

Achieving SQL Server High Availability with Kaminario K2

April 2018

TABLE OF CONTENTS

- 2** Executive Summary
- 2** Introduction to Kaminario K2
- 3** Microsoft SQL Server
- 3** The Need for High Availability in SQL Server
- 6** Kaminario High Availability
- 7** K2 Consideration for SQL Server Installations
- 8** K2 and SQL Server Always-On Availability Group
- 21** K2 Replication with SQL Server
- 24** Appendix A
- 28** About Kaminario

Executive Summary

The Kaminario K2 all-flash array is built to support modern workloads that require efficient use of resources, high performance and flexibility in scale. Virtualized environments, that are used to increase the utilization of servers and increase productivity, are an ideal workload for the K2's all flash architecture. One such environment is SQL Server deployed via VMware vSphere using ESXi. Kaminario has a unique value proposition for this implementation, and it is recommended that best practices and reference guides are followed in order to maximize capacity and performance efficiencies.

This whitepaper highlights the several options available to achieve high availability (HA) of SQL Server instances and databases. The Kaminario K2 has several built-in methodologies to ensure HA, each one is covered in detail in this whitepaper.

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 SQL Server on the Kaminario K2 with VMware ESXi environments.

Introduction to Kaminario K2

The K2 all-flash array is an appliance which is a combination of tested and packaged hardware, software and services. K2's Gen6 hardware platform is based on leading off-the-shelf enterprise components that enable K2's software-defined architecture and software stack. The K2 runs Kaminario VisionOS™, the next-generation flash operating system stack, that provides the core software and advanced data services framework.

VisionOS enables modular components and services that demonstrate a superior value proposition across a real scaleout storage platform, both in innovation and in ease of use:

- **DataShrink** – Data reduction features and capabilities are mandatory for economics of flash storage. With differentiating inline, global, adaptive and selective deduplication, together with inline byte aligned compression, thin provisioning and zero detection, Kaminario is able to establish itself a cost-efficiency leader of flash storage.
- **DataProtect** – Kaminario values its customers' data more than anything. Native array based snapshots and replication allow for returning to any point in time in any site. Data-at-rest AES256 encryption makes sure that data is kept private and safe at all times. A highly resilient design of no single point of failure, non-disruptive upgrades (NDU) and a robust RAID scheme facilitate 99.999% of data availability.
- **DataManage** – The K2 can be managed by various means. Internal management includes an intuitive web-based GUI, a scriptable CLI and a fully programmable RESTful API platform.
- **DataConnect** – K2's RESTful API allows for external applications of the IT eco-system to easily integrate and seamlessly manage the K2. This eco-system is constantly growing and includes: VMware vSphere, Microsoft VSS, OpenStack, Flocker (containers) and Cisco UCS director.

Microsoft SQL Server

Microsoft SQL Server is one of the most widely used databases in the world. Most organizations often run both production and revenue-generating applications on redundant instances of SQL Server. SQL Server is often preferred by developers for its simplicity and ease of use, enabling the use of a single database to perform tasks such as Online Transaction Processing workloads (OLTP), data mining, and Online Analytical Processing workloads (OLAP). One limiting factor about SQL Server is its host utilization - needing to utilize all of the available CPU, memory and disk. This often leads to either underutilized or a lack of resources based on host configuration. The growth of virtual environments has helped overcome this hurdle. Virtualizing SQL Server allows for the optimization of resources and the ability to add resources as needed and migrate to newer hardware without incurring expensive costs. Storage is another critical component in running SQL Server, with different workloads requiring different I/O patterns and levels. Supporting VMware hosts with scalable storage is just as important for achieving the desired performance from the SQL Server environment.

The Need for High Availability in SQL Server

Databases like Microsoft SQL Server often support both revenue-generating customer transactions and internal business intelligence activities that drive insights and innovations, thus playing a critical role in data processing and storage infrastructure for most organizations. In most cases SQL Server needs to be globally available around the clock, making the database servers and their supporting storage infrastructure the ideal components for implementing and using different high availability solutions. When considering database platforms, performance is often the most important factor. However, uptime and availability also need to be considered when designing and implementing solutions. SQL Server provides different high availability solutions to choose from, depending on environment requirements. The Kaminario K2 supports all the different SQL Server HA solutions.

LOG SHIPPING

Log shipping operates at the database level. Log shipping maintains one or more warm standby databases as a backup database for the primary production database. Log shipping automatically sends transaction log backups from a primary database on a primary server to a backup database on a backup server.

PROS

- Provides a disaster-recovery solution for a primary database at one or more backup databases in separate SQL Server instances.
- Allow read-only access on the backup databases.
- Restore time of the transaction log on the backup databases is configurable, allowing a safe time to detect user errors before they are reflected on the backup databases.

Cons

- Clients are disconnected from the backup databases during transaction log restores.
- Manual failover.
- High maintenance cost for administration in case there are multiple backup database servers or more than a single database.
- Log shipping owns the backup chain, any other log backup will cause chain issues.
- Additional licensing cost in different log shipping scenarios.

DATABASE MIRRORING

Database mirroring is a solution for increasing the availability of a SQL Server database. Database mirroring maintains two copies of a single database that must reside on different instances of SQL Server. One server instance serves the database to clients (the principal server), while the other acts as a hot/warm standby server (the mirror server). When the standby server is synchronized, it supports a rapid failover without data loss of committed transactions. If the server is not synchronized, the mirrored server acts as a warm standby server with possible data loss. Mirroring is implemented on a per-database basis and works only with databases that use the full recovery mode.

PROS

- Supports automatic and manual failover.
- Supports automatic page repair for protection against page corruption.
- Supports rolling upgrades for production environments.

Cons

- Automatic failover supported only in synchronous mirroring which is not always supported by the infrastructure.
- In synchronous mirroring, a transaction is committed when it is committed to the mirrored database as well.
- Database mirroring will be removed in future versions of Microsoft SQL Server (based on SQL Server 2016 documentation).

ALWAYS ON FAILOVER CLUSTER INSTANCES

Always-On Failover Cluster Instances leverages Windows Server Failover Clustering (WSFC) functionality. FCI provides local high availability through redundancy at the server-instance level. FCI is a single instance of SQL Server installed on multiple nodes (on the same failover cluster). Each node in the cluster has its virtual IP but it appears in the network as one IP of a single server. The cluster disks are shared disks and they should be in a SAN or SMB. When a single node, which acts as the active SQL Server at that time fails, another node at the cluster takes ownership of the resources and acts as the active SQL Server node.

PROS

- Protection at the instance level.
- Supports automatic and manual failover.
- No configuration is needed in client level during failover - one IP for all nodes.
- Configurable and flexible failover.
- Resources can be throttled during failover.

Cons

- High cost on passive nodes.
- High maintenance costs.
- Storage acts as single point of failure.

ALWAYS ON AVAILABILITY GROUP

Always-On Availability Group leverages Windows Server Failover Clustering (WSFC) functionality. It provides high-availability and disaster-recovery solutions at the enterprise level to database mirroring. Always-On Availability Group provides a failover environment to a specific set of databases known as availability databases that failover together. In a specific availability group, there is one read/write primary node and between one and eight secondary nodes. The secondary nodes can act as read only nodes and for backup operations. Failover to a replica happens when there are instance level issues and not a per database issue or a user issue such as deletion of a database, corruption of transaction log or loss of a datafile.

PROS

- Supports up to nine replicas (one primary and eight secondary's).
- Supports synchronous and asynchronous modes.
- Supports automatic and manual failover.
- Supports read-only and backup preferable secondary replicas.
- Supports one listener for client connection - one IP for all nodes, no configuration is needed in client level during failover.
- Supports automatic page repair for protection against page corruption.
- Supports encryption and compression.

Cons

- Additional licensing cost.
- User error on a primary database is reflected automatically to the secondary replicas.

