

Accelerating SAP Performance and Scalability with Kaminario K2

By Eyal Markovich
VP of Field Engineering

May 2016

TABLE OF CONTENTS

- 2** Executive Summary
- 3** The Role of the Database in SAP Performance and Scalability
- 5** Introduction of SSDs and the Kaminario K2
- 8** Identifying the Storage and Database Tier Contribution for Response Time
- 9** Identifying I/O Wait Time for SAP Database
- 11** The Results That Can Be Expected
- 12** Kaminario K2 Return on Investment (ROI)
- 12** Conclusion
- 13** About Kaminario

Executive Summary

Enterprises in every industry rely on fast access and processing of business-critical information stored in their databases to stay competitive and grow their business. SAP, being one of the most critical enterprise applications, is no different. Slow or unavailable SAP access is business-impacting. Companies around the world spend millions monitoring SAP performance and following best practices to achieve maximum performance for their SAP business applications. While SAP applications frequently consist of multiple tiers, the database tier is the foundation of the system architecture. It is not a surprise that the majority of performance problems with SAP can be traced back to slow database processing.

The different types of operations performed by databases can stress underlying servers and storage systems. Growing data volumes and compute-intensive applications are pushing databases to their limit, often slowing down the operations and increasing response time to the users. A study of the performance of hundreds of applications released by Kaminario showed 90% of slow database performance was attributed to storage bottlenecks. These findings are also true for SAP. SAP application workloads are known to stress storage and create I/O bottlenecks. Accelerating storage performance (and thus database performance) is crucial to success.

This white paper will help the reader answer the following questions:

- Does my SAP application suffer from I/O-related performance problems?
- What is the Kaminario K2?
- How can Kaminario K2 accelerate my SAP database performance?

This document is intended for SAP architects and SAP Basis, database and storage administrators that are responsible for the design, deployment and high availability and performance of SAP implementations.

The Role of the Database in SAP Performance and Scalability

The SAP Enterprise Architecture, which comprises multiple tiers (see Figure 1), allows for remote access to SAP services. As a result, the monitoring of SAP performance has become more and more complex. For example, a problem in one tier can easily impact the other tiers as well. Companies are therefore utilizing SAP and a variety of third-party performance tools to monitor and identify transaction time breakdowns across the SAP tiers.

