

Best Practices for Running Oracle on Kaminario

June 2018

TABLE OF CONTENTS

- 2** Overview
- 2** Oracle Database
- 3** Oracle Database Storage Best Practices
- 5** Additional Best Practices



Overview

The Kaminario K2 all-flash array is built to support modern workloads that require efficient use of resources, high performance and flexibility in scale. Oracle databases are widely used in the industry, and are an ideal workload for the K2's all flash architecture. Kaminario has a unique value proposition for Oracle implementations, and it is recommended that best practices and reference guides are followed in order to maximize capacity and performance efficiencies.

This best practices guide describes in detail all integration points of the K2 all-flash array and Oracle databases. The document also includes examples and screenshots that facilitate easy understanding of the different settings and clearly demonstrate the tradeoffs between different options of configurations. Content in this document is quite technical and is intended to be used by pre-sales engineers, system and storage administrators and customers who want to deploy Oracle databases on the Kaminario K2.

Oracle Database

The Oracle Database is a relational database management system (RDBMS) from the Oracle Corporation. One of the most trusted and widely-used relational database engines worldwide, Oracle is a fully scalable database and it is often used by global enterprises to hold their most crucial and critical data. It is known to be a very popular database for running online transactional processing (OLTP), data warehousing (DW) and mixed (OLTP & DW) database workloads.

Oracle Database Storage Best Practices

The following table presents the traditional Oracle storage best practices and K2 recommendations. While the recommendations cover both file system and ASM installations, please note that Kaminario strongly recommends the use of ASM in Oracle installations.

Best Practice	Oracle Recommendation	Kaminario Recommendation	Comments
Raid Configuration	Oracle recommends to stripe and mirror as much as possible, depends on the used storage.	K-RAID (used by default).	No RAID configuration is needed at the ASM or filesystem levels.
Separation of File Types	Allocate different file types on different LUNS for best spindle performance and administration (Redo Archive, Data).	Separation of different file types for different LUNS is unnecessary for best performance.	Separation of different database files is recommended when using K2 snapshots.
Number of LUNs	For ASM, four LUNs for each ASM Disk Group. For filesystems, according to storage subsystem.	One LUN is sufficient for filesystems. For ASM it is recommended to use at least four LUNS for each ASM Disk Group.	For Filesystems, there are certain cases where more than one LUN can improve performance depending on workload.
Number of Files	Several data files might increase performance. Depends on the workload.	No need for more than a single data file for best performance.	Redo Logs - It is recommended to have a large number (e.g. 8) of well-sized redo logs.
ASM Allocation Unit Size	Default	Default	There is no performance benefit to changing the Oracle ASM default allocation unit size.
ASM Redundancy	Normal	External	Each ASM disk group is protected by K2's array-based K-RAID.

For additional information, please refer to [Oracle documentation](#).

RAID CONFIGURATION

https://docs.oracle.com/cd/E11882_01/install.112/e47689/install_overview.htm#LADBI1064

- When deploying Oracle on Kaminario K2, there is no need to define any RAID configuration.
- Kaminario K2 uses its own RAID configuration, called K-RAID. K-RAID is a highly efficient scheme with an 87.5% utilization rate, achieved by deploying efficient erasure coding on each of its 24 SSDs.
- K-RAID is an optimized version of RAID6 protection that does not compromise on protection or performance.
- K-RAID consists of two logical RAID groups, each one with a parity (P1 and P2), and an additional parity for the two groups (Q), as show in Figure 1, below:

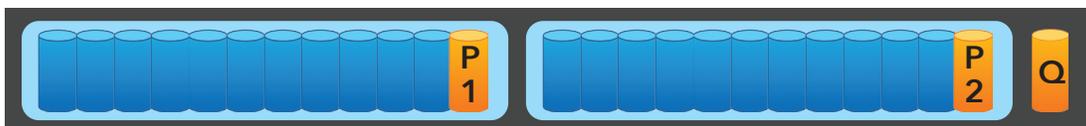


Figure 1: K-RAID logical layout of two RAID groups and a Q parity for both groups

SEPARATION OF FILE TYPES

- As Kaminario K2 is built from Solid State Drives (SSDs), there is no issue of moving parts and there is no performance influence if data is random or sequential.
- Kaminario K2 implements a global adaptive block size algorithm instead of using a fixed I/O block size. It allows the K2 to support different I/O sizes, all running concurrently.
- There is no real performance benefit in K2 to allocate multiple LUNs to store different database file types.
- If you plan to utilize K2 snapshots, it is recommended to use separate physical storage LUNs and volume groups for user's data, redo log and archive files.

