

# Lab Validation Report

## Kaminario K2 Hybrid Solid-state SAN

Simply Fast Storage with DRAM and MLC Flash Capacity

*By Brian Garrett*

September 2011

## Contents

Introduction .....	3
Background.....	3
Source: Enterprise Strategy Group, 2010.....	3
Kaminario K2 Hybrid Solid-state SAN .....	4
ESG Lab Validation .....	5
Fast .....	5
Safe .....	9
Easy.....	10
ESG Lab Validation Highlights .....	12
Issues to Consider .....	12
The Bigger Truth .....	13
Appendix.....	14

### ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. 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. This ESG Lab report was sponsored by Kaminario.

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 from time to time. 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.

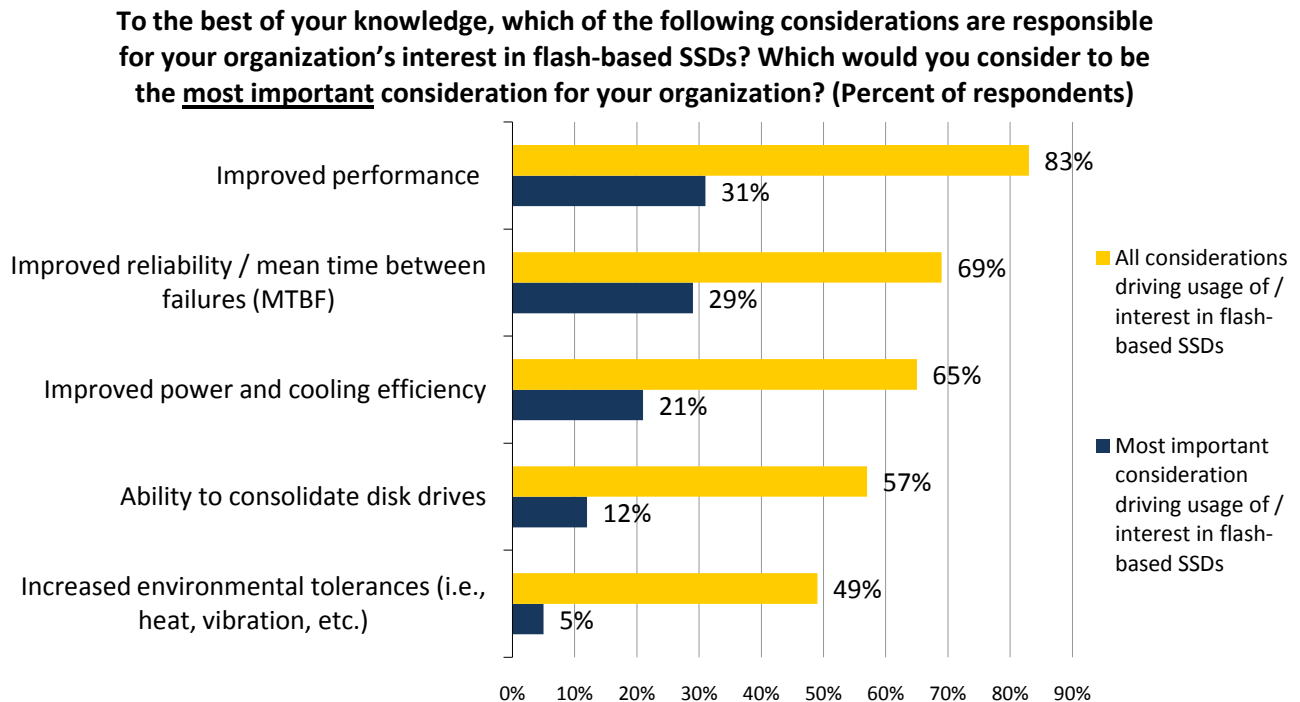
## Introduction

This ESG Lab Validation report documents the results of hands-on testing of the K2 Hybrid Solid-state storage area network (SAN) from [Kaminario](#) with a focus on the value of a highly available storage solution that leverages a mix of dynamic random access memory (DRAM) and multi-level cell (MLC) flash technology to simply solve IO performance problems in business critical database environments.