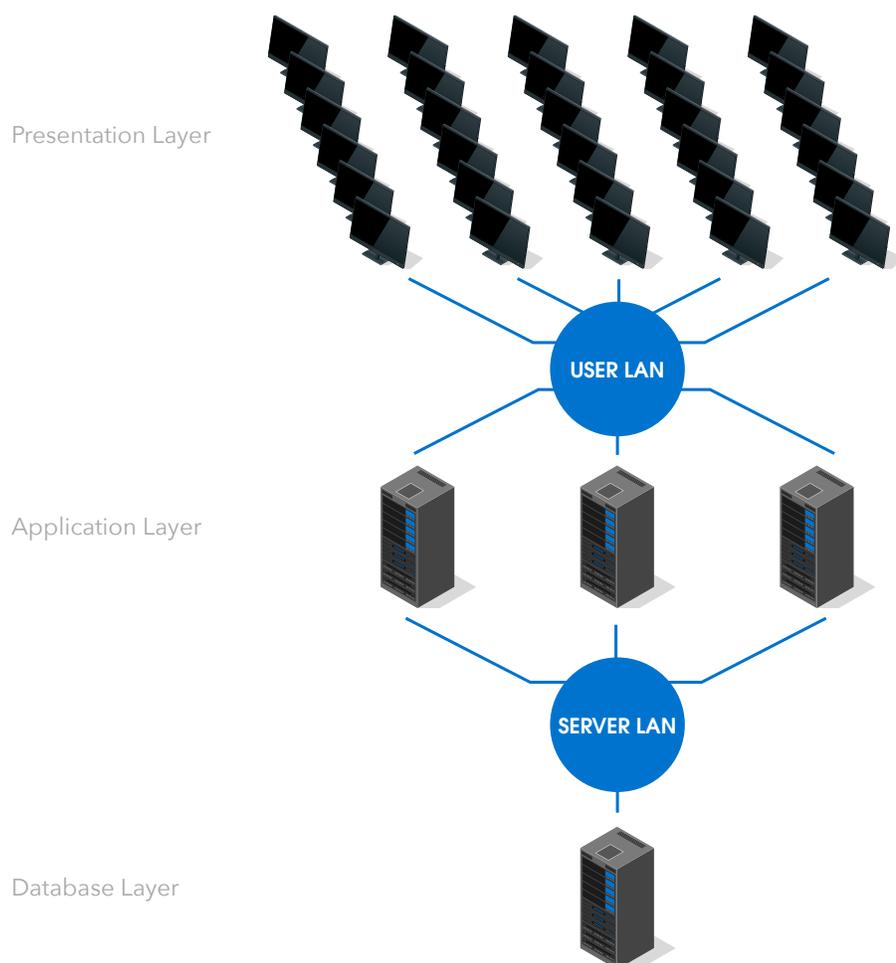


Figure 1: Three-layer model of the SAP Enterprise Architecture

The best practice in SAP performance management is to look at SAP performance from two different perspectives:

- **End user performance:** How much time does it take to complete a single step or display a single screen of online user activity? An SAP system administrator should differentiate between fast and slow Dialog steps, as we have found a different performance profile between the two. For slow-running steps, we have found that almost always the majority of the time is spent in database processing and that the storage system binds the performance of the application. Processing of fast-running steps is usually evenly distributed among the UI rendering (including network), SAP application and the database. It is therefore critical to ensure that the storage and database time is improved in order to improve end user response time.
- **Batch data processing:** The application's capability to process data in parallel is in most cases tied directly to the storage and database performance. We have found the database storage performance to be a major bottleneck for many long-running programs.

It is no surprise then that slow database performance is the dominant cause of SAP performance problems. Yet following SAP best practice will not necessarily guarantee that transaction time will be optimum. Best practices are likely to result in well-tuned database queries (no indexes missing, queries plans are optimal, database settings are correct, etc.) but the database may continue to be the number one contributor to lengthy transaction times.

Database performance is also the critical limit for SAP scalability. For many sites, even if performance today is acceptable the system will not be able to scale up to expand in size or number of users because the database will not be able to scale and support the extra workload.

So, why is the storage a bottleneck for many SAP systems?

- **Almost all queries are using indexes and there are thousands of indexes in each SAP system.** This results in a high amount of random access to disk (index seeks are random in nature) with small block sizes. With that workload, storage latency is a key to assure SAP performance and most storage arrays cannot deliver sub ms latency like the Kaminario K2.
- **Large data sets result in high physical I/O.** It is common to see SAP databases reaching several TB in size. The bigger the dataset, the more likely that I/Os will not be served from memory (logical I/O) but require physical access to disk. As a result, on a large SAP system it is common to see thousands of IOPS which most storage arrays can't support without sacrificing latency. (Unlike the Kaminario K2 that can support hundreds-of-thousands of IOPS with sub ms latency.)

- **Large sequential scans are common.** It is common for ETL, batch SAP programs and reports to trigger large table scans. It is therefore a critical requirement from the storage to support ultra-high throughput—which most storage arrays can't do.
- **SAP is a mixed workload application with large sequential scans together with online user activities.** Most storage can't support just the online user activity or the batch processing. Obviously, when combining the workloads, these storage arrays become a major bottleneck. SAP-optimized storage arrays must serve high throughput without affecting low latency in random I/O. This is also the reason why an SAP implementation using flash solutions that can't handle mixed workload well is likely to still experience high wait time on the storage.

Introduction of SSDs and the Kaminario K2

The introduction of solid-state drive (SSD) technology has changed the game for storage database performance. Solid-state technology is designed to solve the problem of I/O wait time by offering faster access times (low latency) and more IOPS. With many users and many large data files, the reduction of latency is where SSD helps the most.

Solid-state disks have the high bandwidth to sustain random data streams while traditional platter disks can only sustain high bandwidth numbers with sequential data streams. I/O bandwidth can be conceptualized as the width of the highway between the device and the database. A traditional disk is a one-lane dirt road while an SSD storage device is an eight-lane superhighway. Sustaining high bandwidth for random I/O is critical for SAP database performance, as we will show next.

SSD technology comes in different formats and solutions. The common solutions are PCI or SAS/SATA drives, Hybrid SAN arrays that include spinning platter disks and SSDs, and all-flash SAN arrays. PCI or SAS/SATA drives lack the throughput required for large scans because they are limited to a single device. These solutions can't offer SAP applications the features of high availability and cost. Scale-out is also limited. Hybrid solutions on the other hand suffer from poor performance. The natural solution for SAP applications is an all-flash SAN array that can handle a mixed workload, allow effective scalability and offer good ROI.

