

Kaminario K2 Best Practices For SAP HANA

October 2016

TABLE OF CONTENTS

- 2** Executive Summary
- 2** What is SAP HANA?
- 3** Introduction of SSDs and the Kaminario K2
- 4** Kaminario K2 Scalability
- 5** Kaminario K2 Reliability and Availability
- 6** Kaminario K2 Specifications
- 7** Solution Overview
- 19** Kaminario K2 Sizing and Scalability
- 19** Conclusion
- 20** 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 HANA, being one of the most critical enterprise databases, is no different. Slow or unavailable SAP HANA access is business-impacting. The combination of the Kaminario K2 all-flash storage array and SAP HANA allows enterprises to analyze and access its most critical data in a reliable and fast way.

This best practices document will help the reader to:

- Understand the different components of SAP HANA and Kaminario K2
- Configure SAP HANA to work with Kaminario K2

What is SAP HANA?

SAP HANA is an in-memory computing platform that enables enterprises to accelerate business processes, deliver more intelligence, and simplify the IT environment. By providing the foundation for all data needs, SAP HANA removes the burden of maintaining separate legacy systems and siloes of data, allowing real-time business decisions in the new digital economy.

SAP HANA TDI Overview:

SAP HANA tailored data center integration (TDI) was released in November 2013 to offer an additional approach to deploying SAP HANA. While the deployment of an appliance is easy and comfortable for customers, appliances impose limitations on the flexibility of selecting the hardware components for compute servers, storage, and network. Furthermore, operating appliances may require changes to established IT operation processes. For those who prefer leveraging their established processes and gaining more flexibility in hardware selection for SAP HANA, SAP introduced SAP HANA TDI.

More information on SAP HANA TDI, please refer to the following link:

<http://go.sap.com/documents/2016/05/827c26ba-717c-0010-82c7-eda71af511fa.html>

Introduction of SSDs and the Kaminario K2

The introduction of solid-state drive (SSD) technology has changed the game for storage performance. Solid-state technology is designed to solve the problem of throughput and I/O wait and time by offering faster access times (low latency) and more IOPS. With many users and many large data files, the increasing of throughput and 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. Scale-out is also limited. Hybrid solutions, on the other hand, suffer from poor performance. The natural solution for SAP HANA applications is an all-flash SAN array that can handle mixed workloads, allow effective scalability and offer good ROI.

For SAP HANA, the natural solution is a SAN array, since performance and high availability are critical. 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 a cost-effective all-flash storage array with better cost than HDD storage. There is no compromise on enterprise resiliency, which is gained via full hardware redundancy, native snapshot and replication features, high availability (HA) and non-disruptive upgrades (NDU).

K2's unique scale-out and scale-up capabilities drive business agility to the maximum, enabling independent linear growth of capacity and performance according to datacenter needs and specifically to SAP HANA's scale-out model.

The adaptive block size algorithm makes Kaminario K2 the perfect fit for multiple applications or an application such as SAP HANA. The ability to handle any IO size allows Kaminario K2 to deliver consistent low latency for every write.

From a software perspective, K2 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. Kaminario K2 also includes a complete software stack of enterprise resiliency features such as HA, NDU, snapshots, replication and cloud-based Kaminario Clarity™ for enterprise serviceability and app-aware analytics.



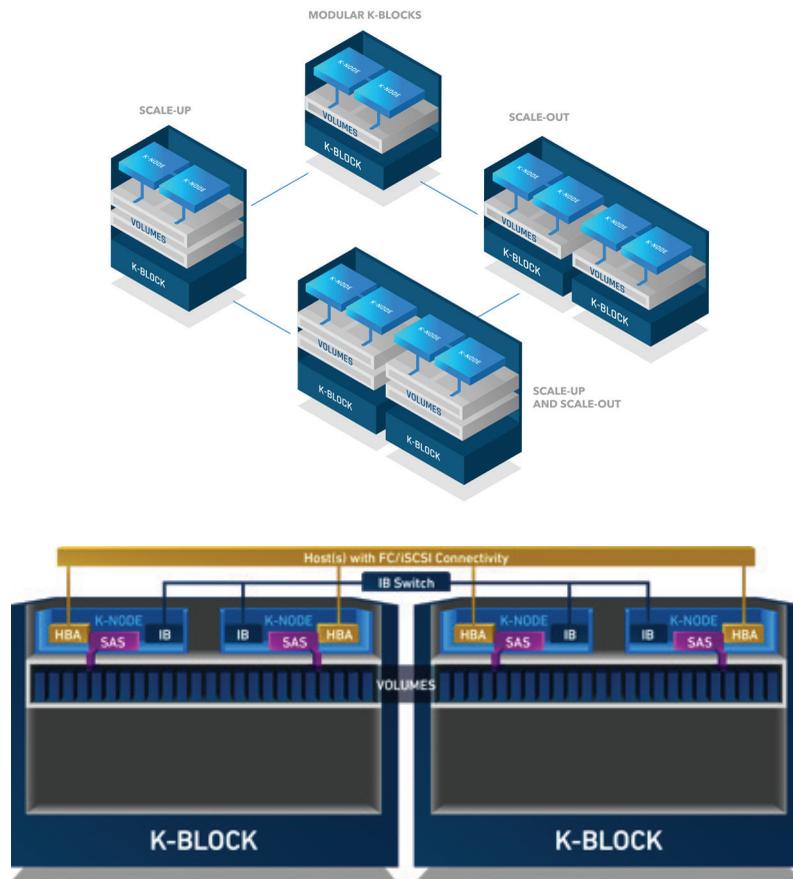
Figure 1: The Kaminario K2 All-flash Array