## Background

There has been a significant increase in the number of organizations considering flash-based solid-state technology for data center storage in recent years. When ESG asked IT managers their reasons for assessing this option, improved performance topped the list, with 83% of respondents citing it as a consideration and 31% citing it as the most important consideration (see Figure 1). A significant majority of those surveyed also mentioned improved reliability (69%), improved power and cooling efficiency (65%), and the ability to consolidate disk drives (57%). While no single consideration outweighed the others as the most important reason, performance and reliability were mentioned as most important by 60% of respondents.<sup>1</sup>

Figure 1. Reasons for Interest in Flash-based Storage Solutions



Source: Enterprise Strategy Group, 2010.

While solid state drives (SSDs) have been used in consumer products such as MP3 players for a decade or more, data center deployments are new in the last couple of years. Solid state solutions have provided organizations with orders of magnitude improvements in IO performance compared to traditional hard disk drives (HDD). A single SSD can deliver performance that might take hundreds of HDDs as it eliminates storage bottlenecks for applications that require fast performance and low latency. Today's SSDs are also highly reliable with a life expectancy around ten years, and are much more power efficient than mechanical, spinning HDDs. These features make solid state solutions attractive from an operational perspective, but the high acquisition cost of solid state capacity compared to traditional hard drive capacity has focused early adoption within businesses whose revenue depends strictly on application performance including transactional databases, and trading and other financial applications.

<sup>1</sup> Source: ESG Research Brief, [Enterprise SSD Usage is Not a Flash in the Pan](#), June 2009.


## Kaminario K2 Hybrid Solid-state SAN

Kaminario K2 Hybrid Solid-state SAN storage system is a high-performance solid state storage solution that uses a mix of high-speed DRAM and cost-effective MLC flash to eliminate IO bottlenecks and accelerate the performance of business critical applications. The K2 Hybrid brings high-performance, low-latency storage to a wide range of applications, workloads, and budgets while maintaining high availability, ease of use, and lower cost per IO than a traditional disk-based storage solution.

DRAM is extremely fast with low latency, but as volatile memory it requires the use of batteries or traditional hard drives to retain data when power is off. MLC Flash technology is faster than hard disk drives, but not as fast as DRAM; however, it is less expensive than DRAM and is non-volatile, retaining data without battery backup. The K2 Hybrid leverages the advantages of both with DRAM to provide the ultra-fast performance required by write-intensive, latency-sensitive application workloads and MLC flash to offer a more cost effective performance boost for read-intensive workloads. This hybrid approach is ideally suited for a mix of application types, as well as different file types within an application (such as database temp files, redo logs, etc.).

As shown in Figure 2, the fully redundant Kaminario Scale-out Performance Storage Architecture (SPEAR) within each K2 system was designed for enterprise-class levels of performance, scalability, and high availability. Two or more IO directors are used to create a fault tolerant Fibre Channel connection with servers. Host addressable volumes are striped over three or more hybrid data nodes with each data node containing a mix of DRAM and MLC flash capacity. Data residing in volatile DRAM memory is stored on mirrored disk or solid state MLC flash disk in the unlikely event of a power failure in the data center. A redundant internal 10 Gigabit Ethernet network is used for internal communication, automated failover, and load balancing.

Figure 2. Kaminario K2 Hybrid Solid-state SAN




K2 solutions are delivered in a standard 42" rack with enclosures for up to 16 blade servers, dual power supplies, and three types of switches: Fibre Channel connecting K2 ports to external hosts, 10GigE interconnects for sending data between blades, and 1 GigE interconnects for management communications.

A K2 system can be deployed in a matter of hours. Appearing to operating systems, applications, and end-users as a traditional block-based storage device, K2 Hybrid capacity can be deployed with no changes needed to existing applications, servers, or storage environments.

## ESG Lab Validation

ESG Lab tests were performed over the course of three days in a Kaminario lab with a goal of examining the performance, high availability, and ease of deployment of the K2 Hybrid. The test bed used during ESG Lab testing is shown in Figure 3. A single server with eight Intel Xeon processor cores and 4 GB of RAM was connected via Fibre Channel to a K2 Hybrid configured with two IO directors and six hybrid data nodes. Each of the six data nodes was configured with a mix of DRAM and flash capacity.<sup>2</sup> A high-performance N+1 topology was automatically created as host volumes were striped over five of the six data nodes.

Figure 3. ESG Lab Test Bed



### Fast

ESG Lab testing began with a goal of assessing the speed, low latency, and scalability of a K2 Hybrid solution. A multi-user stock brokerage workload was tested with a Microsoft SQL Server application. The industry standard Iometer utility was used to measure performance scalability as additional IO directors and data nodes were used to increase the capacity and performance of a K2 Hybrid solution.