NUMBER OF LUNS

- K2 can provide the necessary performance needed by Oracle with a single physical LUN.
- Different workloads might need more than one LUN in Oracle file system installations. There is no issue to define more than one LUN for Oracle.
- Operating system *I/O Scheduler* queues can become a bottleneck in high I/O workloads. Thus, it is recommended to have at least four equally-sized ASM disks (i.e. storage volumes). If higher-performance is needed, this value can be increased to eight ASM disks per diskgroup. Filesystems typically cannot achieve the same high levels of performance as ASM and therefore do not see these issues at the I/O Scheduler level.

NUMBER OF FILES

- No need for more than a single data file for best performance.
- The number of files should be subject to administration needs and environment limitations. K2 is agnostic to the number of files in the environment.
- Because of the high performance available from K2, redo logs often rotate faster than on traditional storage systems. For this reason, it is recommended to increase the size and number of redo logs in order to avoid the possibility that the database has to wait for a log to finish being archived before it can be used. Consider at least 8 redo log groups.
- There is no need to multiplex needed for redo log files as K2 K-RAID ensure the logs redundancy.

ASM ALLOCATION UNIT SIZE

- There is no performance benefit to changing the ASM default allocation unit size.

ASM REDUNDANCY

- K2 makes use of Kaminario's K-RAID protection to ensure that data remains available even during multiple component failures.
- Kaminario recommends using ASM's EXTERNAL redundancy option to avoid additional overhead of double or triple mirroring on top of K-RAID.

Additional Best Practices

There are additional best practices which are not entirely Oracle related and are recommended to implement for achieving best performance.

SCSI UDEV DEVICE RULES

I/O SCHEDULER

- It is recommended to set the Linux I/O Scheduler property to [noop] to avoid the waste of CPU on scheduling I/O.
- To configure the I/O scheduler to [noop] use the following udev rule.
 - Create/Add the rule file `"/etc/udev/rules.d/62-io-schedulers.rules"` with the following contents:

```
ACTION=="add|change", SUBSYSTEM=="block", ENV{ID_SERIAL}=="20024f400*", ATTR{queue/scheduler}="noop"
ACTION=="add|change", SUBSYSTEM=="block", ENV{DM_UUID}=="mpath-20024f400*",
ATTR{queue/scheduler}="noop"
```
 - Run the following udev command:

```
udevadm trigger && udevadm settle
```
 - Verify that the scheduler is configured to noop:

```
multipath -ll | grep KMNRIO | awk '{print "cat /sys/class/block/" $3 "/queue/scheduler"}' | bash
```

RQ_AFFINITY

- Subject to testing, changing the `rq_affinity` to 2 can benefit performance by targeting an I/O to the same CPU core that originated the process.
- To configure the I/O scheduler to [noop] use the following udev rule.
 - Create/Add the rule file `"/etc/udev/rules.d/62-io-schedulers.rules"` with the following contents:

```
ACTION=="add|change", SUBSYSTEM=="block", ENV{ID_SERIAL}=="20024f400*",
ATTR{queue/rq_affinity}="2
ACTION=="add|change", SUBSYSTEM=="block", ENV{DM_UUID}=="mpath-20024f400*",
ATTR{queue/rq_affinity}="2"
```
 - Run the following udev command:

```
udevadm trigger && udevadm settle
```
 - Verify that the `rq_affinity` is configured to 2:

```
multipath -ll | grep KMNRI0 | awk '{print "cat /sys/class/block/" $3 "/queue/rq_
affinity"}' | bash
```

ASMFD VS. ASMLIB

- ASMLib and ASMFD, both aim to simplify the management of Oracle files.
- Both ASMLib and ASMFD are compatible with K2.
- Kaminario recommends using ASMFD as it has TRIM (UNMAP) support.



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 and the Kaminario logo are registered trademarks of Kaminario, Inc. K-RAID, and Perpetual Array are trademarks of Kaminario, Inc.

Product specifications and performance are subject to change without notice.

The Kaminario ForeSight program is subject to terms and conditions.