Kaminario K2 Scalability

The K2 all-flash storage array is built up from K-Blocks. Each K-Block contains two K-Nodes, which are the storage controller processing units interconnected via InfiniBand and SAS-connected to one or more SSD shelves. Each K-Node provides redundant Fibre Channel and Ethernet iSCSI ports for host connectivity as well as dual hot-swap power supply and battery backup 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 an expansion shelf 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.



For more information about the K2, its architecture and its capabilities, please [visit kaminario.com](http://www.kaminario.com).

Kaminario K2 Reliability and Availability

Reliability and Availability

High Availability (HA) – or Continuous Availability – is a commodity with enterprise storage arrays. Through Kaminario software, the K2 gains five 9s of availability, and its scalability does not impact this availability since every building block is treated as an independent failure domain. Kaminario K2 provides the following:

K-RAID

Aside from being highly efficient, the RAID is extremely robust. It can sustain two concurrent SSD failures and up to three SSD failures within a single SSD shelf. As the K2 scales capacity, the number of system-wide SSD failures that the system can sustain scales as well. The RAID has a dual parity protection that adapts according to the failure at hand. An SSD failure is quickly recovered thanks to efficient metadata and real-time system health monitoring, and there is no performance impact on the array's performance once the rebuild is completed.

No Single Point of Failure (SPoF)

Kaminario K2 maintains a double-everything approach. However, the K2 does not have passive or idle components in the array; all resources are utilized at all times. There is full redundancy of every component in the system, and there is no single component that can fail and cause unplanned downtime or data loss. Each building block is a standalone failure domain, which means that the entire array can sustain more failures as the system scales.

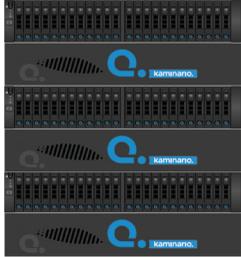
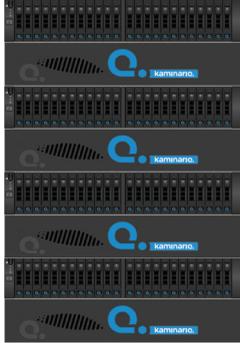
Power loss

A storage array must be able to sustain a power outage in the datacenter and still keep the data intact and available for when power returns. Any metadata and/or data that are accounted for within the storage controllers before being stored to the RAID are saved in two places for redundancy. Each storage controller in the K2 is equipped with an internal battery that provides sufficient power for the controller to destage any inflight data that hasn't been stored to the RAID. These batteries do not occupy any space within the rack since they are internal. Data that is already stored to the RAID is kept persistent and is sustainable through power cycles.

Proactive Support

The K2 continuously monitors all of its components, and is able to detect in real time any error that might occur. The K2 also collects system-wide information that is sent daily to Kaminario for further inspection. This information is reported back to Kaminario Support, where it is analyzed and processed to recognize issues of performance, configurations, capacity management and more. Based on these reports, the support team can proactively deliver support tailored to customer needs.

Kaminario K2 Specifications

	1 K-Block	2 K-Block	3 K-Block	4 K-Block
				
Footprint K-Blocks	4U	10U	14U	18U
Footprint +1 Exp	6U	14U	20U	26U
Footprint +2 Exp	8U	18U	26U	34U
Usable Capacity*	7TB-360+TB	14TB-720+TB	21TB-1.08+PB	28TB-1.44+PB
Media (encrypted)	480GB 3D MLC SSD / 960GB 3D MLC SSD / 1.92TB 3D TLC SSD			
Density TB/U	Up to 60TB			
IOPS	Up to 250K	Up to 500K	Up to 750K	Up to 1M
Bandwidth	Up to 3.2GB/s	Up to 6.4GB/s	Up to 9.6GB/s	Up to 12.8GB/s
Latency	0.35mSec			