Next table summarizes the pros and cons of the different SQL Server HA options:

Solution	Hardware Cost	Maintenance Cost	License Cost	Support Auto Failover
Log Shipping	Low	High	High	X
Mirroring	Low	Low	High	✓
Failover Cluster Instance	High	High	Low	✓
Always-On Availability Group	Low	Low	High	✓

Table 1: Pros and Cons of Different SQL Server HA Solutions

Kaminario High Availability

The Kaminario K2 is architected and built to meet the requirements of the most sensitive enterprise applications. The enterprise resiliency starts by deploying only enterprise grade hardware components and continues with High Availability (HA) throughout K2's design, scalability of fault domains and providing the right features for building an enterprise product.

SNAPSHOTS

Kaminario's patented snapshot architecture follows VisionOS's guidelines of storage efficiency, performance and scalability. Snapshots are created instantly, with no performance impact and they do not take up any capacity. Using Kaminario VSS provider, Kaminario can take an application-consistent snapshot of SQL Server databases with no performance impact and no need for additional capacity. Kaminario K2 snapshots can be mounted for read/write purposes, which serve to create additional working environments such as QA, Test&Dev, analytics, backup and more, all at a very low cost of storage capacity. Read/write snapshots deliver the same performance of the production volumes, without any impact on the production volumes. The duration of creating a snapshot has no dependencies on the number or size of the volumes being snapped or how big the array is. Using the snapshot's restore functionality for recovery purposes is done without losing any of the snapshot history and is allowed at any time. The snapshots can be accessed from any of the storage controllers of the K2, without bottlenecks or load balancing of affinity to a specific controller.

REPLICATION

Replication provides site resiliency that completes the enterprise datacenter resiliency requirements. Kaminario leverages its snapshot architecture to facilitate asynchronous snapshot-based replication between K2 AFAs. Since the replication is based on the snapshot architecture, there is no impact on the production environment's consistent high performance; deltas of the replicated copies are captured with no dependencies on the link speed. SQL Server snapshots using Kaminario K2 are application-consistent, replication of such a snapshot can be mounted in a remote site with minimum additional configuration and can be used as a production like environment. VisionOS storage efficiency features of deduplication and compression are also used to significantly reduce the amount of data sent between arrays.

K-RAID™

Aside from being highly efficient, the K-RAID is extremely robust. It can sustain two concurrent SSD failures and up to three SSD failures within each separate SSD shelf, without loss of data. As the K2 scales capacity, so does the number of system-wide SSD failures that the system can sustain. The K-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. It has minimal performance impact during the rebuild and no performance impact on the array's performance once the rebuild is completed. The K-RAID is fully automatic and does not require any configuration or human intervention, which means another IT task is offloaded to the Kaminario K2. With the use of K-RAID, there is no need to configure any special RAID for SQL Server environment, Kaminario K-RAID provides the best performance and best protection for any SQL Server need and workload.

POWER LOSS

A storage array with enterprise capabilities must have the ability 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 was already acknowledged by the storage controllers (K-Nodes) before being stored on the K RAID are saved in two distinct K-Nodes for redundancy. Each K-Node in the K2 is equipped with non-volatile memory that will allow the controller to de-stage any inflight data that has not been stored to its K-RAID. Data that is already stored to the K-RAID is kept persistent and is sustainable through power cycles. In case of a power loss any SQL Server committed transaction is safe within the K2.

NO SINGLE POINT OF FAILURE (SPOF)

VisionOS supports a double-everything approach for Kaminario's hardware components and data. All data and metadata at rest is protected by the dual-parity K-RAID. However, the K2 does not have passive or idle components in the array; all of its resources are being utilized at all times. There is full redundancy of every component in the system and there is not a single component that can fail and cause unplanned down time or data loss. Each K-Block is a standalone failure domain, which means that the entire array can sustain more failures as the array scales and the MTBF of the array stays constant.

More information regarding K2 HA and general architecture, can be found at *K2 All-Flash Array Architecture White Paper* - http://info.kaminario.com/k2_architecture

K2 Consideration for SQL Server Installations

A deployment or installation of SQL Server environment on the K2 is an easy and fast task with almost no special configurations or implications to consider. Having said that, there are few aspects to take into consideration when deploying SQL Server on K2:

VOLUME GROUPS AND VOLUMES

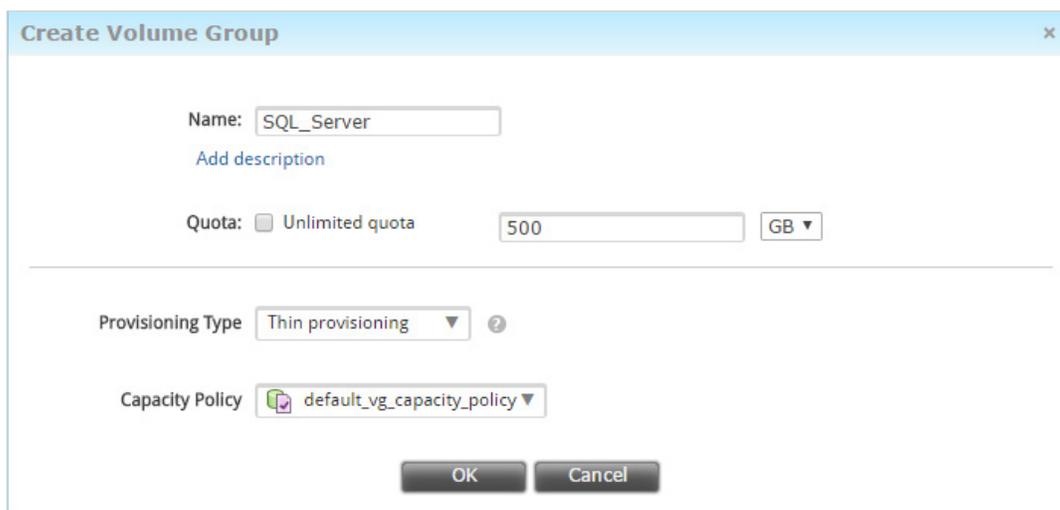
In the K2, a volume group acts as a consistency group when taking a snapshot. A volume group should contain the volumes that support the SQL Server environment. Using the Kaminario VSS Provider and Microsoft Diskshadow, it is possible to take a consistent-application snapshot of a SQL Server database. If it is planned to utilize K2 snapshots and VSS, it is recommended to use a separate physical storage LUN for the user's data and log files and for the TempDB and system databases files. LUN is the smallest granularity of a snapshot, a single K2 snapshot should include only user data and not TempDB or system data. Separation of two or more different databases files to different physical storage LUNs is recommended in case it is planned to use K2 snapshots and VSS. Having more than a single database in a physical LUN will lead to multiple databases in a single snapshot. The restore operation on the LUN, restores all databases in the LUN, an operation that is not always necessary.

DEPLOYING SQL SERVER IN A VMWARE ENVIRONMENT

If you plan to deploy your SQL Server environment on a VMware VM, it is recommended to decide in advanced if the VM disks will be configured as VMDK files, as part of a bigger datastore, or as RDM disks. It is possible to use any of SQL Server HA with either RDM or VMFS disks (with additional configurations needed when working with virtualized environments). In terms of performance there is no performance difference when using either RDM or VMFS disks, Kaminario provides the best performance with either of the options. In case of utilizing Kaminario snapshots, in order to take consistent-application snapshots, the database user files should reside on RDM disks and not on VMFS disks.

DEDUPLICATION

When creating a volume group in Kaminario, the default Provisioning Type chosen by the K2 is Thin provisioning with dedupe. It is recommended to change the provisioning type to Thin provisioning only. SQL Server database might benefit from deduplication in terms of capacity, however in most cases the performance penalty of having a SQL Server instance on a deduplication LUN overcomes the capacity benefit and we suggest having a SQL Server instance on a LUN without deduplication. SQL Server will benefit from the K2's inline compression in any case.



