

# 100TB Data Warehouse Fast Track Reference Architecture for Microsoft SQLServer 2017 Using Emulex Gen 6 Fibre Channel HBAs on HPE DL380 Gen10 and Kaminario K2

October 2018

---

## TABLE OF CONTENTS

- 2** Executive Summary
- 3** Solution Components
- 7** Broadcom Data Warehouse Fast Track reference architectures for SQLServer 2017
- 5** Recommended Reference Architecture
- 8** Windows Server 2016 Configuration
- 9** Dwft for Sqlserver 2017 Certification
- 10** Summary
- 10** References
- 11** About Kaminario

## Executive Summary

This Broadcom®, Kaminario, and Microsoft® SQL Server Data Warehouse Fast Track (DWFT) Reference Architecture provides guidelines and principles to assist customers in designing and implementing a balanced configuration for Microsoft SQL Server® data warehouse workloads to achieve out-of-the-box high performance. These database reference architectures enable each of the components in the database stack to provide optimal throughput to match the database capabilities of the specific setup. This paper describes the designs, configuration and guidelines used to achieve an optimally balanced 100TB DWFT reference architecture for SQL Server 2017 using Emulex Gen 6 HBAs by Broadcom, a Kaminario K2 all-flash array and HPE DL380 Gen10 servers providing at least 12,800MB/s of throughput.

The target audience for this reference architecture includes database administrators, business intelligence architects, storage administrators, IT directors, and data warehousing users seeking sizing and design guidance for business intelligence solutions with SQL Server 2017.

### GOALS AND OBJECTIVES

**Broadcom and Kaminario undertook this testing to provide a reference architecture capable of:**

- Achieving rock solid, scalable, sustainable performance for a demanding data warehouse environment
- Targeting moderate to large size data warehouses of 100+TB
- Showing mixed workload performance for those customers who would like to leverage a common storage array for both OLTP and OLAP workloads

## Solution Components

### Microsoft SQLServer 2017: Data with Improved Column-Store Technology

Microsoft SQL Server 2017 has made significant improvements in data warehousing technologies and performance, including column-store features as well as many other improvements. Column-store indices offer great advantages over traditional row stores for analytics and data warehousing queries. They are ideally suited for the star schemas, and tables with billions of rows which have now become commonplace. Among their advantages for analytics are:

#### Up To 10X Compression in Data Size:

Data warehouses are very large by nature and the compression offered by column-store index technologies delivers both space and cost savings as well as significantly increased performance. These benefits are possible due to the dramatically reduced I/O requirements given by the compression and coupled by the ability to only scan the specific columns required by each query. This compression also reduces the amount of memory required to hold a given number of rows from the source data warehouse.

#### Additional Indices:

SQL Server 2017 also provides the capability to add (B-Tree) indices to column store-based tables, which enables efficient single-row lookup. In addition to these architectural features, Microsoft has further optimized the processing of queries in column-store indices in the following ways:

- **Operator Pushdown:** Pushdown refers to moving both filter and aggregation query operations closer to the data, so that many of the filters and calculations can be done in the scan operators, dramatically reducing the volume of data that needs to be handled further on in query processing.
- **Batch-Mode Processing:** SQL Server 2017 includes enhancements in batch-mode processing that handles many rows at a time rather than serially doing calculations on each individual row. These batch operations are further optimized by leveraging Single Instruction Multiple Data (SIMD) vector processing CPU instructions in the Intel® architectures.

## Emulex Gen 6 HBAs

An essential element to the performance of applications and all-flash storage is the selection of the Fibre Channel Host Bus Adapter (HBA). Because HBAs provide the connectivity for data transfer between devices in the SAN, from host servers to switches and flash storage, their performance and reliability is pivotal to the performance of the entire SAN.