The database application used during ESG Lab testing was designed to emulate the activity of users in a typical online brokerage firm as they generated trades, performed account inquiries, and conducted market research. The workload was composed of ten transaction types with a defined ratio of execution. Four of the transactions performed database updates and the rest were read-only. The workload was deployed on a mix of K2-H devices. The transaction log and tempdb files were stored on a DRAM volume while the database files were stored on a flash volume. The results are summarized in Table 1.

<sup>2</sup> For more configuration details, please refer to the Appendix.

Table 1. Brokerage Database Application Performance Analysis

Customers	User Load	DB IO Response Time (ms)	Log IO Response Time (ms)	Tempdb IO Response Time (ms)	Application Response Time (ms)	CPU Utilization
20,000	1	0.8	0.4	0.8	11	32%
20,000	5	0.8	0.4	0.8	12	60%
20,000	10	0.8	0.4	0.8	17	81%
20,000	15	0.8	0.4	0.8	26	97%
20,000	20	0.8	0.4	0.8	34	100%

### What the Numbers Mean

- The Microsoft SQL Server R2 database was initialized to support 20,000 simulated brokerage customers.
- The number of customers using the system at the same time was increased from 1 to 20.
- A very low host IO response time of less than one millisecond was recorded by the Microsoft perfmom utility for the database files and tempdb volumes. The transaction log volume (stored on DRAM) reported a 0.4 ms response time.
- All response time values include K2 internal latency as well as external (host fiber to storage) latency.
- Low IO response times delivered excellent application response times during low levels of database activity.
- Fast K2 Hybrid IO response times avoided the common problem of running into a storage performance bottleneck which limits the number of users that a multi-user interactive database application can support.
- As user load increased, application response times increased and CPU utilization became the bottleneck.
- As user load increased, K2-H response times were not affected.

Having exhausted the CPU of the dual quad core server that was used during Microsoft SQL Server testing, ESG Lab switched to the industry standard open source Iometer utility<sup>3</sup> with a goal of measuring the performance of the K2-H with a larger number of simulated database users. The Iometer utility was chosen due to its low CPU overhead and its ability to emulate real-world workloads.

The ESG Lab 8 KB OLTP Iometer workload was used during this phase of simulated database testing.<sup>4</sup> The performance of two configurations was compared:

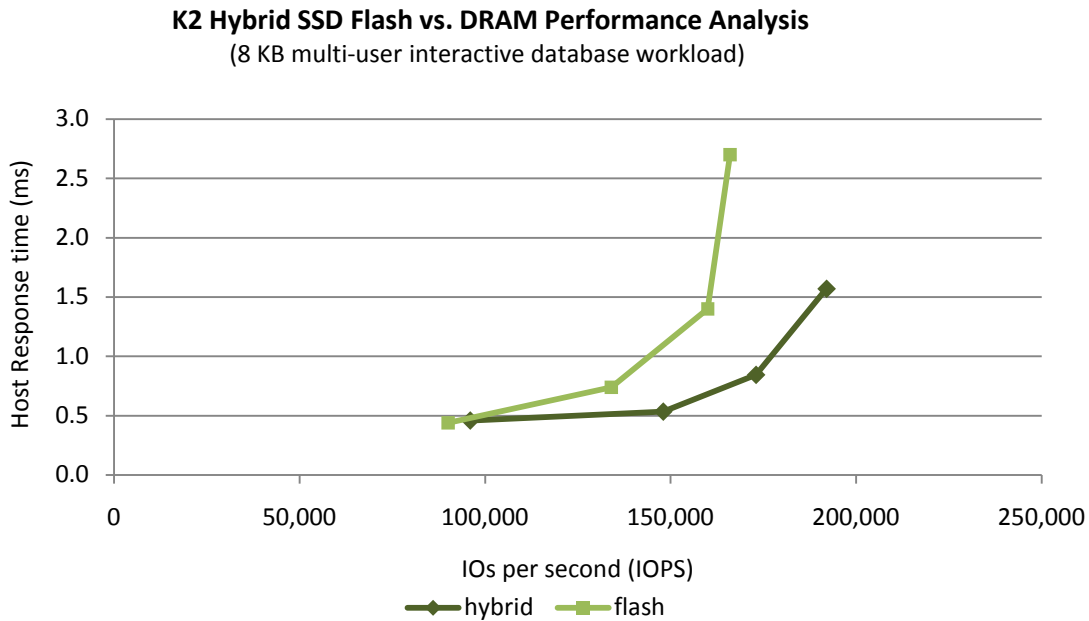
1. A hybrid mix of six 10 GB DRAM LUNs and four 600 GB flash LUNs
2. A flash configuration with ten 600 GB flash LUNs