The screenshot shows a 'Create Volume Group' dialog box with the following fields and options:

- Name:** SQL_Server
- Quota:** Unlimited quota, GB
- Provisioning Type:** Thin provisioning
- Capacity Policy:** default_vg_capacity_policy

Buttons: OK, Cancel

Figure 1: Create Volume Group - Provisioning Type

K2 and SQL Server Always-On Availability Group

Another, often used HA option with SQL Server is Always-On Availability Groups. Always-On Availability Groups provide continuous database availability via a failover environment using a discrete set of databases that fail-over together. The following tests demonstrates how the K2 supports the use of Always-On Availability Groups and how it is possible to consolidate the work of more than one SQL Server environment on a single K2 all-flash array. To demonstrate the environment a three-phase test was conducted:

1. Standalone SQL Server 2016 Test - Before demonstrating the work of a SQL Server Always-On on K2 environment, we would like to demonstrate the workload used for the tests on a standalone SQL Server environment working with Kaminario K2 as the storage.
2. Always-On Availability Group with SQL Server Failure Test - Running the same workload as in the first test on an Always-On Availability Group environment while failing the primary node.
3. Always-On Availability Group with K2 components Failure Test - Running the same workload as in the first test on an Always-On Availability Group environment while failing different K2 components.

STANDALONE SQL SERVER 2016 TEST

HammerDB (www.hammerdb.com) a known database load testing and benchmarking tool for Linux and Windows, was used to create an OLTP workload on a standalone SQL Server 2016 environment. A database of 1600 warehouses was created while running a Timed Test Driver Script of 5 minutes with 1000000 transaction and 160 users. The test presents the workload created on the SQL Server environment and on the K2 array. This workload is the baseline workload for the following HA tests.

As seen in Figure 2 below, the workload produced in HammerDB was an average of 1.4M TPM.

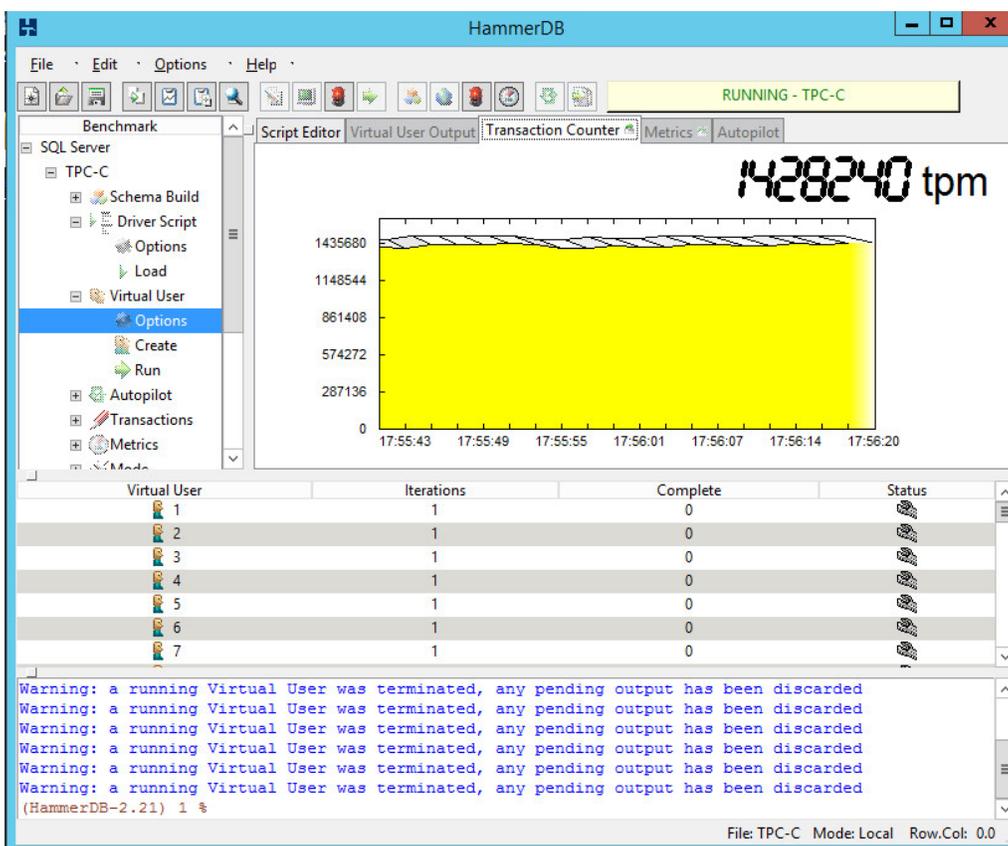


Figure 2: HammerDB TPM in a Standalone SQL Server Test

The storage workload seen on the K2 during through the entire test, averaged at 230 MB/s throughput, 29,000 average IOPS and latency below 0.5 ms. As shown in Figure 3 below.

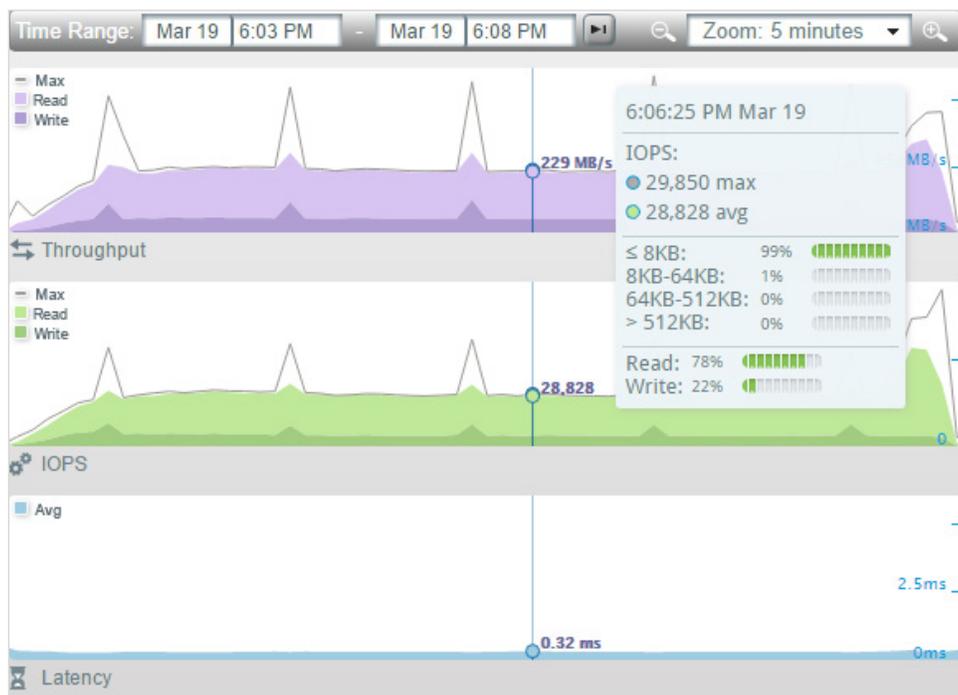


Figure 3: K2 Performance on a Standalone SQL Server

ALWAYS-ON AVAILABILITY GROUP WITH SQL SERVER FAILURES TEST

In this test, the HammerDB workload was used as in the first test. The workload ran on a SQL Server Always-On environment of three nodes. Two out of the three nodes were connected to a dual K-Block system while the third node was connected to another, single K-Block system. The primary node called sol-always08, together with a secondary synchronized node sol-always09 were connected to a dual K-Block system. The asynchronous node is sol-always10 which is connected to a different single K-Block system. Figure 4 below shows the Always-On configuration from the SQL Server Management Studio.

