

Kaminario K2 Balances an IT Seesaw — Matching I/O Rate to Processor Characteristics

Analyst: Jim Baker

Management Summary

One of the barriers to putting primary data on cloud infrastructures can be characterized as an IT Seesaw. You remember the seesaw (or center-hinged balance beam) from your childhood. If you were on one end and a Big Kid (likely from the next higher grade) sat on the other end, you could be left dangling helplessly in the air until he gave up and allowed you to return safely to the ground. This particular IT seesaw has processors and memory on one end (you) and lots and lots of storage on the other end (sort of like the heavy fourth grader). Until and unless there is a solution to getting needed data from storage to system processors and back again really quickly, you risk making expensive processors wait for their data or worse timing out your I/O or even timing out the entire application. And, if the processors are waiting, you can bet that their users are waiting too.

Kaminario, Inc. of Newton, Massachusetts, has a solution to this problem of a mismatch between the speed at which data arrives versus the speed at which the processor needing that data can handle it as part of the compute cycle. It is a DRAM-based appliance that acts like a standard block device sitting behind a server as very fast primary storage for all of the organization's data. Alternatively, it could be configured as a peer to traditional storage handling only the most critical data (such as DBMS indexes) while the traditional storage handles other less important data. The product name is *Kaminario K2*.

The choice of K2 as a product name is insightful because of what K2 represents. K2 is the name of the second tallest peak on Earth (behind Mt. Everest). As an ultra tall peak, it likely penetrates cloud formations and may even create its own weather system (because it is so high). With public and private clouds forming all over the Internet, perhaps it just might be the time to step back and take a fresh look at a cloud-penetrating technology like Kaminario's K2.

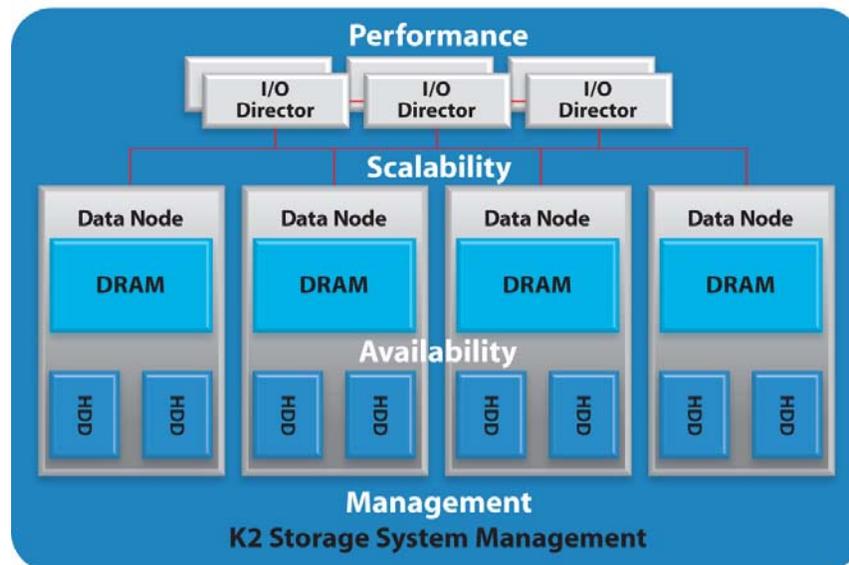
Kaminario's K2 does not just penetrate public and private clouds. It penetrates tricky performance bottlenecks anywhere in the I/O stack. Performance issues are many and varied regardless of the application under review. Fixing them is like working through the layers of an onion. Some layers are thin and easy to solve, whereas others are thick and harder to fix. When you knock down one issue, you do not know whether the next one will be trivial or complex. What you do know is that – inevitably – there is yet another issue to be addressed. That is why it is important to take advantage of the expertise of Kaminario as they examine the I/O path and seek to speed it up.

The Kaminario K2's claim to fame is that it speeds the I/O cycle because it is in the physical (or virtual) data path and provides necessary information to hungry, very fast processors. It does so without the traditional disadvantages of Dynamic Random Access Memory and/or Flash Memory and/or non-cached rotating storage. All right, Class, by show of hands, please indicate which of you are NOT interested in better performance at a reasonable price? Hmmm, I see no hands. Class, please read on to learn more about the K2 DRAM Appliance from Kaminario.