The HBAs key role in talking to both servers and storage, means that HBA innovation must be in lock-step with innovation in both storage and server design. On the server side, multi-core processors, new memory technologies, higher VM loads and the latest PCIe3.0 bus have placed increased demands on the network. On the storage side, new all-flash arrays such as the Kaminario K2 have solved storage performance bottlenecks. Installing the latest Emulex Gen 6 HBAs in both the K2 and in the servers it's connected to, enables the network with the highest horsepower, industry-leading reliability and the latest innovations required to meet the needs of the most demanding all-flash environments. Emulex Gen 6 HBAs provide a number of features:



- **Optimizes Flash & Application Performance** - 2.5x more IOPS in active-standby mode than others thanks to Emulex's exclusive Dynamic Multicore Architecture which can dynamically apply all ASIC resources to any port that needs it - either to a single active port or across both active ports in a dual-port configuration, as demanded by the workload. This ensures that performance is delivered when and where needed, to meet service level agreements (SLAs). Real workload testing shows that Emulex HBAs significantly boost application performance, delivering up to 46% greater OLTP transactions per minute (TPC-C order entry dataset) simply by upgrading to an Emulex Gen 6 HBA.
- **Alleviates SAN Congestion** - The exclusive Emulex congestion management feature throttles outstanding IO via queue depth assignment to avoid over subscription (slow drain). Emulex's queue depth changes require no reboots. The BB Credit recovery feature enables lost credit recovery via periodic synch-up.
- **Improves Security** - The exclusive rogue firmware protection feature uses digital signature verification prior to HBA firmware download to prevent security breaches caused by malicious firmware.
- **Future-proofs Data Center Investments** - Emulex leads the industry with the first shipping NVMe over Fibre Channel solutions on multiple Operating Systems. Concurrent support for NVMe over Fibre Channel and traditional SCSI over Fibre Channel, means there is no rip-and-replace required for existing 16 & 32GFC networks which provides datacenters with an easy transition path to NVMe all-flash storage.
- **Simplifies SAN Management** - Emulex supports Brocade Fabric Vision to maximize uptime, simplify management and provide unprecedented visibility.

### Kaminario K2 All-Flash Array

Flash and all-flash arrays have become synonymous with performance. However, they have also been regarded as expensive and lacking the full set of enterprise features, and are better used as a solution for a specific pain point or a single application.

Flash prices have gone down considerably over the last few years and flash as a storage media has gained maturity and so has the all-flash array market. Customers are looking for all-flash arrays to replace HDD or hybrid-based arrays, however the market is crowded with all-flash products that are based on legacy architectures not fit for flash media, or products that are capable of servicing only a single application, not an entire datacenter.



In its sixth generation, the K2 manages to realize Kaminario’s vision of an all-flash datacenter making the K2 the right choice for IT departments of the on-demand world, supporting applications such as relational databases, virtual environments, private-cloud implementations, NoSQL databases, containerized applications and more. Kaminario’s scale-up and -out building block architecture (shown below) allows a customer to scale up and scale out in defined increments. If a customer wishes to scale up capacity, Kaminario can add disk shelves with predetermined amounts of capacity to the K-Block without charging for unrequired increased performance capability. If a customer wishes to scale out and increase performance, then the customer adds another K-Block. The K2 offers scalability from 60TB-4PB with throughput scalability from 3.2Gb/s to 25.6Gb/s and I/O scalability from 420K to 1.6M IOPS.

With its automatic variable adaptive block architecture, the K2 can handle any kind of workload natively, whether it be small, random 4K online transaction processing or large 128K analytical processing or any kind of mix in between, along with virtual server support for hundreds of different applications at the same time. No manual tuning of any kind is needed to do this, the K2 does it all automatically. No complex setup, configuration, or maintenance is required.