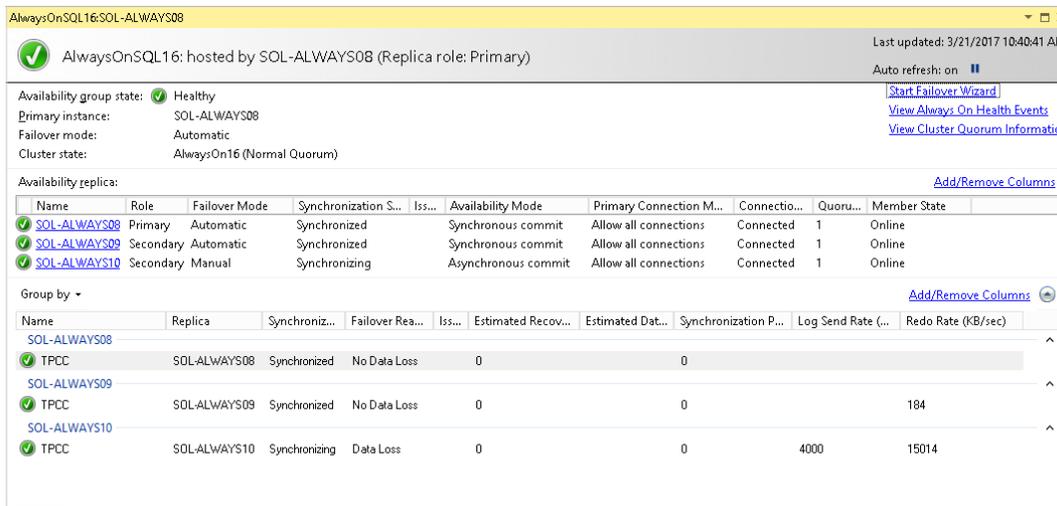


Figure 4: SQL Server Availability Group Dashboard

The test demonstrated the performance of the SQL Server Always-On environment during a SQL Server failover. A constant workload stream of read/write database I/O was routed using HammerDB through the Availability Group listener to the primary SQL Server database. A network failure was simulated on the primary database (sol-always08) to trigger a failover of the primary node to the secondary synchronized node (sol-always09). The workload on all SQL Server instances stopped during the failover process until the secondary node (sol-always09) took over and became the primary node. After a few minutes, sol-always08, the original primary node was reconnected to the network, and then became secondary node, resynching the data it lost when it was not connected. The following images show the entire process from an SQL Server perspective and K2 perspective.