IN THIS ISSUE

➤ Kaminario's DRAM Solution	2
➤ Questions for Kaminario – and Their Answers	5
➤ First Steps	5
➤ Conclusion	5

Exhibit 1 — Kaminario's Scale-out Performance Architecture (SPEAR)



Source: Kaminario

Kaminario's DRAM Solution

Long the staple of the server, Direct Random Access Memory has been promoted to yet another use in the *K2 Appliance* from Kaminario.

Describing the Kaminario K2 Appliance

The *Kaminario K2 Storage System* is a high-performance, block storage array with the following features:

- SCSI over Fibre Channel interface (SCSI-3)
- No single point of failure
- Hard Disk Drives behind DRAM to provide redundancy and backup to the data that resides in the DRAM
- Fault tolerant with failure analysis and automatic failover
- Multiple host access
- Automated Management software, including command line interface (CLI) and graphical user interface (GUI), and automated data distribution over all the data nodes
- Remote configuration management
- Notifications of events through email

The Kaminario K2 runs a grid of industry-standard hardware components: blade servers, network cards, and switches. In general, Kaminario's K2 is a block device designed for high-performance purposes, utilizing DRAM for the entire data set, which serves all read and write

requests and is mirrored by a set of low-performance HDDs (Hard Disk Drives) for *persistence*. (More about persistence to follow.)

Scale-Out Architecture

One of the unique features of the Kaminario K2 is its Scale-out Architecture. When you configure a Kaminario K2, you get to size the appliance to exactly what you need. If you need more IOPS (performance), then add more I/O Directors. If you need more capacity, add more Data Nodes. It is unique in high performance appliances to be able to scale performance and capacity independently.

To the host server, the Kaminario K2 looks like an independent storage array managed in its entirety by SPEAR, its storage operating environment. (See Exhibit 1, above.)

Overcoming Classical DRAM Objections

Why DRAM? There are many classical objections to a Direct Random Access Memory device, especially for enterprise-class storage solutions. To handle these objections, Kaminario has engaged in a product-justifying game of *Technical Objection Take Away*. Here are the game rules and a quick play-by-play analysis:

1. Starting in order, list the objections to your solution beginning with the most worrisome or the one that is counter to prevailing wisdom.
2. Take away this objection by solving this technical objection with creative engineering

that is both persuasive and compelling.

3. Move to the next most worrisome objection and solve this one, also.
4. Continue until all objections are answered (taken away) by the solution.
5. Apply for patents to keep your solution's intellectual property within the control of (in this case) Kaminario.
6. Price the product well and offer the solution in a variety of environments to which multiple customers can relate.

Classical Objection #1 – DRAMs have no persistence

“**Persistence** in computer science refers to the characteristic of state that outlives the process that created it. Without this capability, state would only exist in RAM, and would be lost when this RAM loses power, such as a computer shutdown.”¹

For storage, it works like this. In the event of a power failure, all data that was in the DRAM is lost unless it is transferred to a persistent device such as a Hard Disk Drive (HDD). It is a known requirement for enterprise applications that they must exhibit persistence because data cannot be allowed to become lost – or even delayed. Kaminario's solution is to keep “always on” electrical power that is provided to the DRAMs. As an additional precaution, Kaminario's K2 continually writes data to rotating disk that will store it in a recoverable state even when/if power is lost. Kaminario then adds redundant power supplies to handle the unlikely event of a power supply failure. An uninterruptible power supply (UPS) is also standard equipment on the K2.

Classical Objection #2 – DRAMs are expensive so appliances based on them will be expensive too.

At one time, DRAM memory indeed was expensive, but not any more. They now are commodity parts sitting on the (declining) commodity parts price curve. Virtually every electronic device today uses DRAM from high-definition TVs to cell phones to digital game consoles to book readers to computers of all sizes to hand held devices of all kinds. With so many uses, their piece part prices have dropped accordingly.

Kaminario has supplemented the DRAMs themselves with industry standard commodity parts, including blade servers, network cards, and

Exhibit 2 — Examples of Kaminario K2 at Work

Telco Billing queries

- ☛ 81X Improvement (102 minutes query down to 1.2 minutes)

Credit card Marketing analysis (SAS)

- ☛ 40X Time Reduction (181 minutes for report down to 4.5 minutes)

Fortune 10 Products (Cognos BI & product simulation)

- ☛ 20X Time Reduction (21 minutes to process down to 1 minute)

Government Real-time data mining

- ☛ 25X Time Reduction (15 minutes free files text search down to 0.6 minute)

Source: Kaminario

switches. Again, parts costs are minimized.

