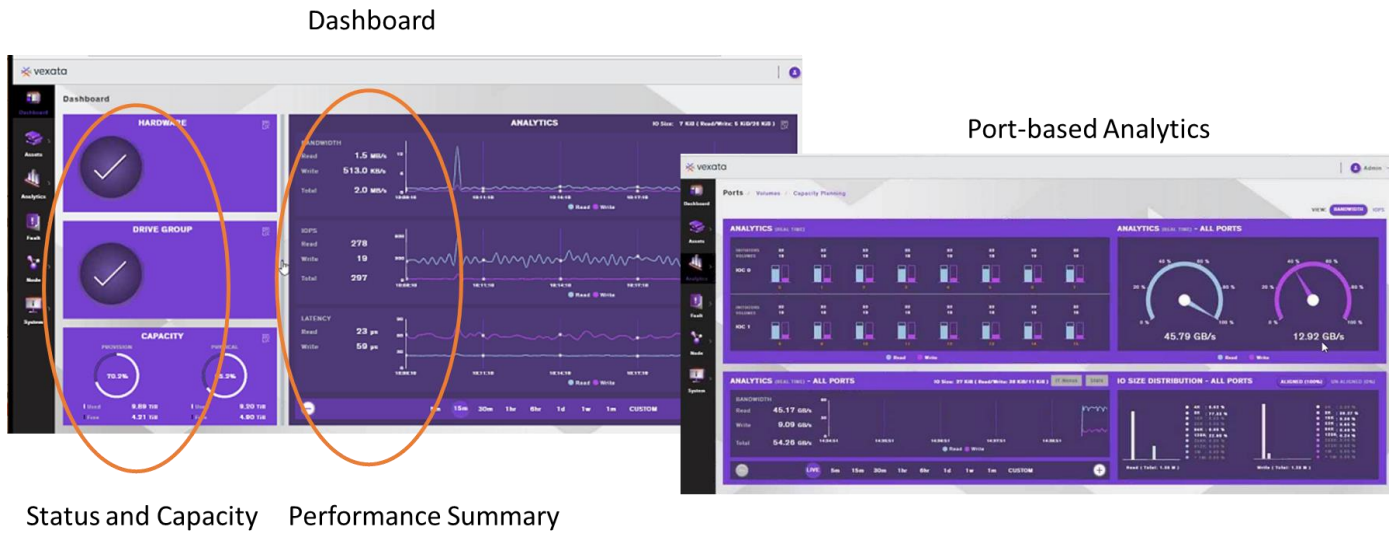
Next, ESG Lab explored the management GUI and introduced failures to verify system resiliency. The Vexata VX-100 systems are designed to deliver high performance while meeting the deployment, manageability, and resiliency expectations of enterprise users.

The VX-100 systems can directly plug-and-play with a wide range of enterprise IT environments. The systems connect to block storage hosts via Fibre Channel and Vexata has certified its Fibre Channel target with Linux, Solaris, Windows, and ESX operating systems and hypervisors. The VX-100 systems can present LUNs to hosts as traditional shared storage over the Fibre Channel connections. ESG Lab verified the Fibre Channel connections to Linux servers in our test configurations. Some of the performance tests used a four node Oracle RAC cluster, which showed that the systems can deliver the shared storage required by Oracle RAC.

The Vexata systems include the requisite management and data services. ESG Lab opened the GUI dashboard, seen in Figure 8, to view hardware and software configuration information and a visual display of key performance metrics. We clicked to an example analytics display, the Ports screen, also shown in Figure 8. The systems provide basic data services including thin provisioning, snapshots, and clones. 256-bit encryption is performed inline in the I/O path. It's important to note that we were unable to measure any impact in our tests as the Vexata systems encrypted data at line rate.

² ESG Brief, [Flash Storage: Growth, Acceptance, and the Rise of NVMe](#), September 2017.

Figure 8. Simple Dashboard and Analytics



Status and Capacity Performance Summary

Enterprise users demand high availability and data protection for the transactional and real-time analytics workloads intended for the VX-100 systems. The systems have no single point of hardware failure, with redundant I/O controllers and power supplies and RAID 5 or RAID 6 protection across the 4-drive ESMs. The systems can be maintained non-disruptively to handle failures, change configuration such as adding capacity, or to update software.

ESG Lab executed a controller failure and restore, examining the impact on the system GUI dashboard. We ran a Vdbench workload while we made hardware changes. In Figure 9, we removed a controller, triggering an alert on the GUI Dashboard. We clicked to the *Hardware Details* screen, which graphically indicated the failed component. The Analytics screen showed the impact of the failure and repair process, with the system throughput dropping about 10%, slowing momentarily while the replaced and old controllers synchronized, and returning to full performance after the re-sync.