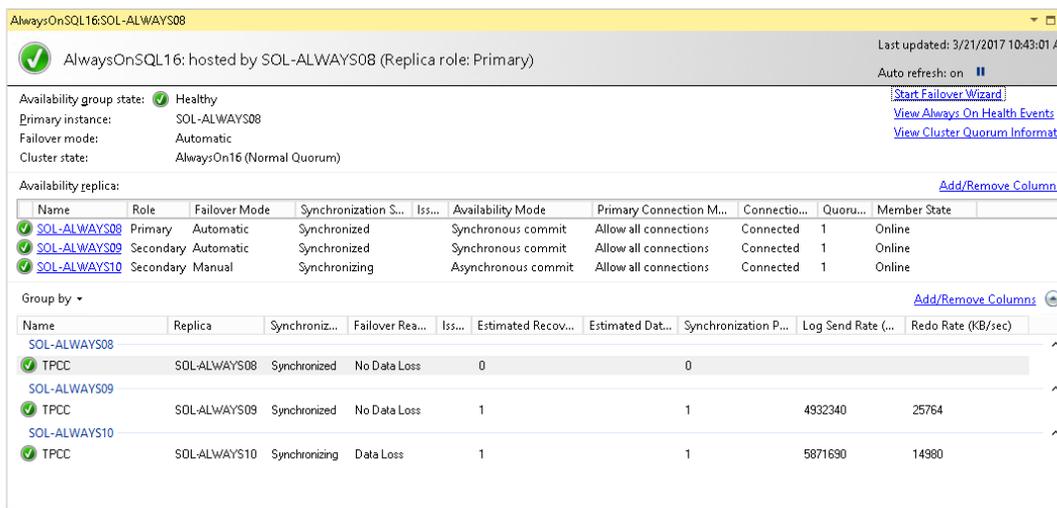


Figure 5: Always-On Dashboard - Workload Started

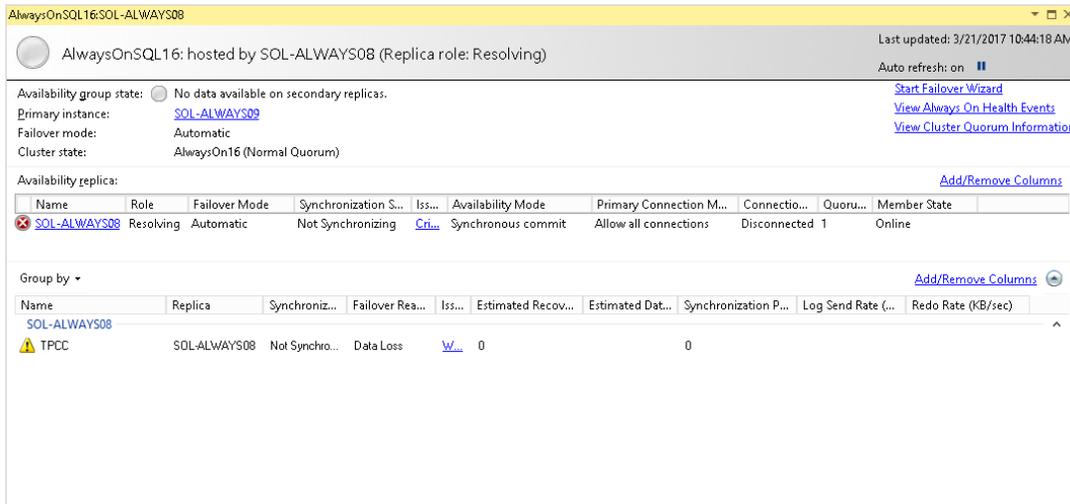


Figure 6: Primary Node Lost Connection



Figure 7: Secondary Node Becomes Primary

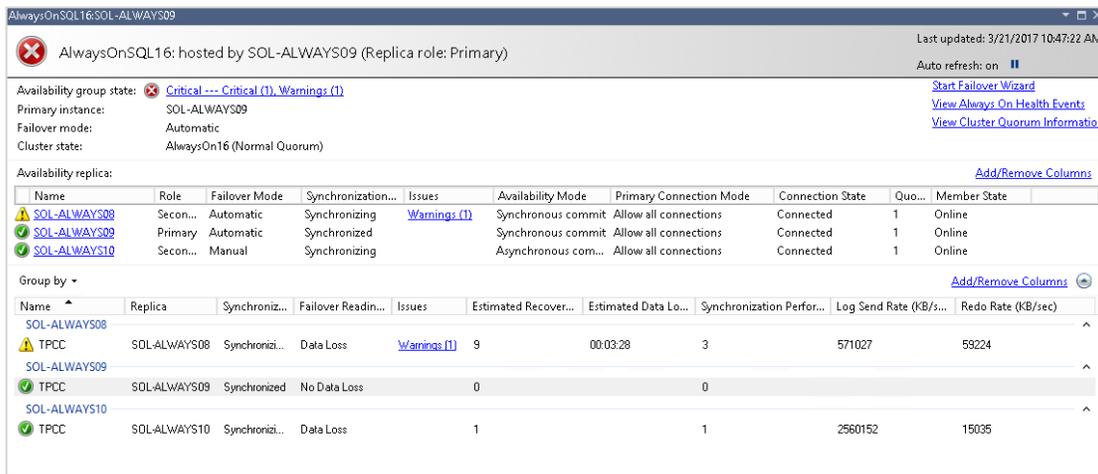


Figure 8: Former Primary Synchronizing as Secondary after Reconnection

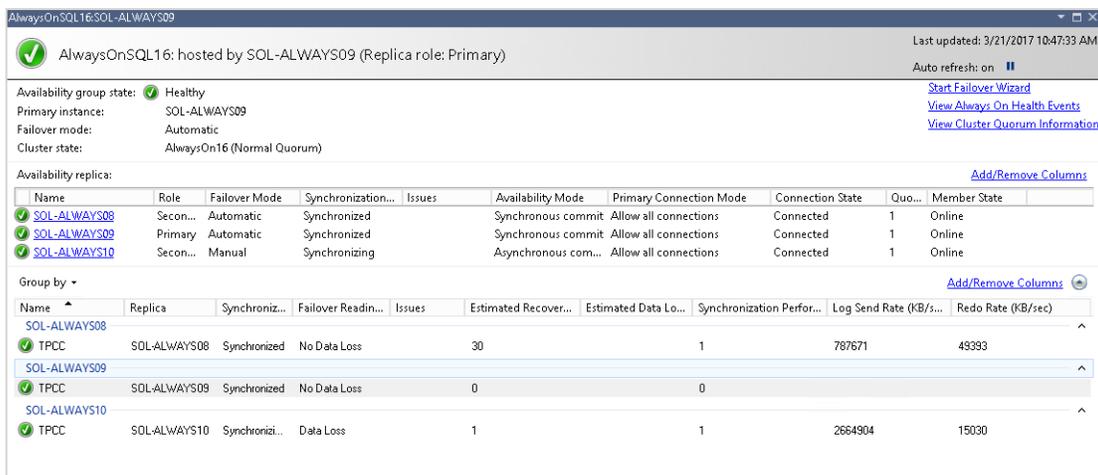


Figure 9: Secondary Synchronization Finished



Figure 10: Dual K-Block Performance during Always-On Workload



Figure 11: Single K-Block Performance during Always-On Workload

ALWAYS-ON AVAILABILITY GROUP WITH K2 COMPONENTS FAILURES TEST

In Order to be able to serve I/O at all times, Kaminario developed its own OS called VisionOS, operating the K-Nodes. VisionOS supports a double-everything approach for all hardware and data. A component might fail on the K2, but the K2 can still maintain its high availability and high performance with no degradation to the workloads. The following tests represent how an Always-On environment keeps working with no performance degradation to the SQL Server when simulating a HW failure on the K2. Two tests were performed simulating HW failures on both K2s of the Always-On environment. During the tests, simulation of faulty components were run by removing the HW component from the K2 while running a HammerDB workload. The two tests were:

1. Simulation of a PSU (Power Supply Unit) Failure - The failure simulation was done by removing a PSU from a K-Node of the dual K-Block array, serving two nodes of the Always-On environment.
2. Simulation of an SSD Failure - The simulation was done by removing an SSD from the shelf of the single K-Block array, serving the asynchronous node of the Always-On environment.

The following set of figures represents the failures on the K2 GUI and how it keeps delivering consistent performance while the components are missing from the array.

SIMULATION OF A PSU FAILURE

In the dual K-Block system a PSU was removed from K-Node01 while the K2 was servicing the HammerDB workload on the Always-On environment. In Figure 12 below the Inventory tab shows that PSU01 was phased out from K-Node01 (simulating the removal of the PSU).

The screenshot shows the 'Inventory' tab in the K2 GUI. At the top, there are navigation tabs: 'General', 'External Configuration', 'Inventory' (selected), and 'Capacity Policies'. A yellow warning banner with an exclamation mark icon states 'There are unhealthy components.' with a 'Show' button. Below this is a table with three columns: 'Name', 'Type', and 'Health Status/Message'.

Name	Type	Health Status/Message
▼ kblock01	K-Block	Healthy. Contains unhealthy components
▼ knode01	K-Node	Healthy. Contains unhealthy components
battery01	Battery	Healthy
battery02	Battery	Healthy
psu01	PSU	Phased out, sensor 'PS1 Status' shows status 'ok'
psu02	PSU	sensor 'PS2 Status' shows status 'ok'
storage01	K-Node local storage	Healthy
storage02	K-Node local storage	Healthy
▶ knode02	K-Node	Healthy
▶ shelf01	Shelf	K-RAID. Healthy
▼ kblock02	K-Block	Healthy
▶ knode01	K-Node	Healthy
▶ knode02	K-Node	Healthy
▶ shelf01	Shelf	K-RAID. Healthy
switch01	Switch	Healthy
switch02	Switch	Healthy

Figure 12: PSU Failure Indication on the Inventory Tab in the K2 GUI

The failure indication is indicated on the K2 Dashboard as in Figure 13 below.

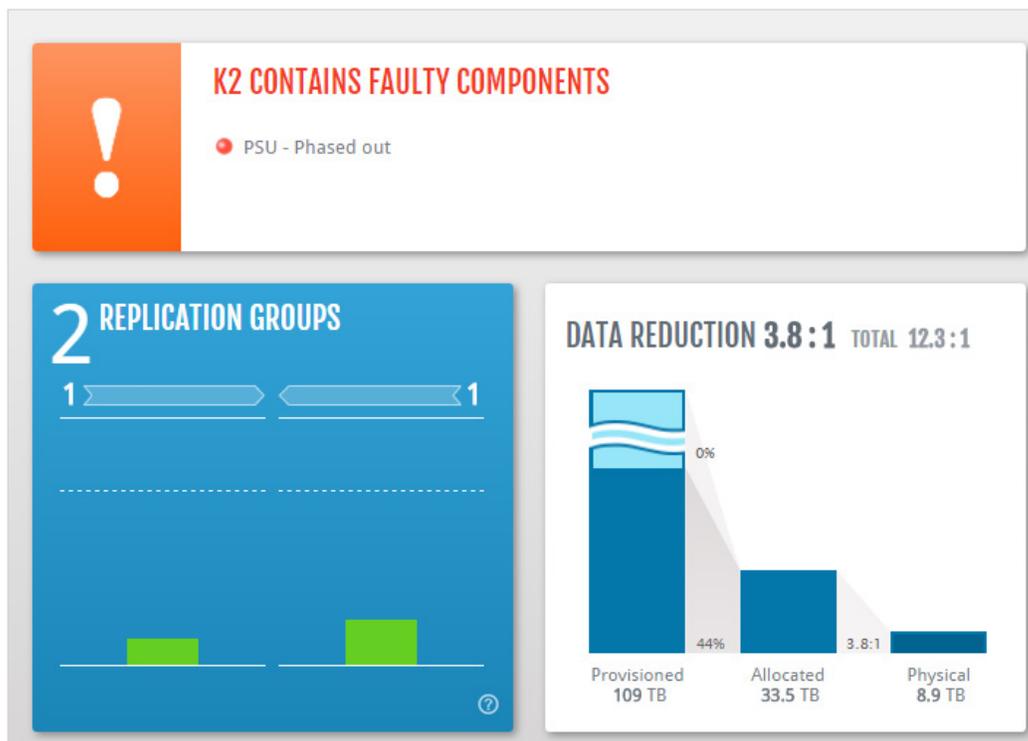


Figure 13: PSU Failure Indication on the Dashboard Tab in the K2 GUI

Figure 14 below shows the PSU failure on the Kaminario K2 Events tab.

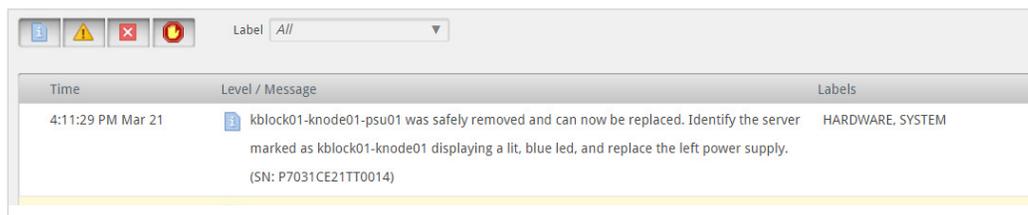


Figure 14: PSU Failure Indication on the Events Tab in the K2 GUI

Pay attention that the failure event happened at 4:11:29 PM. Figure 15 shows the performance of the K2 during the Failure. It is seen that the performance stays consistent and there is no degradation. The Figures shows a range of time between 4:10 PM and 4:14 PM.