Classical Objection #3 – DRAMs are not reliable.

Maybe they were not reliable way back in the late 1980s, but such is not the case now. When was the last time you heard of a DRAM failure or the need to “fence off” bad sections of memory or cache? Not for a long, long time. In fact, EMC used to require a duplex version of their storage cache in the data path preceding the round, brown spinning disks where data ultimately resided; but no longer. The reliability of today's DRAMs plus embedded reliability check values and mechanisms makes them extraordinarily reliable. Perhaps the better question ought to be *why not DRAMs?*

The Case FOR DRAMs

Storage systems from the early 1990s have not fundamentally changed since their initial breakthroughs. Storage Level cache has been a feature of enterprise-class storage for decades now. As such, it is susceptible to innovation breakthroughs like the Kaminario K2. Why?

- DRAMs are the fastest form of storage. (See Exhibit 2 above to see how really fast DRAMs can be in various work environments served by the Kaminario K2 appliance.)
- Performance profile of DRAMs is predictable
- Performance profile of DRAMs is achievable
- DRAMs offer consistent results no matter what the type of workload

¹From Wikipedia, accessed November 2, 2010; see [http://en.wikipedia.org/wiki/Persistence_\(computer_science\)](http://en.wikipedia.org/wiki/Persistence_(computer_science)).

- DRAMs are dropping in price

Let's look at various technologies that could compete with Kaminario K2.

Solid State Devices (SSDs)

Solid State Devices (SSDs; also known as Enterprise Flash Drives or EFDs) possess wear-leveling issues. The problem with SSDs – in short – is that an SSD has a fixed duty cycle that limits its life. If the same memory in an SSD is written to repeatedly, then it could be prematurely worn out. When used in a DBMS application (such as the index portion of the file), there are continual data changes, meaning that SSDs make poor caching targets compared to DRAMs. To address this “wear leveling” issue, those providers selling SSD solutions have taken two steps. First, they rotate (or “round robin”) writes to make sure certain devices do not get “hit” more than others do. Second, they implement recommendations from the SSD manufacturers not to have all-SSD solutions. Instead, they suggest *Rules of Thumb* calling for a low percentage of SSDs vs. HDDs in most systems (10 to 20%). These are typical conservative steps taken while a technology is still early in its life cycle, especially as compared to DRAMs. For these reasons, SSDs are not appropriate to the massive caching that is done in the Kaminario appliance.

Application Specific Appliances

Several major companies have taken a different approach than Kaminario by developing or acquiring special-purpose appliances. IBM is in the process of buying Netezza² as a DBMS-specific appliance maker. EMC offers the Greenplum appliance³ that is also specialized. While these special-purpose appliances may address application specific problems, such as analytics or table manipulation, they do not necessarily help solve I/O acceleration issues in the way that Kaminario's K2 does. Kaminario's approach is general purpose and is therefore more universally applicable.

The Kaminario K2 is definitely NOT special purpose and that is its biggest attraction. It is operating system, database, and application agnostic. This is a huge advantage for the following reasons.

- **Wide Applicability** – The Kaminario K2 can be used in any situation suffering from an I/O performance problem. There are no special

versions that vary by the context into which the product is placed.

- **Investment Protection** – Suppose your performance problem is corrected. Maybe your DBMS vendor has a new feature or a patch that addresses your particular performance issue. The Kaminario K2 can be taken out of service for this issue and placed into service on your next I/O performance issue. In terms of investment protection, the Kaminario K2 is reusable totally in another scenario.
- **Extensibility** – The Kaminario K2 can scale out with the addition of more DRAM or HDDs in the appliance. If the problem is related to the sheer size of the data being used, this extensibility is a valuable advantage.
- **Flexibility** – Not only can the Kaminario K2 be employed to find the root cause of an I/O stack problem, it can be used as a diagnostic tool to *rule out* such a problem. When you know that the problem is not in the I/O chain, then you can investigate other causes such as problems with scripts, procedures, loops on the processor, even operator-induced wait times.
- **Support infrastructure and training** – As a general-purpose appliance, operator training and support infrastructures (including parts) can be leveraged across virtually every server to which it is attached.