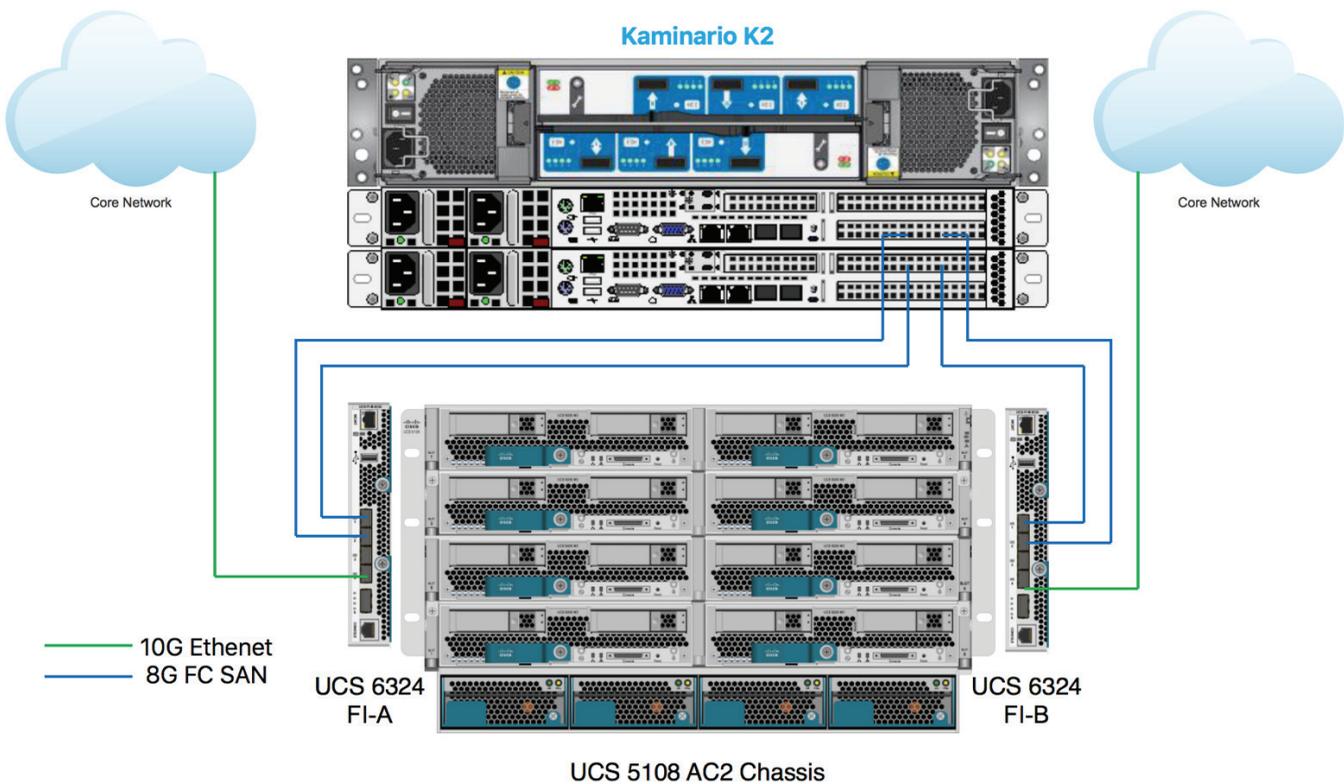
*Capacity is subject to drive size and the application data reduction ratio. For some datasets such as VDI, the range will be higher

Solution Overview

SAP has certified that Kaminario K2, the leading all-flash storage company, meets all the functionality and performance requirements for SAP HANA TDI solutions. Following the guidelines in this document, and all SAP HANA tailored data center integration best practices, provides a full supported environment.

Topology and Connectivity

For the certification, the configuration described below was used.



Servers: 3 x Cisco UCS B260 M4

RAM: 1TB

CPU: Intel® Xeon® CPU E7-8880 v3 @ 2.30GHz

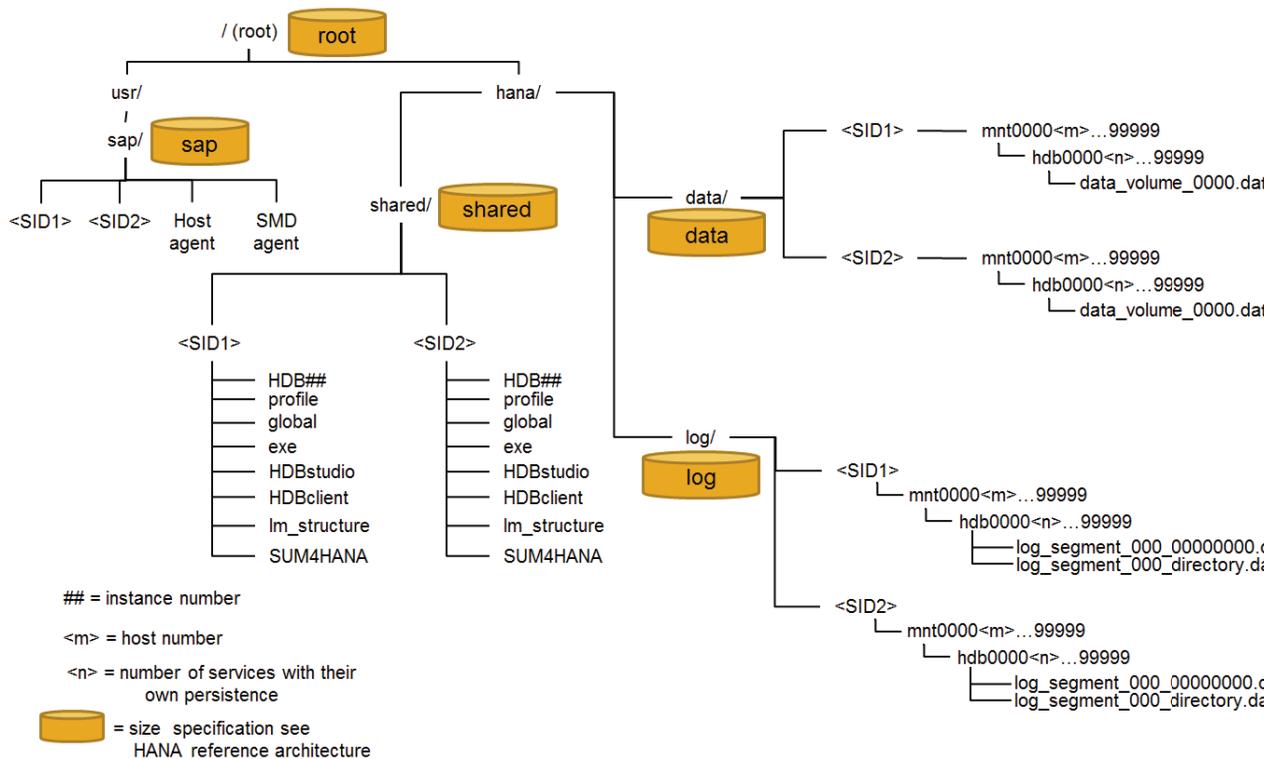
Ethernet/SAN switches: Cisco UCS Fiber Interconnect 6324