For SAP, since performance and high availability are critical, the natural solution is a SAN array. Kaminario K2 is the leading all-flash array and offers the best performance, scalability and resiliency.

The Kaminario K2 All-Flash Storage Array is designed to harness the power of flash with the right storage efficiency features, such as global inline selective deduplication, inline compression, thin-provisioning, an efficient and robust Kaminario RAID scheme and highly efficient metadata management. The outcome is cost-effective all-flash array storage, with better cost than HDD storage. There is no compromise on enterprise resiliency, which is gained via native snapshot and replication features, high availability (HA) and non-disruptive upgrades (NDU).



Figure 2: The Kaminario K2 All-Flash Storage Array

K2's unique scale-out and scale-up features drive business agility to the maximum, enabling independent linear growth of capacity and performance according to datacenter needs. Combined with a global adaptive block size algorithm, single management pane and full VMware integration, K2 is able to sustain the performance of multiple environments and mixed workloads while keeping it simple and easy to manage.

The K2 All-Flash Storage Array is built up from nodes called K-Blocks. The minimum system contains 2 K-Node processing units interconnected via InfiniBand and SAS-connected to a single 24-slot shelf of SSD storage. Each K-Node provides redundant Fibre Channel and Ethernet iSCSI ports for host connectivity as well as dual hot-swap power supply and battery back-up units. Volumes and metadata are automatically distributed between all SSDs in the array and can be accessed from every K-Node in the array.

Scale up expansion is accomplished by adding 24-slot expansion shelves to a K-Block. The expansion increases the capacity density and reduces the cost/GB of the array. The SAS-connected expansion shelves are added non-disruptively with no downtime or decrease in performance, and existing volumes are automatically redistributed between all the SSDs in the array with no need to change any host connections or definitions.

Scale out increases the number of K-Blocks in the array, adding more capacity and compute power. Redundant InfiniBand switches interconnect all the K-Nodes in Multiple K-Block systems, and expansion linearly increases the capacity and performance of the array. Existing volumes are automatically redistributed between all the SSDs in the array and can be accessed from every K-Node in the array.

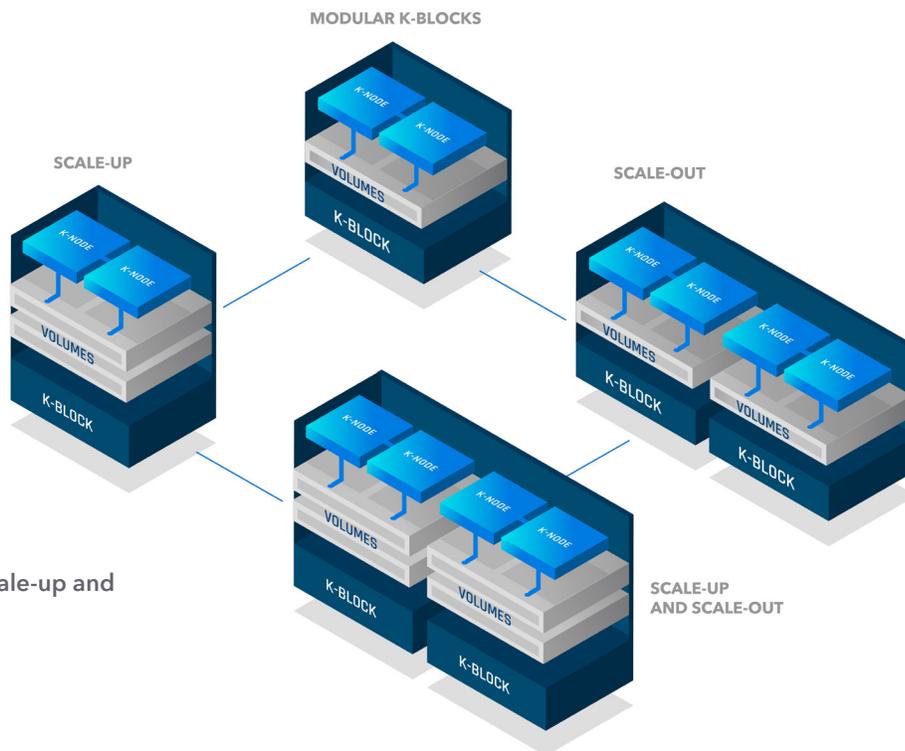


Figure 3: K2 Array scale-up and scale-out

From a software perspective, K2's SPEAR (Scalable Performance and Resilience) architecture provides the Kaminario K2 AFA with a comprehensive set of storage efficiency features, such as global inline selective deduplication, inline compression, thin-provisioning and K-RAID™ data protection. SPEAR also includes a complete software stack of enterprise resiliency features such as HA, NDU, snapshots, replication and cloud-based HealthShield™ for enterprise serviceability.