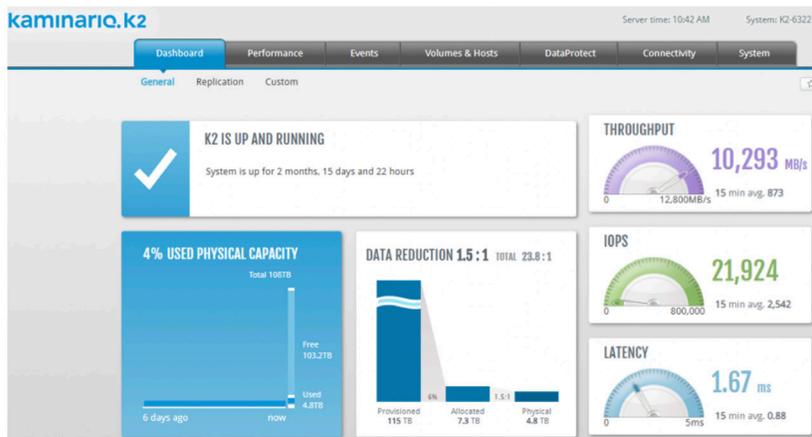
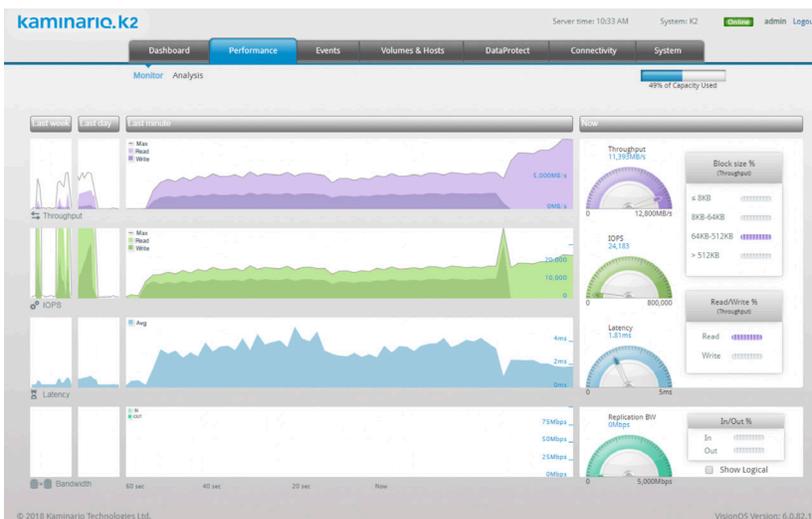
The K2 array is an appliance which is a combination of tested and packaged hardware, software and services. The K2 Gen6 hardware platform is based on leading off-the-shelf enterprise components that enable K2’s software-defined architecture and software stack. The K2 runs Kaminario VisionOS™, the next-generation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

#### K2: Cost Effective, Full Featured, General Purpose All-Flash Array for Consolidation of Moxed Workloads



VisionOS enables modular components and services that demonstrate a superior value proposition across a real scale-out 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 as the 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.



## Broadcom Data Warehouse Fast Track Reference Architectures for SQLServer 2017

Broadcom, Kaminario and Microsoft DWFT reference architectures contain the latest technology advancements in database, server, and storage technology. Broadcom's Emulex Fibre Channel HBAs, Kaminario K2 all-flash arrays, HPE servers with Intel® Xeon® processors, and Microsoft SQL Server 2017 are the latest additions to the list of reference architecture components. The Broadcom/Kaminario DWFT reference architectures for SQL Server 2017 are engineered jointly by Broadcom, Kaminario and Microsoft. The hardware and software optimizations are tested by Broadcom and Kaminario and the performance results are crosschecked by Microsoft. This approach presents a fast time to value using integrated, balanced, and verified architectures.

This specific Broadcom/Kaminario 100TB data warehouse reference architecture is an example in storage performance optimization. Key to the observed high IO performance is the use of Emulex 32GFC storage controller technology on both the database server as well as the K2, which provided more than 3200MB/s of IO throughput per port. By using multiple 32GFC data paths from the SQL Server host to the K2 we were able to achieve peak observed storage IO rates that exceeded 12,800MB/s.

### Single Server Configuration

The server hardware used was an HPE DL380 Gen10 with 2x Intel Xeon Scalable processor Gold 6148 CPUs running at 2.4Ghz each with 20 physical cores and 80 total logical cores with Hyperthreading enabled. System memory used was 128GB. System BIOS Configuration:

Setting Name	Selection
Workload Profile	IO Throughput
Intel Hyper-Threading	Enabled
Power Regulator	Static High Performance Mode

### Gen 6 Fibre Channel IO Connectivity

In the server we used three Emulex LPe32002 dual port 32GFC HBAs with all six ports connected to a Brocade Gen 6 X6-4 director class switch. The K2 was configured with two K Blocks each with dual active/active controllers. Each K2 controller contains an Emulex LPe32002 Gen 6 dual port 32GFC controller, the total number of Gen 6 32GFC target ports connected to the Brocade switch is eight. The Brocade Gen 6 Fibre Channel switch was configured to use fabric zoning to isolate SAN visibility to just the SQL server HBA ports and the ports of the Kaminario K2 array. Brocade Network Advisor was used to monitor SAN traffic flows for path fairness during the Fast Track benchmark. MPIO was used to create a high availability, load balanced IO connectivity that is capable of driving 19,200MB/s of I/O into the host server.