OS: SUSE Linux Enterprise Server for SAP Applications 12 SP1. Kernel 3.12.62-60.62-default

SAP HANA: Version -1.00.120.00.1462275491

SAP HANA File System Requirements and Layout

SAP HANA uses storage for several purposes. The layout of the different file systems is displayed in the diagram below.



The SAP HANA installation - In a SAP HANA scale-out configuration, SAP has certified Kaminario K2 with the OCFS2 clustered file system for HANA compute nodes with the SUSE Linux Enterprise Server operating system. This type of file system must be mounted on /hana/shared and exposed to all the nodes in the cluster. The HA package included in the SUSE Linux Enterprise Server for SAP Applications provides all the necessary services for such a configuration.

DATA file system - Access to the data file system is primarily random, with various block sizes ranging from small (4K) to large (64M) blocks. The data is written asynchronously with parallel I/O to the data file system. During normal operations, most of the data file system I/O operations are writes. Data is read from the data file system only during a database restart, HA failover or column store table load.

LOG file system - To ensure the recovery of the database with zero data loss in case of faults, SAP HANA records each transaction in the form of a so-called redo log entry. Each SAP HANA service separately writes its own redo- log files. Typical block-write sizes range from 4KB to 1MB.

In a **SAP HANA** scale-out cluster configuration, the volumes containing the file systems for DATA and LOG must be exposed to all the nodes in the HANA cluster, as described in the "Kaminario configuration" section of this document.

The size for HANA SHARED, DATA, LOG file systems must be calculated according to SAP requirements described on [SAP HANA TDI - Storage Requirements \(http://scn.sap.com/docs/DOC-62595\)](http://scn.sap.com/docs/DOC-62595).

SAP HANA Network Requirements

SAP HANA uses multiple network channels to transfer data between different components and modules. It is recommended practice to have a well-defined network topology.

The components belonging to SAP HANA communicate over several logical network zones:

Client Zone - Different clients, such as SQL clients on SAP application servers, browser applications using HTTP/S to the SAP HANA XS server and other data sources (such as BI) need a network communication channel to the SAP HANA database.

Internal Zone - The internal zone covers the communication between hosts in a distributed SAP HANA system as well as the communication used by SAP HANA system replication between two SAP HANA sites.

Storage Zone - As described in the previous section of this document, Kaminario K2 solution is certified with OCFS2 file system for the SAP HANA installation. For this use case, two main components are part of the storage zone:

- FC-SAN configuration: to learn about the recommended practices on how to setup the SAN network using Kaminario K2 please refer to the next section, 'Connectivity' under 'Kaminario Configuration'.
- OCFS2 file system networks: To ensure that OCFS2 traffic network including its heartbeat is not routed through the Internal Zone network, which is exclusively for inter node communication, it is required to have separate networks for the internal zone, the OCFS2 and OCFS2 heartbeat related traffic.