In addition, the Kaminario Management Service delivers ease of use and flexible management abilities. With scale-out and scale-up capabilities, the same ease of use and management simplicity is gained for any array size, with performance and capacity to suit the datacenter needs—driving business agility to the top.

Kaminario K2 offers SAP the perfect solution with its:

- Performance. K2 can support millions of IOPS and thousands of MB/s throughput with sub-millisecond latency. K2 support for adaptive block size allows servicing a mixed workload that is typical to SAP at premium performance.
- HA. Kaminario K2 was designed with no single point of failure and with self-healing components that insure consistent operation of SAP applications. All components are hot swappable with non-disruptive-upgrade. K2 guarantees performance under hardware failure with no more than 25% overhead, considerably more than competing all flash solutions. SAP applications will continue to operate without noticing performance degradation while K2 is recovering from a hardware failure.
- Scale-out. Kaminario K2 is a true scale-out solution, which offers SAP sites the opportunity to start small and grow both capacity and performance as needed.
- Data Reduction services. Kaminario K2 offers data reduction services via its compression and online deduplication, which allow for extra saving and reduce the cost of flash deployment while offering a quick ROI.
- Snapshots. Kaminario K2 offers the best performance snapshots of any SSD vendor. This allow SAP sites to maintain several SAP landscapes without duplicating data across all landscapes.

You can find more technical details about the K2, its architecture and its capabilities, and see a demo by going to <http://kaminario.com>

Identifying the Storage and Database Tier Contribution for Response Time

As discussed previously, we usually find the storage and the database tier to be responsible for slow SAP transaction response time. When monitoring and troubleshooting SAP performance, it is mandatory that the tier that is responsible for the majority of the transaction response time be identified. You can either use SAP transactions like ST03N or STAD or a third-party tool, but you will need to be able to determine the elapsed time attributable to each of the three tiers: UI, application time and database time.

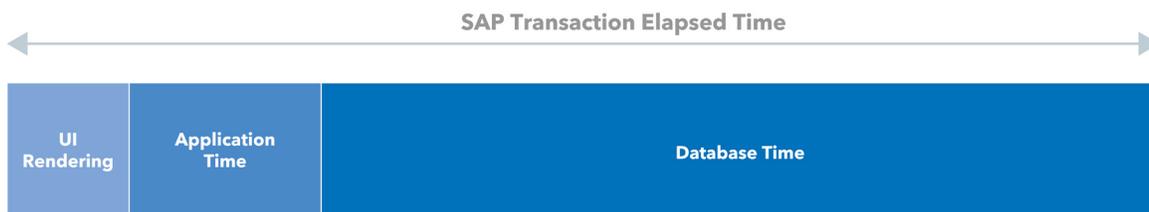


Figure 4: A typical elapsed time breakdown and the fact that the database tier is the main contributor to slow response time

It is very important to perform the analysis of elapsed time per tier by individual transactions. On an active SAP system, this analysis is likely to identify many transactions that are waiting on the database and many that are not database-constrained.

The following table taken from an SAP production database identifies several transactions in which the majority of time is spent in the database.

Transaction	Elapsed Time	DB Time	DB Percentage
/ATL/ABOOKDEPR_NEW	20938.00	9788.00	46.7%
CJ20N	65.73	64.73	98.5%
CL30N	247.76	234.01	94.5%
IP17	135.00	84.00	62.2%
IW37N	270.00	189.00	70%
MB51	1039.53	914.96	88%
MB52	317.22	265.15	83.6%

Table 1: SAP database transaction elapsed time

There is a big difference between SAP written code and custom (home) code. In many cases we find that with the performance of SAP transactions, the database response times contribute to about 1/3 of the total transaction response time, the remaining 2/3s being in the application. For custom code, it is flipped around; 2/3 database time and 1/3 application time.