The number of outstanding IOs per volume was increased with a goal of measuring the performance scalability of a six-node K2 Hybrid solution. The results are depicted in Figure 4.

<sup>3</sup> <http://sourceforge.net/projects/iometer/>

<sup>4</sup> 8 KB IO requests, 70% random reads, 30% random writes. The ESG Lab Iometer workloads are freely available at <http://www.enterprisestrategygroup.com/using-esg-lab-workloads/>

Figure 4. OLTP Database Performance Characterization



#### How to Read the Graph

- Host response time, a measure of how long each IO operation takes, is represented on the Y-axis as the amount of traffic, measured in IOs per second, is increased along the X-axis.
- This graph, known in the industry as a response time throughput curve, provides an excellent visual picture of how well a storage subsystem performs under duress.
- As traffic increases along the X-axis, response times tend to increase along the Y-axis.
- When a system becomes overloaded, it has a hard time keeping up as queues start to fill up. Response times elongate and the system feels slow to end-users. This point of no return, known as the “knee” of the curve, represents the breaking point after which more traffic cannot be handled gracefully.
- The best performing configuration is therefore represented by the lowest curve extending the furthest to the right.
- A storage solution which delivers more sustained IOs per second can be used to support a larger number of database users without running into performance problems.


#### What the Numbers Mean

- The Hybrid and flash configurations both recorded excellent response times of less than 0.5 milliseconds when processing a high rate of 90,000 IOs per second.
- A host IO response time of 0.5 milliseconds is four to ten times faster than a typical hard drive.
- The hybrid configuration scaled to support up to 192,000 IOPS with a low host response time of 1.57 milliseconds.
- It would take more than 1,000 power-hungry disk drives to deliver 192,000 IOPS.
- The K2 Hybrid configuration with a mix of DRAM and flash volumes scaled to support more IOs per second with better response times than the configuration composed entirely of flash volumes.

Next, ESG Lab audited a series of tests performed by the Kaminario engineering organization with a goal of confirming that a K2-H system can easily support more IOPS without an increase in response time.

The number of IO directors and data nodes was increased as the Iometer utility was used to simulate an interactive database workload.<sup>5</sup> Like the tests that were performed during ESG Lab testing, the performance of a Hybrid configuration with six DRAM LUNs and four flash LUNs was compared to a configuration with ten LUNs. The results are shown in Figure 5.

*Figure 5. K2 Hybrid Performance Scalability*



**What the Numbers Mean**

- Host volumes which were automatically striped over all available data nodes using an N+1 algorithm delivered near linear performance scalability as IO directors and data nodes were added.
- The audited hybrid and flash performance correlates well with the results recorded during ESG Lab testing with five and six data nodes.
- A 16 data node K2 Hybrid delivered an extremely high level of 825,344 IOPS for a database workload.

**Why This Matters**

For many transactional databases and business-critical applications, faster application performance means more revenue, better customer satisfaction, and greater productivity. The challenge is to accomplish performance objectives without adding extraordinary costs or resorting to wasteful strategies such as over-provisioning.

ESG Lab validated the ability of the cost-effective K2 Hybrid array to break through storage bottlenecks and deliver extremely fast performance for multi-user database workloads. With response times of under a millisecond and the ability to scale to more than 800,000 IOs per second, ESG Lab has confirmed that a K2 Hybrid solution that fits in a single rack is four to ten times faster than a traditional disk array consuming multiple racks full of power-hungry disk drives.

<sup>5</sup> 4 KB random, 75% read/25% write.




## Safe

The highly available enterprise-class architecture of the K2 Hybrid uses redundant components and distributed load balancing algorithms to ensure continuous operation in the unlikely event of hardware failure.