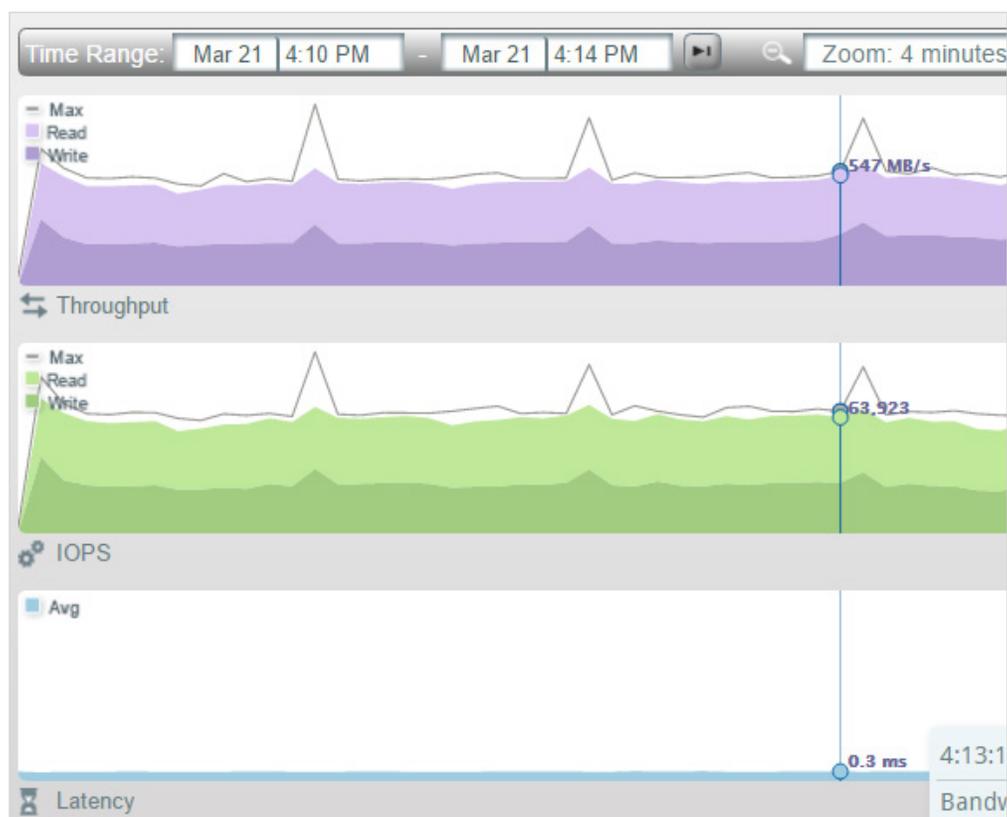


Figure 15: K2 Performance During PSU Failure

SIMULATION OF AN SSD FAILURE

In the single K-Block system an SSD was removed from the K-Block shelf while the K2 was servicing the HammerDB workload of the asynchronous node of the Always-On environment (sol-always10). Figure 16 below represents the Inventory Tab and the indication that storage02 SSD was phased out (simulating the removal of the SSD).

There are unhealthy components. [Show](#)

Name	Type	Health Status/Message
▼ kblock01	K-Block	🟡 Healthy. Contains unhealthy components
▶ knode01	K-Node	🟢 Healthy
▶ knode02	K-Node	🟢 Healthy
▼ shelf01	Shelf	🟡 K-RAID. RAID rebuild progress: 4%
expander01	Expander	🟢 Healthy
expander02	Expander	🟢 Healthy
psu01	PSU	🟢 Healthy
psu02	PSU	🟢 Healthy
storage01	SSD	🟢 Healthy
storage02	SSD	🔴 Phased out
storage03	SSD	🟢 Healthy
storage04	SSD	🟢 Healthy
storage05	SSD	🟢 Healthy
storage06	SSD	🟢 Healthy
storage07	SSD	🟢 Healthy
storage08	SSD	🟢 Healthy

Figure 16: SSD Failure Indication on the Inventory Tab in the K2 GUI

Every failure or procedure on the K2 is reflected on the K2 Dashboard. Figure 17 below shows the failure indication on the **Dashboard** tab.

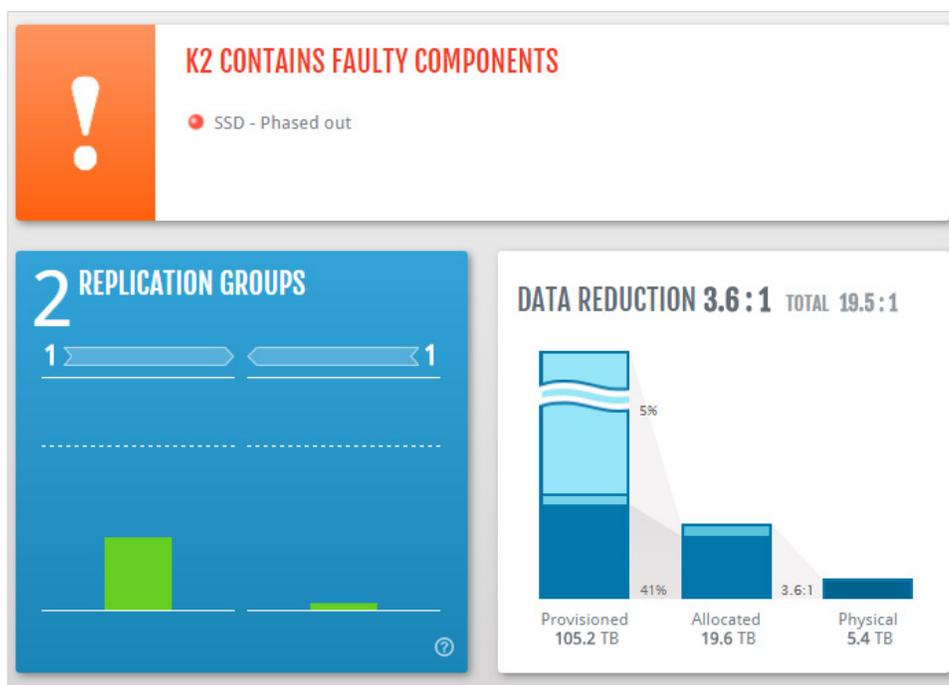


Figure 17: SSD Failure Indication on the Dashboard Tab in the K2 GUI

An event is immediately shown on the Events tab to indicate there is a problem. Figure 18 shows the event on the **Events** tab.

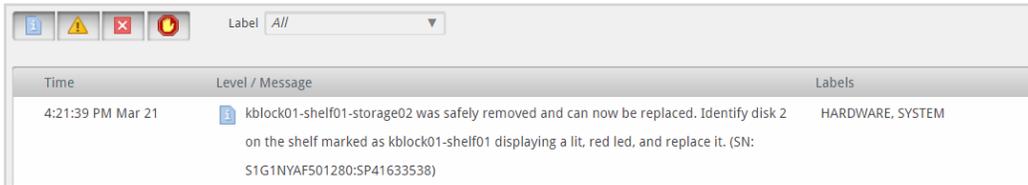


Figure 18: SSD Failure Indication on the Events Tab in the K2 GUI

As shown in the PSU failure, the performance of the K2 didn't change and stayed constant. There was no degradation and no issues were indicated from the Always-On node. Figure 19 presents the constant performance. Pay attention that the SSD was removed from the system at 4:21:39 PM.