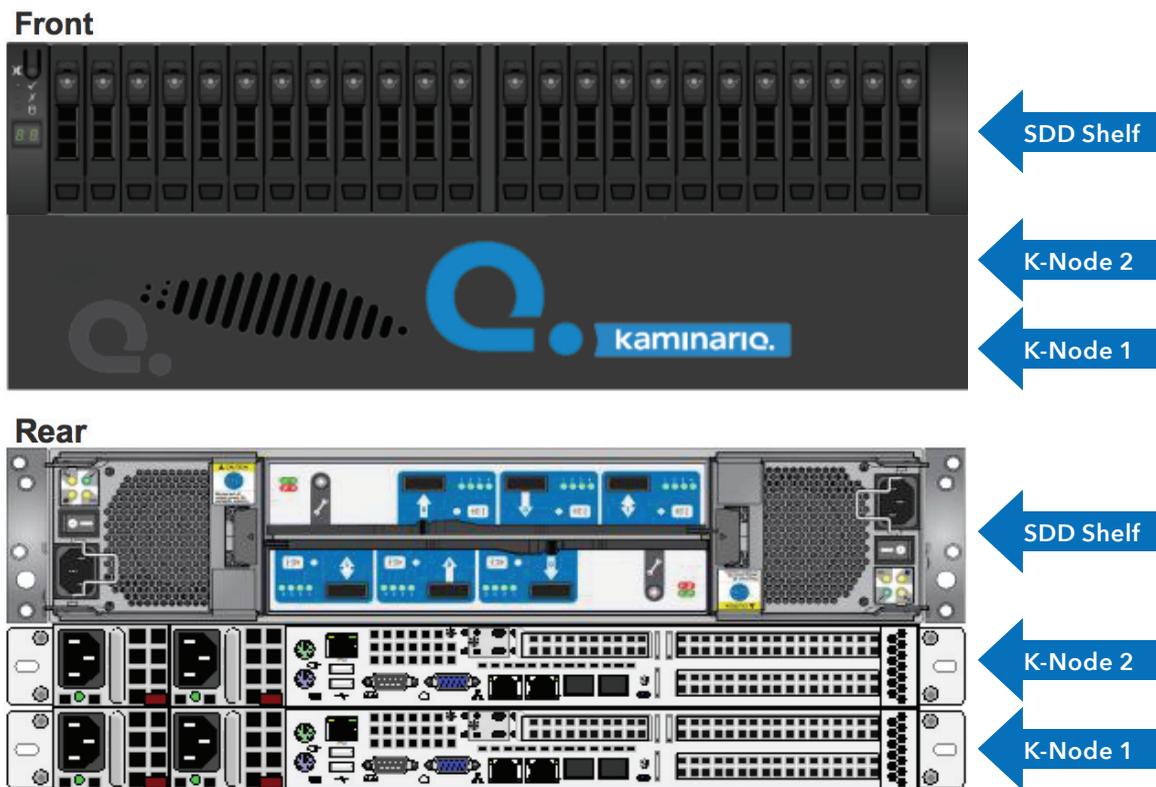
Each communication zone needs special attention with regard to configuration, setup, security and performance requirements. For a complete guide to the different network configuration options please refer to the [SAP HANA Network Requirements](http://www.sdn.sap.com/irj/scn/index?rid=/library/uuid/e0bf7501-e4bd-3210-9e82-c7e200e7943a&overridelayout=true) document (<http://www.sdn.sap.com/irj/scn/index?rid=/library/uuid/e0bf7501-e4bd-3210-9e82-c7e200e7943a&overridelayout=true>).

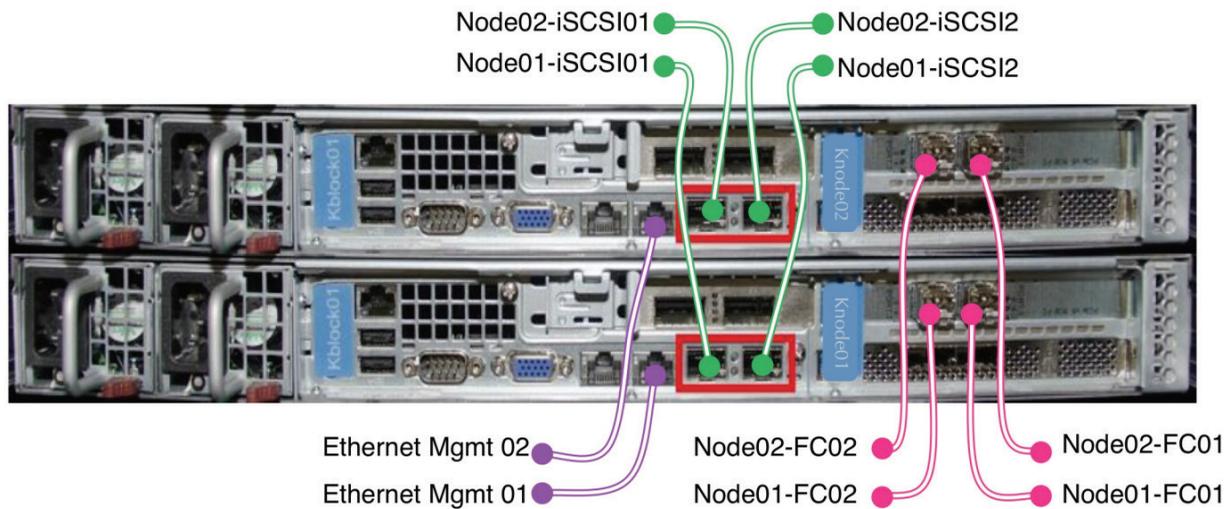
Kaminario Configuration

Connectivity

- Kaminario K2 connectivity relies on the K2 array configuration. Each K-Node enables 2 x 8Gb Fibre Channel ports.
- 1 K-Block (2 x K-Nodes) configuration will deliver 4 x FC 8Gb ports
- 2 K-Blocks (4 x K-Nodes) configuration will deliver 8 x FC 8Gb ports
- K-Blocks (6 x K-Nodes) configuration will deliver 12 x FC 8Gb ports
- 4 K-Blocks (8 x K-Nodes) configuration will deliver 16 x FC 8Gb ports

Each K-Node delivers both internal and external connectivity, as shown in the figure below.