ESG Lab injected errors while a heavy simulated database workload and directory level file copy operation were run on a K2 Hybrid solution with two IO directors and six data nodes. A video stored on a K2 Hybrid volume was also viewed during error injection testing. As shown in Figure 6, a FC host interface and an internal 10 GigE network connection were disconnected. Error injection testing ended with a hard shut down of a data node.

*Figure 6. Error Injection Testing*



The multipath IO (MPIO) driver that's built into the Microsoft Windows Server operating system provided transparent failover when the FC cable was disconnected. K2 Hybrid load balancing and failover algorithms ensured that IO processing continued without interruption when a data node was shut down.

Application workloads continued without interruption during ESG Lab error injection testing. Error messages and system status were monitored via the K2 Hybrid management interface. The ability to configure e-mail notification and SNMP traps for automated alerting was noted.

### ***Why This Matters***

IT organizations expect their data to be safe and secure; today's HDD arrays provide reliability and high availability as standard features. SSD safety concerns include DRAM memory volatility and the tendency of flash to wear out over time; as a result, IT managers charged with keeping data secure and recoverable are often wary of implementing SSD.

The redundant architecture of the K2 Hybrid with N+1 load balancing and failover of SSD storage media provide enterprise-class levels of fault tolerance. The K2 Hybrid provided non-stop data access as a variety of hardware errors were introduced during ESG Lab testing.

## Easy

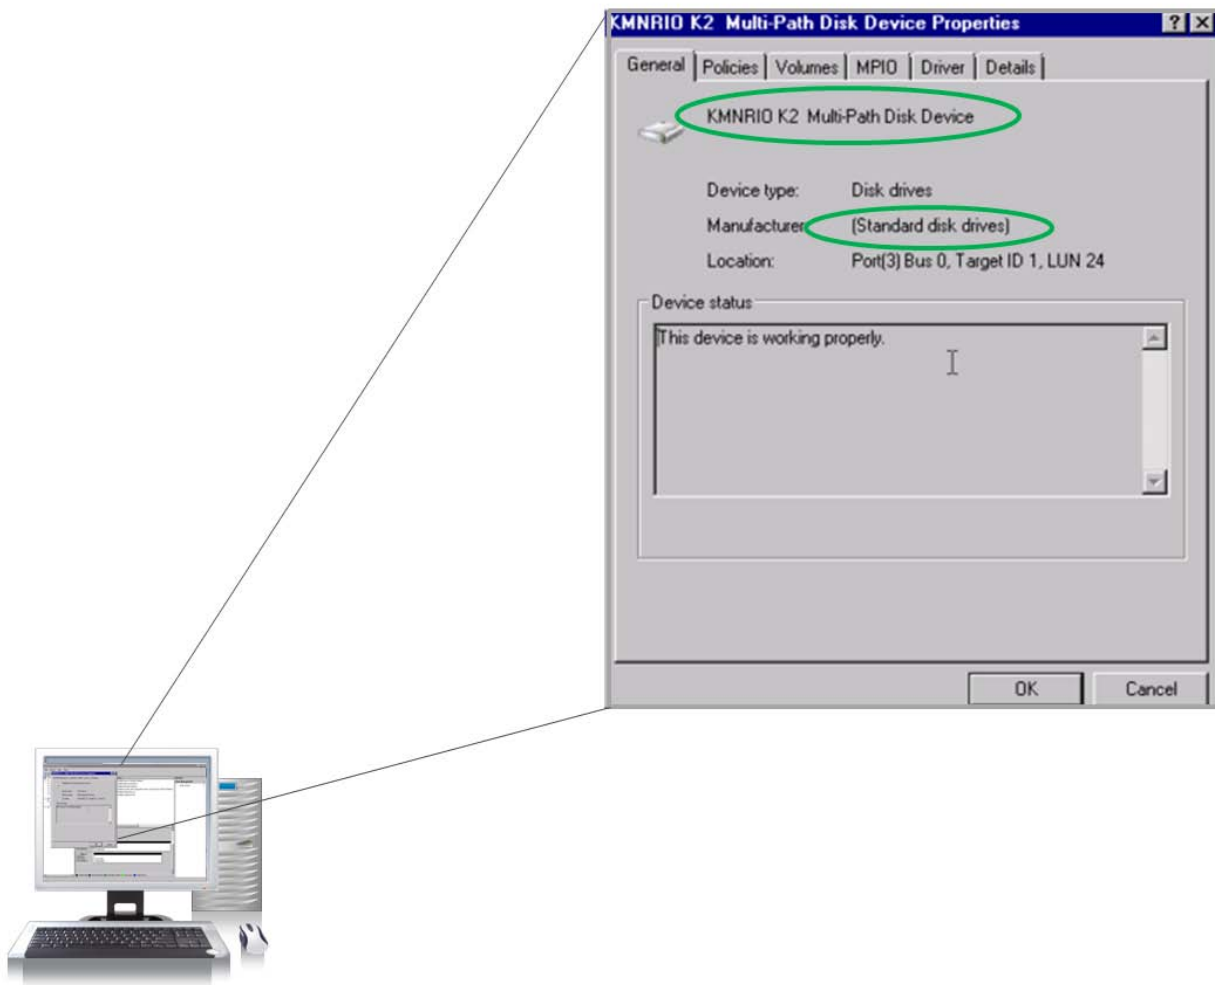
The K2 Hybrid works like a traditional disk array, presenting disk capacity as a standard block-based device. No changes are needed to existing applications, operating systems, and storage environments.