The database tier being the bottleneck is not limited to SAP ERP. We have found significant database time in many SAP SCM batch processes, SAP-BW data loads and queries, and many more. Once you have identified the database as the major contributor of long transaction elapsed time, the next logical step is to identify the I/O wait time within the database.

Identifying I/O Wait Time for SAP Database

I/O wait time is the total time the database (Oracle, SQL Server, DB2, Sybase, etc.) sessions are blocked while waiting for an I/O event to complete. A storage bottleneck means that I/O wait is the dominant wait in the SAP total elapsed time.

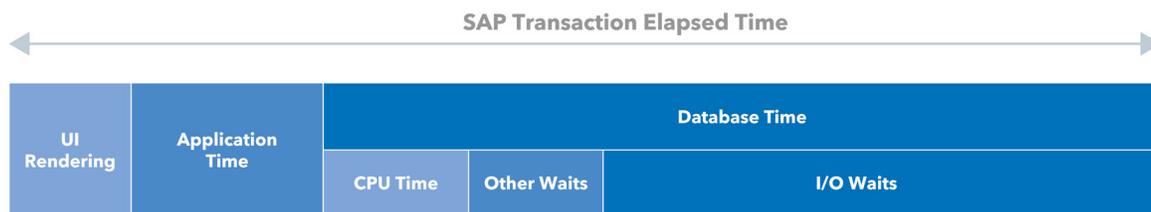


Figure 5: Typical split of database time with traditional storage arrays that shows I/O as the bottleneck

Identifying I/O wait times in the SAP database is different for each database:

- **Oracle:** it is quite straightforward, by using the standard Automatic Workload Repository (AWR) report for the period you want to study. There are several views within the report that report storage and I/O information, but the one to start with is the Foreground Wait Class table. This table identifies the total waits associated with the storage. The next interesting table within the AWR report is the Foreground Wait Events, which lists the actual wait events in Oracle. With the list of wait events, you can better understand the effect of the Kaminario K2 on your system. Several waits are I/O-related waits that should be analyzed. The most common waits are *db file sequential read*, *db file scattered read*, *direct path read*, *direct path read temp*, *direct path write*, *direct path write temp*, *free buffer wait*, *log sync* and *read by other session*.
- **DB2:** DB2 9.7 introduces a new performance snapshot mechanism called monreport. With the monreport it is quite straightforward to identify if the storage is a bottleneck.
- **SQL Server:** Storage I/O performance can be obtained from server Dynamic Management Views (DMV). Especially, `sys.dm_os_wait_stats` gives a wait breakdown of SQL Server times. Several waits are I/O-related waits that should be analyzed. The most common waits are PAGEIOLATCH* waits, WRITELOG, IO_COMPLETION and several others.

For an improvement in application performance, these I/O wait times need to be dealt with. The Kaminario K2 is specifically designed to address these issues.

The Results That Can Be Expected

The expected results from a Kaminario K2 will vary from one application to another. The general rule is that the higher the portion of I/O waits within the application, the higher the observed improvements from the move to the K2 SAN.

Event	Waits		Total Time(s)		Avg Wait (ms)		Wait Class
	Before K2	With K2	Before K2	With K2	Before K2	With K2	
CPU Time					12.0%	62.0%	
DB File Sequential Read	4,338,155	4,868,422	4	0	68.3%	23.9%	User I/O
DB File Scattered Read	802,380	615,694	3	1	11.0%	5.0%	User I/O
Read by Other Session	395,427	420,483	3	0	6.0%	0.4%	User I/O
Log File Sync	44,451	65,488	6	1	1.0%	0.3%	Commit
Log File Parallel Write	46,605	50,717	4	2	1.0%	1.6%	System I/O

Table 2: Top 5 timed events observed before and with Kaminario K2

Table 2 highlights the improvements of using K2 for SAP on Oracle when compared to the wait times before Kaminario. CPU utilization was increased dramatically from 12% to 62% as I/O waits shrank from 87% to 37%. As expected, the improvement can be attributed to the improved latency of the Kaminario K2 compared to the previously used SAN. The following graph displays the performance of SAP production before and after the K2 implementation.