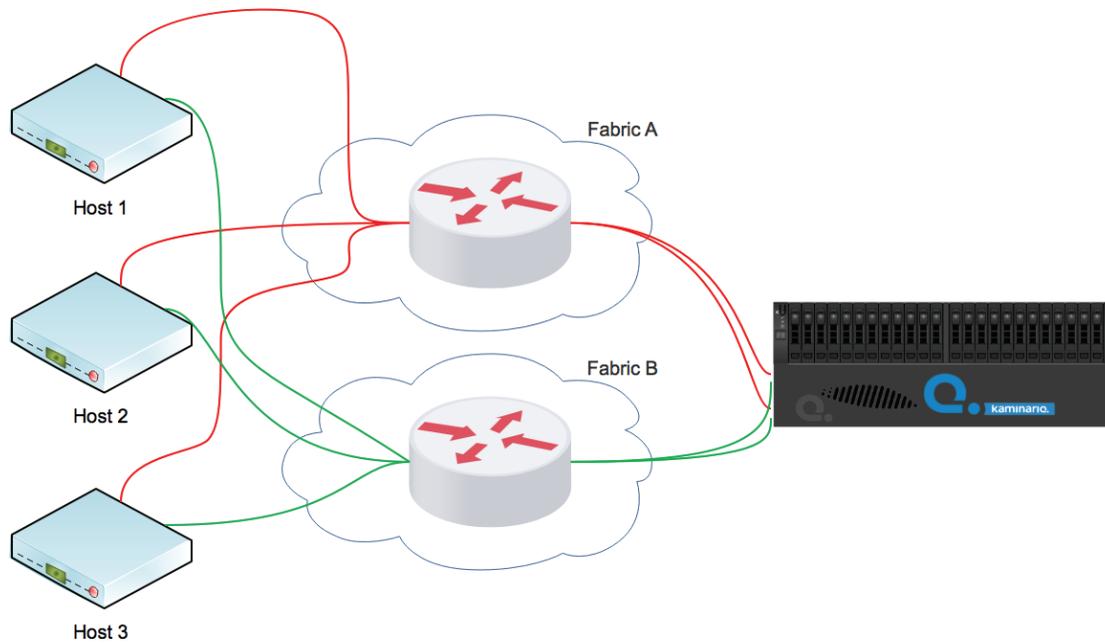
The way an administrator connects the Kaminario K2 system to the data center may significantly impact the K2's redundancy, performance and availability. Following the recommendations on how to connect the Kaminario K2 system to the data center ensures maximum availability, redundancy and performance.

- Avoid creating bottlenecks in the infrastructure network. Bottlenecks can be easily created almost without noticing.
- Utilize multipath solutions for hosts connected to the K2. Multipath solutions can use underutilized paths, thus reducing bottlenecks and improving bandwidth.
- Include enterprise grade switches.
- Connect each K-Node port to a different fabric for high availability.
- For ideal latency performance in a Core-Edge layout, connect the Kaminario K2 directly to the core switch.
- Configure zones in a single initiator - single target configuration per zone. For additional information, review the examples on the next page.

1 K-Block Zoning Example

In a One K-Block configuration, there are a total of four Fiber Channel ports available.

In this example, a three-node cluster (2 +1) is formed out of Host 1, Host 2 and Host 3.



In a One K-Block configuration

Fabric A

Host 1:

- Zone Name: Z_Host_1_Fabric_A_K500_KBLOCK01_KNODE01_P1
 - Member: Host 1 HBA Port A
 - Member: K2 KBLOCK01_KNODE01 Port 1
- Zone Name: Z_Host_1_Fabric_A_K500_KBLOCK01_KNODE02_P1
 - Member: Host 1 HBA Port A
 - Member: K2 KBLOCK01_KNODE02 Port 1

Host 2:

- Zone Name: Z_Host_2_Fabric_A_K500_KBLOCK01_KNODE01_P1
 - Member: Host 2 HBA Port A
 - Member: K2 KBLOCK01_KNODE01 Port 1
- Zone Name: Z_Host_2_Fabric_A_K500_KBLOCK01_KNODE02_P1
 - Member: Host 2 HBA Port A
 - Member: K2 KBLOCK01_KNODE02 Port 1

Host 3:

- Zone Name: Z_Host_3_Fabric_A_K500_KBLOCK01_KNODE01_P1
 - Member: Host 3 HBA Port A
 - Member: K2 KBLOCK01_KNODE01 Port 1
- Z_Host_3_Fabric_A_K500_KBLOCK01_KNODE02_P1
 - Member: Host 3 HBA Port A
 - Member: K2 KBLOCK01_KNODE02 Port 1