ESG Lab tested the simplicity of configuring a new K2 Hybrid volume for use in an online database application. Four simple commands executed from the Kaminario command line interface were used to:

1. Create a new volume
2. Create a new host definition using the World Wide Name(WWN) of the database server
3. Map the new volume to the database server

The Windows disk administrator utility was used to rescan and discover the new volume as shown in Figure 7. The volume was initialized and an NTFS file system was created from the disk administrator console.

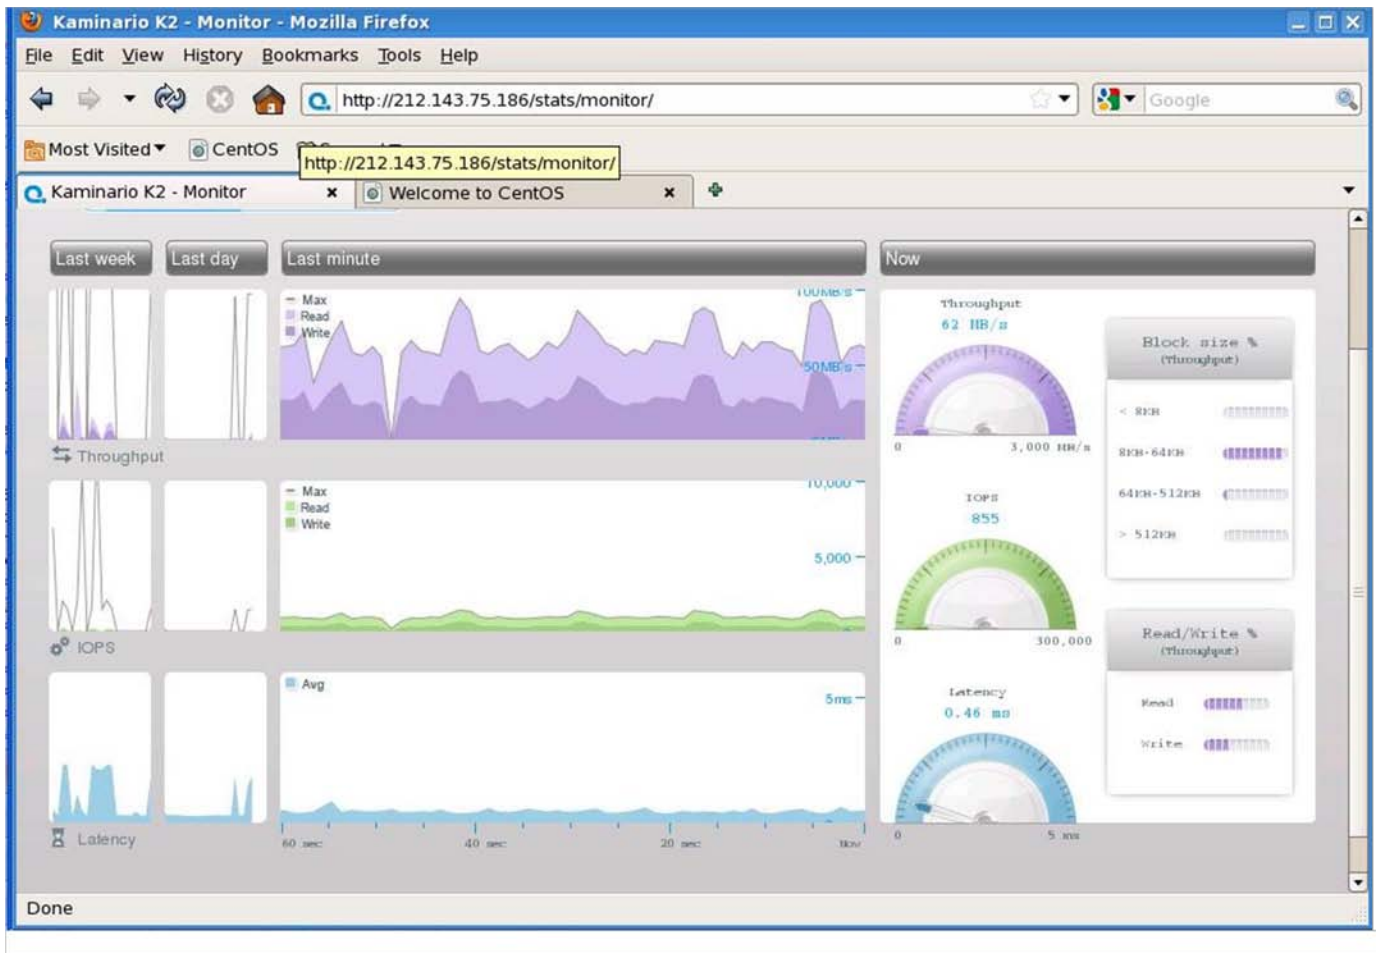
*Figure 7. Configuring a New LUN*



Three minutes after getting started, files were being copied to a new K2 Hybrid volume which appeared to users and applications as a “standard disk drive” accessed a Windows drive letter: the K: drive.

The Kaminario K2 Monitor GUI was used to monitor performance during ESG Lab testing. As shown in Figure 8, the graphical interface shows current and historical performance in an intuitive format. The dials toward the right summarize key metrics that are needed to understand the current performance of a storage solution (throughput, IOPs, latency). Trended performance for the last minute, last day, and last week are shown toward the left.

Figure 8. Monitoring K2 Hybrid Performance Statistics



Graphical presentation of storage performance is supported by some, but not all, traditional disk array vendors. Often, this data must be collected from a command line interface. Graphically displaying storage performance trending over time often requires the use of an additional licensed software package. As a growing number of organizations deploy solid state solutions to eliminate IO bottlenecks and improve application performance, storage performance monitoring tools that are easy to deploy and interpret are a valuable addition to an administrator's toolbox. Of the dozens of storage solutions that ESG Lab has tested over the past eight years, the Kaminario K2 monitor had the best look and feel, and trending capabilities of a built-in performance monitor.