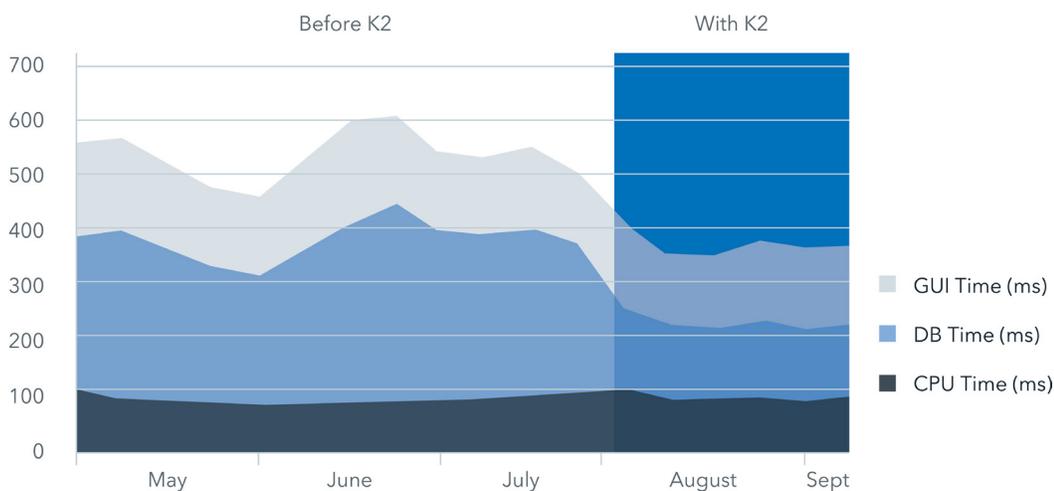


Figure 6: K2 effect on SAP production

The reduction of the database time denoted in Figure 6 is clear, while the UI and application times were not changed.

Kaminario K2 Return on Investment (ROI)

For SAP systems having a storage bottleneck, the ROI from a K2 implementation is very high and very quick:

- When SAP suffers so does the business. With K2, companies can expect major performance improvements across all SAP programs with an immediate return on the K2 investment.
- When an SAP application is not performing, enterprises invest millions of dollars squeezing efficiency out of every bit of each SQL statement. Companies spend thousands of hours optimizing ABAP code and controlling program/reports execution. In many cases, the dollars invested in these pursuits could be significant but the return on that investment could be disappointing. Indeed, SQL tuning can help improve performance to some extent, but even the best-written SQL cannot compensate for the poor storage I/O problems. It is common to see up to 30% of SAP Basis team manpower allocated for SAP performance and optimization. It is also common for companies to control access to SAP functionality such as reports or data processed by large batch programs. With Kaminario K2, the overall performance of the SAP system will be dramatically faster. Reports and batches will complete much faster and companies can allocate SAP resources to other areas beside performance.
- Server and processor problems. One of the first things IT shops do when performance wanes is add processor or memory to the database server. This approach can help if I/O wait is not the dominant problem. It is however, a catalyst for discovering I/O bottlenecks since with faster servers, more time is spent waiting for the storage. This problem is resolved with Kaminario K2

Conclusion

Today, businesses around the world spend significant amounts of time and money optimizing their environments to meet growing storage performance requirements. In this white paper, we discussed the Kaminario K2 as a game changer and how it is addressing this challenge for businesses with SAP environments.

Kaminario K2 All-Flash SAN Storage is built from off-the-shelf hardware and architected from the ground up for solid-state technology. It delivers breakthrough performance in a fast, simple and easy way. It integrates easily into existing architectures, without the need to re-architect the storage infrastructure or application. If you are an SAP customer that needs performance acceleration at economic levels, Kaminario K2 is your answer.



Contact

Contact a business development representative to answer any questions you may have.



Schedule a Demo

Schedule a demo with an engineer and learn if Kaminario's solution works for you.



Request a Quote

Request a quote for your application from our business development team.

About Kaminario

Kaminario, the leading all-flash storage company, is redefining the future of modern data centers. Its unique solution enables organizations to succeed in today's on-demand world and prepares them to seamlessly handle tomorrow's innovations. Only Kaminario K2 delivers the agility, scalability, performance and economics a data center requires to deal with today's cloud-first, dynamic world and provide real-time data access -- anywhere, anytime. Hundreds of customers rely on the Kaminario K2 all-flash array to power their mission critical applications and safeguard their digital ecosystem. Headquartered in Needham, MA, Kaminario works with an extensive network of resellers and distributors, globally.

For more information, visit www.kaminario.com

Kaminario, the Kaminario logo and SPEAR are registered trademarks of Kaminario Inc. K-RAID, Perpetual Array and HealthShield are trademarks of Kaminario, Inc. Product specifications and performance are subject to change without notice.