Fabric B**Host 1:**

- Zone Name: Z_Host_1_Fabric_B_K500_KBLOCK01_KNODE01_P2
 - Member: Host 1 HBA Port B
 - Member: K2 KBLOCK01_KNODE01 Port 2
- Zone Name: Z_Host_1_Fabric_B_K500_KBLOCK01_KNODE02_P2
 - Member: Host 1 HBA Port B
 - Member: K2 KBLOCK01_KNODE02 Port 2

Host 2:

- Zone Name: Z_Host_2_Fabric_B_K500_KBLOCK01_KNODE01_P2
 - Member: Host 2 HBA Port B
 - Member: K2 KBLOCK01_KNODE01 Port 2
- Zone Name: Z_Host_2_Fabric_B_K500_KBLOCK01_KNODE02_P2
 - Member: Host 2 HBA Port B
 - Member: K2 KBLOCK01_KNODE02 Port 2

Host 3:

- Zone Name: : Z_Host_3_Fabric_B_K500_KBLOCK01_KNODE01_P2
 - Member: Host 3 HBA Port B
 - Member: K2 KBLOCK01_KNODE01 Port 2
- Zone Name: Z_Host_3_Fabric_B_K500_KBLOCK01_KNODE02_P2
 - Member: Host 3 HBA Port B
 - Member: K2 KBLOCK01_KNODE02 Port 2

Volume Creation and Mapping

Volumes, or logical units (LUs), are the basic unit of the Kaminario K2 appliance. It is not necessary to perform any optimization or apply any consideration when volumes are created. The system offers load-balancing, which is performed regardless of the physical setup and enables the system to work optimally across all K-Nodes.

A volume group is an entity that defines a group of volumes with common properties such as an application, departmental information, and so on. A volume group has three functions:

- Capacity management
- A volume group for snapshots
- Replication sessions can only be defined on volume groups

Each volume group has a capacity quota with either a default or user-defined capacity threshold policy profile. When the capacity reaches one of the thresholds, the action stated in the properties of the capacity policy is triggered.

The volume group capacity contains the volumes and snapshots of that volume group. When taking a snapshot, the entire volume group is captured in the snapshot.

The volume group maintains a special threshold for snapshot capacity allocation and sends event notifications to the storage admin when snapshot usage exceeds the threshold.

The volume group is the basic consistency group for replication. When replication is set for a specific volume group, its volumes are replicated to a different K2 system using snapshots.

Several volumes must be configured in the K2 for SAP HANA. The operating system image can be installed on a local disk or on a Kaminario LUN and booted from the SAN. Sizes are calculated based on [SAP HANA TDI - Storage Requirements \(http://go.sap.com/documents/2015/03/74cdb554-5a7c-0010-82c7-eda71af511fa.html\)](http://go.sap.com/documents/2015/03/74cdb554-5a7c-0010-82c7-eda71af511fa.html)

- HANA-DATA-Node1 and HANA-DATA-Node2, which are part of the volume group HANA-Data, are the volumes for the 'DATA' file systems
- HANA-LOG-Node1, which is part of volume group HANA-Log-node1, is the volume for the LOG file system mounted initially on node1 of the cluster
- HANA-LOG-Node2, which is part of volume group HANA-Log-node1, is the volume for the LOG file system mounted initially on node2 of the cluster
- HANA_Shared_1, which is part of volume group HANA-Shared-volumes, is the volume for the '/hana/shared' file system

All volume groups are created with the provisioning type "Thin provisioning".

Provisioning of volumes and volume groups is executed in the Kaminario K2 management GUI using the "Volume" section, located in the "Volumes & Hosts" tab.

To create a new volume group, click the "+ Volume Group" button and complete the details as shown in the example below:

Create Volume Group

Name:

[Add description](#)

Quota: Unlimited quota GB

Provisioning Type:

Capacity Policy:

To create a new volume, click the "+ Volume" button and complete the details as shown in the example below:

Create Volume(s)

Multiple Volumes

Name:

[Add description](#)

Volume Group:

Provisioned Size:

VMware support: Enable VMFS (SCSI-2 Reservations)

All volumes must be mapped to all three hosts:

kaminario.k2 Server time: 4:12 PM System: K2-5552 Online admin Logout

Dashboard Performance Events **Volumes & Hosts** DataProtect Connectivity System

Hosts Volumes 0% of Capacity Used

Host	Mapped to Volumes	Size	LUN
<ul style="list-style-type: none"> Kaminario-SAP Kaminario-node1 Kaminario-node2 Kaminario-node3 + Add... 	HANA-Data-Node1	0 of 6TB	LUN: 1
	HANA-Data-Node2	0 of 6TB	LUN: 2
	HANA-Shared	0 of 3TB	LUN: 3
	Log-node1	0 of 1TB	LUN: 5
	Log-node2	0 of 1TB	LUN: 6
	Shared-SBD-3G	0 of 3GB	LUN: 4

[+ Map...](#)