### ***Why This Matters***

Deploying, expanding, and tuning a storage infrastructure can be disruptive and costly, requiring application changes, performance tuning, manual mapping of LUNs to hosts, modifications of data protection processes, etc.

ESG Lab testing validated the fast and easy configuration and deployment of the K2 Hybrid with three commands in three minutes, without tuning or changes to existing applications or processes. Most customers will purchase the K2 primarily to accelerate performance; Kaminario makes it easy for them to monitor with a graphical dial displaying both instantaneous and historical performance, providing knowledge of performance trends as well as current conditions.

## ESG Lab Validation Highlights

- ☑ Testing with industry standard multi-user database workloads confirmed that K2 Hybrid flash devices can be used to break through storage performance bottlenecks. Testing with volumes striped over flash devices in six storage nodes yielded up to 190,000 IOs per second with extremely fast response times of under a millisecond—that's four to ten times faster than a power-hungry disk array with more than 1,000 spinning hard drives.
- ☑ Testing with a variety of performance sensitive workloads yielded high levels of performance and throughput that scaled predictably as storage nodes were added. Testing with a 4 KB OLTP database workload peaked at an extremely high rate of 825,344 IOPS for a K2 Hybrid solution with 16 data nodes.
- ☑ ESG Lab confirmed that the K2 Hybrid architecture is fully redundant, highly available, and safe. Applications ran without interruption as errors were injected in host ports, internal networking ports, and storage nodes.
- ☑ Allocating K2 capacity was easy. Three minutes and three commands after getting started on a pre-wired and configured system, a Windows server had recognized freshly configured K2 Hybrid capacity as a standard block device that works with existing applications without tuning.
- ☑ The intuitive user interface of the K2 performance monitoring user interface made it easy to monitor current and historical storage performance.