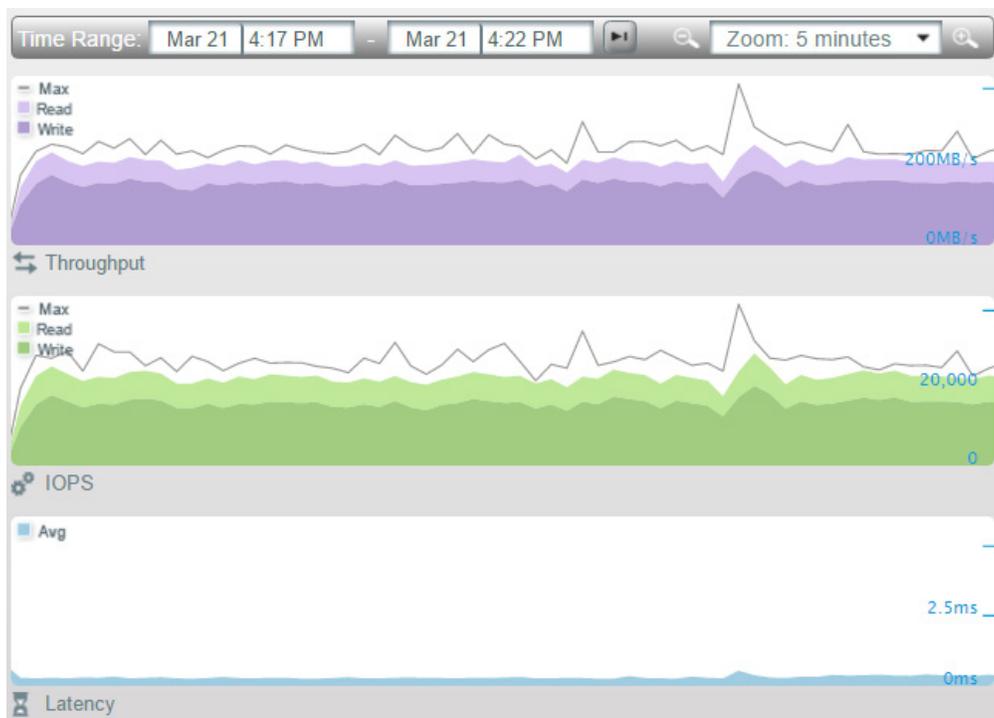


Figure 19: K2 Performance During an SSD Failure

As seen from the tests above, VisionOS with the double-everything approach makes sure that there is no impact on performance when an issue occurs on the K2 array and it is able to serve constant I/O at all times.

K2 Replication with SQL Server

The HA picture is not complete without talking about Kaminario's array-based replication. The Kaminario K2 implements an asynchronous replication, based on its mature native snapshot technology. By using asynchronous snapshot-based replication, the K2 can replicate data over long distances with minimal impact on performance, retain greater control over application consistency, and reduce the network overhead of streamed replication protocols. K2 replication can be utilized for several uses:

1. Utilization of the K2 replication as part of a greater DR solution.
2. Additional copy of the production database on a different K2 system.

Having a copy of the production database on a different K2 can benefit an organization in numerous ways. The replicated database can be used to create additional working environments such as QA, Test&Dev, analytics, offload backup and more, all at a very low cost of storage capacity.

The following tests show the utilization of the K2 native replication for establishment of an additional working environment on a second K2. This test uses the same environment as in the preceding tests; except for one change, ; the replication session ran from the single K-Block system to the dual K-Block system. A VMware virtualized VM established on the single K-Block was being replicated to a waiting VMware VM on the dual K-Block array. The source SQL Server instance was a 2016 SQL Server installation with the same HammerDB database of 1600 warehouses as used in previous tests. The database was being replicated to the destination K2 which was connected to an ESX with a waiting SQL Server 2016 VMware VM. The waiting VM accepts the new disks and starts the SQL Server engine with a view created from the snapshot of the HammerDB database. In order to accomplish the process the steps below should be followed:

1. Take an application-consistent snapshot on the source K2 using VSS and Diskshadow.
2. Validate that the snapshot was replicated and arrived to the destination K2.
3. Create a new view out of the replicated snapshot and map it to the waiting VM's ESXi.
4. Rescan the ESXi datastores so the ESXi recognizes the new LUNs.
5. Add the new LUNs as RDM disks to the waiting SQL Server VM.
6. Modify and adjust the new disks on the waiting SQL Server VM.
7. Start the SQL Server database on the waiting SQL Server VM.

SETUP AND CONFIGURATION

IMPLEMENTING SQL SERVER REPLICATION WITH K2 REPLICATION

Take an application-consistent snapshot on the source K2 using VSS and Diskshadow

As K2 replication is based on Kaminario’s mature snapshots solution, the first step is taking an application-consistent snapshot. In order to have an application-consistent snapshot the snapshot should be triggered from the SQL Server host using VSS and diskshadow. For more detailed instructions on how to take an application-consistent snapshot of SQL Server database please refer to the **SQL Server Backup and Restore using Kaminario K2 Snapshots** whitepaper. Make sure that during the installation of the Kaminario VSS Provider, the REPLICATION_MODE parameter is set to true. Figure 20 below shows running the VSS script in diskshadow.

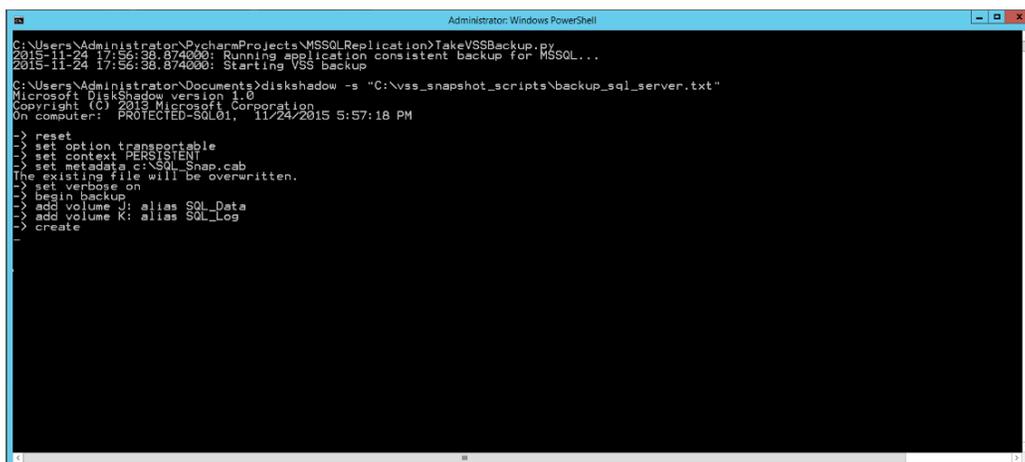


Figure 20: Taking VSS Snapshot

Validate that the snapshot was replicated and arrived at the destination K2

To validate that the snapshot has arrived at the destination K2, a ✓ sign should appear in the Has Peer column in the Snapshot tab in the K2 GUI as shown in Figure 21 below.

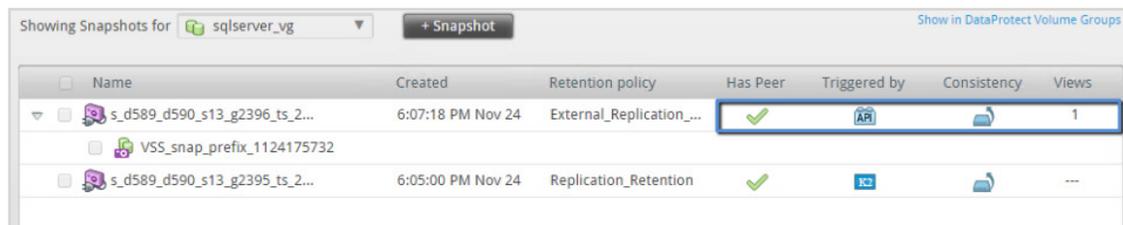
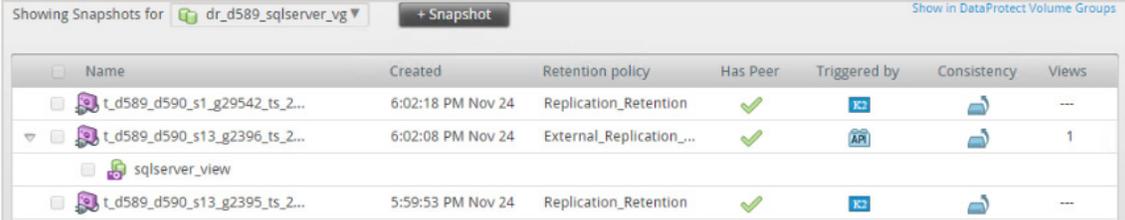


Figure 21: K2 GUI - Snapshots Tab with Replication Validation

Create a new view out of the replicated snapshot and map it to the waiting VMs ESXi

On the K2, a snapshot can't be mapped to a host, in order to map a snapshot to a host a view should be created. Instructions on the different options of creating a view can be found in the K2 User Guide. During the creation of the view, it is possible to map the view to the waiting VM ESXi host. Figure 22 below shows the replicated view on the destination K2.