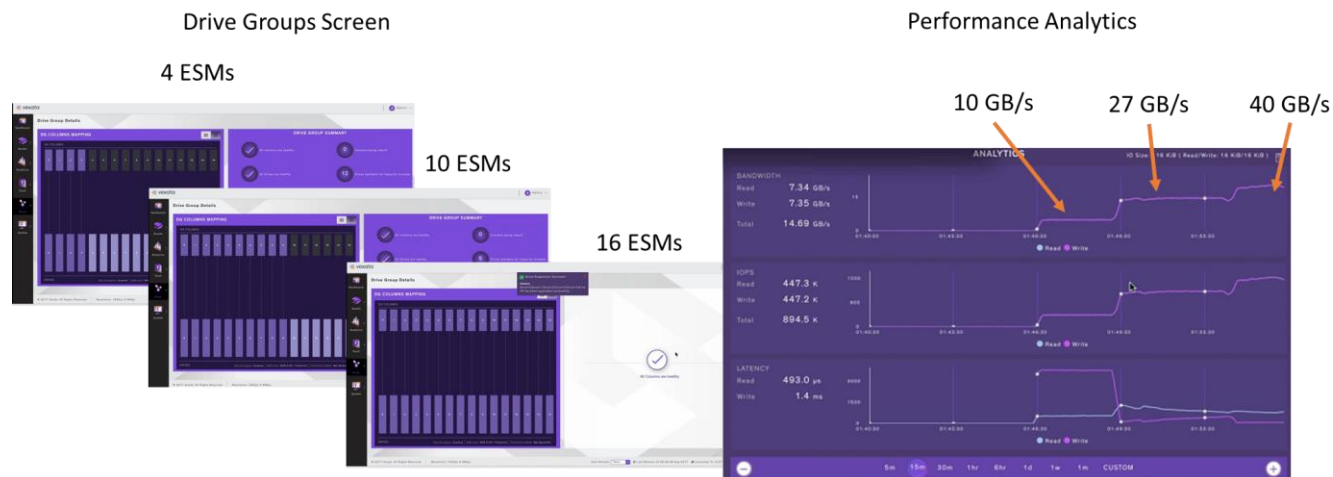
Figure 9. I/O Controller Failure and Recovery



In a similar manner, ESG Lab added capacity to a VX-100 system while it was running a Vdbench workload (see Figure 10). Using the Drive Group Details screen, we added Columns, i.e., ESMs consisting of 4 drives, to the VX-100 chassis. The system initially had four ESMs. We then added six ESMs and finally added another six ESMs. The Analytics screen showed that

Vdbench is a drive-limited workload, because performance increased each time we added ESMs, from 10GB/s throughput with four ESMs, to 27GB/s with 10 ESMs to 40GB/s with 16 ESMs in the system.

Figure 10. Online Non-disruptive System Expansion and Upgrade



ESG Lab also removed an ESM and watched the system rebuild the drive and return to full performance, while running the Vdbench workload.



Why This Matters

Enterprise users remain responsible for the mundane, but mission-critical, task of making IT infrastructure work, even as vendors battle over features and business considerations keep pressure on budgets. Minimal requirements for all infrastructure, old and new, include interoperability, manageability, and resiliency. After a product meets those requirements, users can understand and evaluate other product characteristics. New products that interoperate with already installed products save money by avoiding integration costs and increase the ROI of already purchased gear. ESG research also shows that data protection remains the most commonly identified challenge across storage environments.³

The Vexata VX-100 all-flash storage systems deliver the interoperability and robustness required by enterprise users. The VX-100's Fibre Channel interface and support for Linux, Solaris, Windows and ESX means it will integrate into any 21st century data center. ESG Lab found as little as 10% performance degradation under some failure conditions, almost linear performance improvement for drive-bound applications as capacity was added, and no measurable impact of encryption at line rate. The VX-100's straightforward GUI and ability to continue high performance processing under hardware failures, and to repair faults and upgrade online, helps enterprise users meet SLAs.

³ ESG Brief, [2017 Storage Trends: Challenges and Spending](#), August 2017.

The Bigger Truth

Legacy infrastructures were not built for today's applications, data sets, or even storage media, and they are simply not able to keep up with modern performance requirements. Organizations are using in-memory computing and applications like NoSQL, Hadoop, and Splunk to query petabytes of data that are stored on SSDs. Adding flash to traditional arrays or servers delivered incremental improvements—but only with a redesign of both hardware and software can you make full use of the performance capabilities of NVMe and Optane flash.

While “application latency” may sound like a boring IT metric, in fact latency is the key to application performance. Speeding up processing of large data sets and the ability to process them in place can mean significant savings. As an example, the average airline flight generates 500 GB of data, so major airports that handle upwards of 1500 daily flights need to process 750 TB (or more) of data each day to maintain flight safety and keep things moving.

Vexata VX-100 offers the benefits of both lower server-side latency *and* shared storage. It's designed for workloads that require high transactional performance and leverage large data sets. It delivers the performance of NVMe and Optane storage to hosts via industry-standard networks and protocols, with enterprise-class data protection. This enables organizations to extract value from their production data sets—in place—and to do more analysis with less tuning. It means organizations can get to the answers they are seeking much faster, with fresher data.

ESG Lab validated that the Vexata VX-100M can leverage its massive bandwidth to run complex OLTP and analytics against large data sets without having to extract production data for a data warehouse. The VX-100 demonstrated that it can perform completely non-disruptive, online capacity upgrades and survive a controller failure with minimal impact to I/O.

ESG Lab was extremely impressed with the Vexata VX-100 series. The hardware and software architectures are built specifically to leverage NVMe and Optane SSDs to provide remarkably low response time; as a result, the VX-100 enables levels of performance and data protection that other all-flash or hybrid solutions simply cannot match. Vexata VX-100 provides greater problem-solving capabilities than have been possible in all-flash arrays by orders of magnitude. If your organization is looking to optimize transactional processing and data analytics while reducing your infrastructure footprint, ESG Lab recommends taking a close look at Vexata VX-100 series.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.