## Issues to Consider

- ☑ The K2 Hybrid solution supports a mix of DRAM and flash capacity to meet a wide variety of price and performance requirements, but it doesn't currently support automated online migration between DRAM and flash. Policy-based migration for a whole volume or chunks within a volume (a.k.a., sub-LUN tiering) would be a valuable enhancement to the K2 Hybrid product line. Automatically tiering up to automatically alleviate performance problems and tiering down to optimize costs would enhance the price/performance advantages of a K2 Hybrid solution.
- ☑ While the command line interface used to configure and deploy K2 Hybrid capacity was intuitive and straightforward, a wizard-driven graphical user interface would make the solution even easier to deploy with existing applications and storage environments.
- ☑ While the redundant architecture of the K2 Hybrid delivers enterprise-class levels of high availability, it does not yet support extended functions that are supported by legacy enterprise-class disk arrays (e.g., policy-based quality of service (QoS)). Kaminario has advised ESG that the development of enterprise-class extended functions is underway and planned for an upcoming release.

## The Bigger Truth

While solid-state storage solutions have grown in popularity for data center implementations, many organizations have resisted the call. For performance- and latency-sensitive workloads, solid state storage solutions can provide orders of magnitude improvements in effective speed, but it has been difficult to justify the cost. For some, the answer has been to use SSDs on a limited basis as a cache or as a storage tier, but the difficulty of predicting when and where to move data has often restricted the benefit. Other solid state products are designed for particular workloads and tend to be smaller, single-controller solutions that are not fault tolerant. An affordable, enterprise-class, highly available, all-SSD storage solution has been missing.

Interest in solid state storage solutions has been gaining. While only 16% of respondents to a recent ESG research survey had already deployed systems with flash-based SSDs, an additional 60% were either evaluating systems or had interest in them.<sup>6</sup> Also, speed of performance is not the only reason; SSDs are much more power efficient and can significantly reduce data center costs. This energy efficiency is an important driver: in 2010, 22% of ESG survey respondents indicated that purchasing more power efficient storage hardware would be an area of significant IT investment over the next 12 to 18 months.<sup>7</sup>

IT managers looking to speed application performance may be wary of solid state technology and want answers to two key questions:

- Can they be assured that data will be safe and recoverable?
- Will they have to manually tune their environments to incorporate solid state technology?

Kaminario has a simple answer: the K2 Hybrid. It is fast, safe, and easy.

Combining the cost effective speed of flash and the blazingly fast performance of DRAM devices in a SAN-attached solution with enterprise-class high availability, ESG Lab has confirmed that the K2 Hybrid is ideally suited to meet the needs of performance-hungry database applications. Testing with hefty database workloads yielded response times that were four to ten times faster than a traditional disk array with more than a thousand disk drives. Performance and throughput scaled predictably as storage nodes were added. The redundant, highly available architecture kept data safe and operations continuing while components were both removed and re-integrated. While most management remains CLI-based, installation was easy and required no tuning or modifications to the environment. The K2 Hybrid appeared to existing applications just like a standard block device.

The K2 Hybrid enables organizations to take advantage of the speed of DRAM and the affordability of MLC flash as it extends the benefits of solid state usage to multiple workload types simultaneously. The K2 Hybrid eliminates IO bottlenecks and accelerates business applications as it makes an organization more productive, efficient, and “green” without breaking the bank.

---

<sup>6</sup> Source: ESG Research Brief, [Enterprise SSD Usage is Not a Flash in the Pan](#), June 2009.

<sup>7</sup> Source: ESG Research Report, [2010 IT Spending Intentions Survey](#), January 2010.

## Appendix

Table 2. ESG Lab Test Bed

Test Bed Configuration Details	
Storage	Kaminario K2-H 2 IO directors 6 data nodes (one serving as a spare node for high availability) 6.1 TB usable (ten 600 GB flash LUNs, ten 10 GB DRAM LUNs) K2 v2.0 software
Server	Intel Xeon 5420 @ 2.5 GHz, 4 GB RAM Microsoft Windows Server 2008 Microsoft SQL Server 2008 R2
Storage Area Network	8 Gbps Fibre Channel



Enterprise Strategy Group | **Getting to the bigger truth.**