Name	Created	Retention policy	Has Peer	Triggered by	Consistency	Views
t_d589_d590_s1_g29542_ts_2...	6:02:18 PM Nov 24	Replication_Retention	✓	IC	☰	---
t_d589_d590_s13_g2396_ts_2...	6:02:08 PM Nov 24	External_Replication_...	✓	API	☰	1
<div style="margin-left: 20px;"> sqlserver_view </div>						
t_d589_d590_s13_g2395_ts_2...	5:59:53 PM Nov 24	Replication_Retention	✓	IC	☰	---

Figure 22: K2 GUI - View Creation on the Destination K2

PAY ATTENTION: The next two steps in the document need to be followed when using SQL Server on a VMware VM environment. They are unnecessary in case of physical SQL Server hosts.

Rescan the ESXi datastores

As the tests in this document were done in a VMware environment, there is a need to rescan the ESXi datastores to reflect the new LUNs in view of the ESXi host. Performing a rescan of the storage on an ESX/ESXi host KB of VMware explains different ways to perform this step.

Add the new LUNs as RDM disks to the waiting SQL Server VM

The new LUNs are known to the ESXi host and can be added to the waiting VM as RDM disks. Add an RDM Disk to a Virtual Machine provides instructions on how to do so.

Modify and adjust the new disks on the waiting SQL Server VM

The disks will appear in offline state at first, read only without drive letters. Admins will need to bring the disks online and assign drive letters as appropriate for SQL Server instances. This can be done using Microsoft DiskPart tool.

Start the SQL Server database on the waiting SQL Server VM

After completing all the steps above the SQL Server VM is ready and it is safe to start the instance.

The entire procedure described in the test above can be automated and accomplished in less than a minute. For more information please contact Kaminario Support.

Appendix A

CONFIGURATION DETAILS FOR HARDWARE CONFIGURATIONS – ARRAYS, SERVERS

The environment described in the document consists of two different K2 arrays. A dual K-Block K2 array and a single K-Block K2 array. A K-Block is the building block of the K2 AFA, each K-Block consists of two active/active controllers with one or more SSD shelves. In this test environment, one shelf was used per K-Block in both K2s.

An out-of-box configuration was used for the K2 array and no tuning was performed for any of the specific tests.

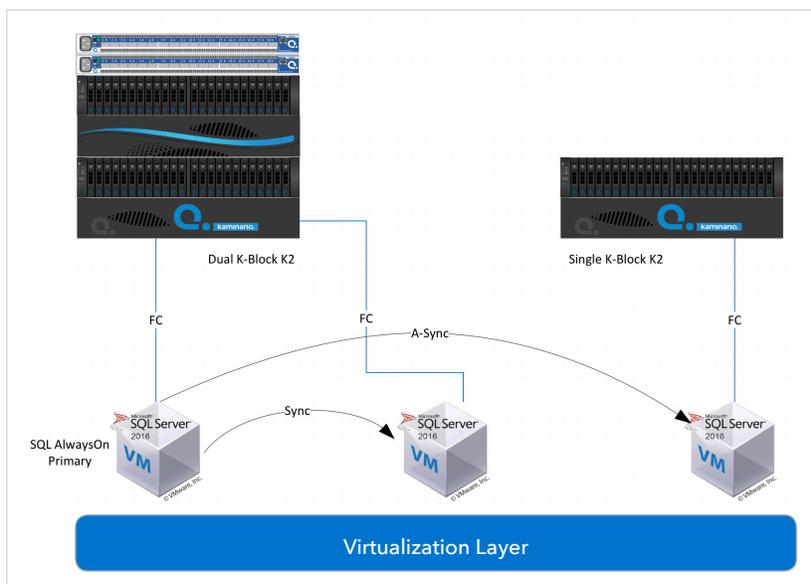


Diagram 1: K2 GUI - SQL Server 2016 Always-On Environment

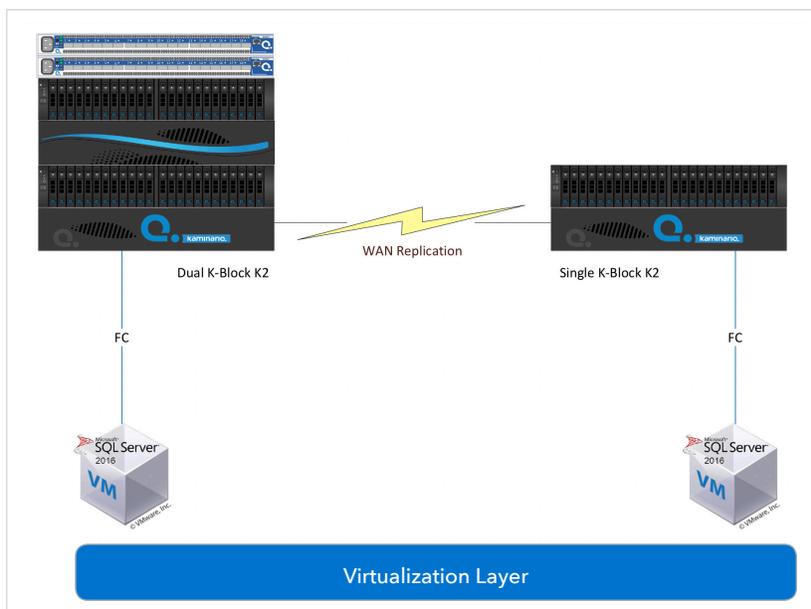


Diagram 2: SQL Server 2016 Replication Environment

SQL Server 2016 Virtual Machines - Always-On

Attribute	Specification
OS	Windows Server 2012 R2 Standard
vCPU	12
vMemory	8,192 MB
Virtual SCSI Controller 0 (OS)	LSI Logical Parallel
Virtual Disk (OS)- VMDK	40 GB
Virtual SCSI Controller 1	VMware Paravirtual
Virtual Disk (Data) - VMDK	350 GB
Virtual Disk (Log) - VMDK	50 GB
Virtual Disk (Temp) - VMDK	150 GB
Virtual Disk (TempLog) - VMDK	50 GB
Virtual Disk (Backup) - VMDK	150 GB
Virtual CD/DVD/Floppy Drives	Removed
Installed Applications	Microsoft SQL Server 2016 (RTM) (X64)

SQL Server 2016 Virtual Machines - Replication	
Attribute	Specification
OS	Windows Server 2012 R2 Standard
vCPU	12
vMemory	8,192 MB
Virtual SCSI Controller 0 (OS)	LSI Logical Parallel
Virtual Disk (OS)- VMDK	40 GB
Virtual SCSI Controller 1	VMware Paravirtual
Virtual Disk (Data) - RDM	350 GB
Virtual Disk (Log) - RDM	50 GB
Virtual Disk (Temp) - RDM	150 GB
Virtual Disk (TempLog) - RDM	50 GB
Virtual Disk (Backup) - RDM	150 GB
Virtual CD/DVD/Floppy Drives	Removed
Installed Applications	Microsoft SQL Server 2016 (RTM) (X64)

MICROSOFT SQL SERVER 2016 SOFTWARE CONFIGURATION:

1. Maximum server memory: 1500 MB
2. Startup parameters and trace flags:
 - E - Increases number of extents allocated for each file in a filegroup.
 - T610 - Controls minimally logged inserts into indexed tables.

HammerDB Virtual Machine	
Attribute	Specification
OS	Windows Server 2012 R2 Standard
vCPU	4
vMemory	4,096 MB
Virtual SCSI Controller 0 (OS)	LSI Logical Parallel
Virtual Disk (OS)– VMDK	40 GB
Virtual CD/DVD/Floppy Drives	Removed
Installed Applications	HammerDB v2.21

ESXI CONFIGURATION

- Supermicro 16 X 2.599 GHz Intel® Xeon® E5-2640 CPUs for a total of 32 hyperthreaded cores.
- 256 GB of memory.
- Dual Port 8 GB QLogic FC HBA used to connect to a dual K-Block K2 array through two redundant 48 Ports Brocade FC Switches.
- Dual Port 8 GB QLogic FC HBA used to connect to the single K-Block K2 array through two redundant 48 Ports Brocade FC Switches.



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.