General Purpose Storage

General-purpose (GP) storage and Kaminario K2 are mutually complementary. Whereas GP storage has to handle virtually every need from archiving to backup and recovery to general-purpose replication to simple I/O support regardless of the application, Kaminario's K2 is designed to improve performance-bleeding applications. Universally, they will be placed into service where GP storage already exists. When the Kaminario K2 improves I/O chain performance, there is a cumulative effect that frees storage cycles to perform other functions, including adding speed to all the other applications in the system.

Endless Application Tuning

This process is futile and frustrating. It is likely how you have been coping for years with only marginal improvements. The Kaminario K2 takes a different approach. By adding this appliance to your infrastructure portfolio, you may be able to fix the performance issues. There probably is going to be a corollary benefit of freeing your staff to work on new or different problems because the headache has been removed. They can

² See <http://www-03.ibm.com/press/us/en/pressrelease/32514.wss>.

³ See <http://www.emc.com/about/news/press/2010/20100706-01.htm>.

concentrate on more strategic issues because of their newly found free time.

Questions for Kaminario – and Their Answers

The Clipper Group asked several questions about the K2 during the preparation of this report. Below are the Kaminario answers.

Why no dedupe capability?

Don't need it and it would negatively affect performance. Kaminario's K2 is in the data path and dedupe overhead would be a performance killer. If you want dedupe (for long term archiving, for example), there are devices in the marketplace addressing that challenge.

Why no thin provisioning capability?

Thin provisioning is a capability for a future release but it is not now part of the Kaminario K2 appliance. In the meanwhile, Kaminario's K2 delivers a ready-made solution to performance problems that arise when the I/O infrastructure does not match the processing power of today's and tomorrow's ever-faster processors.

Why no support for mainframes?

As a small start up company, Kaminario has chosen to focus on open systems server and storage environments only. These environments seriously lack performance tools that will help deliver faster results for their users. Mainframes already have a rich legacy of performance-enhancing products and features. The Kaminario K2 is a major step forward in the Open World, especially as virtualization evolves and clouds form on their horizon.

First Steps

To see for yourself how high your application performance could go, Kaminario offers a web-based analysis tool with expert performance engineers ready to interpret your specific data. They know what to look for because of their experiences with other clients. Additionally, their knowledge base builds daily as more and more users take advantage of the Kaminario K2. This offer from Kaminario may be found on the product pages of their website at www.kaminario.com. There is no charge or obligation for those wishing to have Kaminario look at their "trouble" applications or databases. Kaminario can help you identify I/O performance bottlenecks, quantify your performance improvement potential (wouldn't it be great to know exactly how much headroom you really do have?), and find fast and easy-

to-deploy solutions. This valuable service will enable you to quickly accelerate your application and database response times and deliver extreme performance to your users.

Next Steps

Ask any IT professional and he or she will be able to identify their *Problem Application*. It's the one that never has adequate performance, i.e., the one that wakes them up during the night. This should be your first application to receive the Kaminario K2 boost. Kaminario's experts claim that each K2 appliance can be configured, tested, and placed into service in less than a half day, certainly a reasonable period for silencing your biggest headaches.

Integrating into Established Environments

Kaminario's K2 easily slips into any application environment needing its performance-enhancing services. It is operating system and database management system agnostic. No special versions or special agents are required. Simply insert it into your established SAN block mode environment. If you have made a prior storage buying decision to purchase EMC, no problem. IBM, no problem. NetApp, no problem. Any other block-mode storage array brand, no problem.

Conclusion

With regard to performance, there IS another way. It is called Kaminario K2 and it is affordable. Because the device is shared across many servers, its nominal \$100,000 price per appliance can be recovered quickly. Exact pricing varies by configuration, so you are encouraged to contact Kaminario directly for a quote. Check it out yourself or get your storage team to assess whether and how it might be helpful for your performance-challenged applications.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

➤ ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Jim Baker is a Senior Contributing Analyst for The Clipper Group. Mr. Baker specializes in storage hardware and software, the relationship between technology and its application in business, and emerging technological trends affecting users and vendors in the Information Technology community. He recently joined The Clipper Group after three decades in hardware and software product marketing and product management at companies as large as EMC and Honeywell Information Systems and as small as IPL Systems and Datapoint Corporation. He most recently served as Research Manager for Storage Software at IDC in Framingham, MA. Mr. Baker earned a Bachelor of Science in Marketing degree from the University of Illinois at Urbana-Champaign and an MBA, also from UIUC.

➤ ***Reach Jim Baker via e-mail at jim.baker@clipper.com or at 781-235-0085 Ext. 212. (Please dial "212" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and "*clipper.com*" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.