Multipath and IO Scheduler Configuration

Multipathing must be configured properly to ensure high availability and performance. Edit `/etc/multipath.conf` on each node of the cluster as shown below:

```
defaults {
    polling_interval 1
    # user_friendly_names yes
}
devices {
    device {
        vendor "KMNRI0"
        product "K2"
        path_grouping_policy multibus
        uid_attribute "ID_SERIAL"
        path_checker tur
        path_selector "round-robin 0"
        no_path_retry 0
        features "0"
        hardware_handler "0"
        rr_weight priorities
        rr_min_io 1
        failback 15
        fast_io_fail_tmo 5
        dev_loss_tmo 8
    }
}
blacklist {
    devnode "^(/ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]"
    devnode "^hd[a-z][[0-9]]"
    devnode "^cciss!c[0-9]d[0-9]"
    devnode "^sda[0-9]"
}
```

Edit `/etc/init.d/after.local` and add the line below to set the IO schedule to 'noop':

```
for i in /sys/block/sd*/queue/scheduler; do /usr/bin/echo noop > $i; done
```

SAP HANA global.inir

During the certification process the global.ini was modified as follow:

```
listeninterface = .global

[internal_hostname_resolution]
10.190.42.200 = kaminario-node1
10.190.42.201 = kaminario-node2
10.190.42.202 = kaminario-node3

[persistence]
basepath_datavolumes = /hana/data/KM1
basepath_logvolumes = /hana/log/KM1
use_mountpoint = yes

[storage]
ha_provider = hdb_ha.fcClient
partition_*_*_prtype = 5
partition_1_data__wwid = 20024f40055520001
partition_1_log__wwid = 20024f40055520003
partition_2_data__wwid = 20024f40055520002
partition_2_log__wwid = 20024f40055520004

[system_information]
usage = test

[trace]
ha_fcclient = info
```

The "partition_*_*_prtype" value must be set to "5".

HDPARAM Parameter Vector

The default parameter vectors for hdpparam were used for the certification setup:

```
async_read_submit = ON
```

```
async_write_submit_active = AUTO (Decided based on filesystem)
```

```
async_write_submit_blocks = ALL
```

To confirm that the correct values are configured, the following command can be executed:

```
echo "Using hdbparam -> get actual configuration (short)"; su $SIDADM -c "source hdbenv.sh; hdbparam -p a | grep -A 1000 '1 async_read_submit' | grep -e [[0-9]] -e ' Value' -e ' Index'
```

The output will show the values as in the example below:

1 async_read_submit [9]: "Activate asynchronous read requests"

(1) Value: [string] on

2 size_kernel_io_queue [3]: "Size of kernel IO queue"

(1) Value: [int] 512

3 max_submit_batch_size [6]: "Max. submit batch size"

(1) Value: [int] 64

4 num_submit_queues [1]: "Number of submit queues"

(1) Value: [int] 1

5 async_write_submit_active [7]: "Activate asynchronous write requests"

(1) Value: [string] auto

6 network_log_file [1]: "Log file to store events of the network traffic"

(1) Value: [string] network.log

7 consolidated_channels [1]: "Flag to employ the consolidated network stack model"

(1) Value: [bool] true

8 async_write_submit_blocks [8]: "Activate asynchronous write requests for certain blocks"

(1) Value: [string] all

9 min_submit_batch_size [5]: "Min. submit batch size"

(1) Value: [int] 16

10 num_completion_queues [2]: "Number of completion queues"

(1) Value: [int] 1

11 max_parallel_io_requests [4]: "Max. number of parallel IO requests"

(1) Value: [int] 64

For more information please refer to <https://launchpad.support.sap.com/#/notes/2267798>.

Kaminario K2 Sizing and Scalability

Kaminario's unique scale-up and scale-out capabilities are the perfect fit for the SAP HANA scale-out model. For existing environments or new HANA deployments, Kaminario's flexible scalability allows economical and efficient options for sizing and expanding. The table below shows the maximum number of SAP HANA nodes recommended for each of four Kaminario configurations:

1 x K-Block	UP to 3 SAP HANA nodes
2 x K-Blocks	UP to 6 SAP HANA nodes
3 x K-Blocks	UP to 10 SAP HANA nodes
4 x K-Blocks	UP to 14 SAP HANA nodes

If the Kaminario K2 is being used to run other workloads, it is recommended to run the SAP HANA HWCCT tool to confirm that performance is not impacted for the other workloads.

Conclusion

SAP HANA users rely on fast access and processing of business-critical information stored in it to stay competitive and grow their business. Slow or unavailable SAP HANA access is business-impacting. Consistent low latency, scalability and features such as adaptive block size make the Kaminario K2 a superior choice for SAP HANA deployments.

Kaminario K2 offers SAP HANA 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 HANA.
- **HA:** Kaminario K2 was designed with no single point of failure and with self-healing components that ensure consistent operation of SAP applications. All components are hot swappable with non-disruptive-upgrade.
- **Scale-out:** Kaminario K2 is a true scale-out solution, which offers SAP HANA 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 savings 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 allows SAP HANA sites to maintain several SAP HANA copies without duplicating data across all landscapes.



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.