## Kaminario K2 Configuration

The K2 used has a capacity of 125TB of uncompressed storage. We created a volume group of 8TB in preparation for the SQL Server Fast Track project. A total of four volumes were created, three of 20TB size that were used for Data, Log and TempDB. A fourth volume of 50TB was created for the purpose of hosting the various database backup archives that were created throughout the testing process. A Host was defined for the Windows OS and all available initiator ports were assigned to, in addition to the four volumes being mapped to the host profile.

## Windows Server 2016 Configuration

The server was installed with the full GUI version of Microsoft Windows Server 2016 Datacenter version 10.0.14393 build 14393. The MPIO feature was installed and configured to use the K2. The Windows 2016 power plan was left at its recommended setting of "Balanced". Net Framework 4.6 was installed. PowerShell ExecutionPolicy was set for "Unrestricted". Using Windows Disk Management the four K2 volumes were brought online, partitioned and formatted. The 50TB volume was assigned the drive letter K. A directory was created as c:\mount which was used as the mount point for the three 20TB volumes.

Kaminario Volume	Size	Mount Point	Purpose
Data	20TB	C:\Mount\Data	SQL 2017 Data Files
TempDB	20TB	C:\Mount\TempDB	TempDB files
LOG	20TB	C:\Mount\LOG	Database log files
Backup	50TB	K:\	Backup volumes of various test checkpoints

## SQL Server 2017 Configuration

SQL Server 2017 Datacenter was installed, build 14.0.1000.169.

SQL Server Setting	Settings	Comment
Startup Flag	-E	Allocate multiple extents per file group
Startup Flag	-T1117	Grow data files uniformly
Degree of Parallelism	0	Will use all logical processors (80)
Resource Governor	12%	Memory grant restriction to increase query concurrency
SQL Max Memory	118GB	Required for Fast Track Certification
TempDB	80 files	1 per logical CPU
Data Volumes	1 c:\Mount\Data\	Autogrow_all_files

DWFT for SQLServer 2017 Certification

DWFT Certification #2017-015	Broadcom Emulex / Kaminario DWFT Reference Architecture		Report Date: 5/16/2018		
DWFT Rev. 5.4					
<b>System Provider</b>	<b>System Name</b>	<b>Processor Type</b>	<b>Memory</b>		
	Broadcom / HPE DL380 Gen 10 / Emulex Gen 6	Intel Xeon Gold 6148 2.4 GHz (2/40/80)	2048 GB		
<b>Operating System</b>		<b>SQL Server Edition</b>			
Windows Server 2016		SQL Server 2017 Enterprise Edition			
<b>Storage Provider</b>	<b>Storage Information</b>				
	Kaminario K2 1x 20TB LUN for tempdb Kaminario K2 1x 20TB LUN for Data Kaminario K2 1x 20TB LUN for LOG				
<b>Primary Metrics</b>					
Rated User Data Capacity <sup>1</sup> (TB)	Row Store Relative Throughput <sup>2</sup>	Column Store Relative Throughput <sup>3</sup>	Maximum User Data Capacity <sup>1</sup> (TB)		
100	296	517	139		
<b>Row Store</b>					
Relative Throughput <sup>2</sup>	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/O Throughput (MB/Sec)	Measured CPU (Avg.) (%)
296	350	7,858	9,601	8,729	96
<b>Column Store</b>					
Relative Throughput <sup>2</sup>	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/O Throughput (MB/Sec)	Measured CPU (Avg.) (%)
517	3,359	2,242	N/A	N/A	98
The reference configuration is a 2 socket system rated for 25TB using SQL Server 2014 and the DWFT V4 methodology					
<sup>1</sup> Assumes a data compression ratio of 5:1					
<sup>2</sup> Percent ratio of the throughput to the row store throughput of the reference configuration.					
<sup>3</sup> Percent ratio of the throughput to the column store throughput of the reference configuration.					
* Reported metrics are based on the qualification configuration which specifies database size and SQL Server memory.					

## Summary

This paper describes a reference architecture using Broadcom's Emulex Gen 6 Fibre Channel Host Bus Adapters, a Kaminario K2 all-flash array and HPE DL380 Gen10 servers. By implementing Data Warehouse Fast Track for SQL Server 2017 design principles, this configuration achieved a 100TB rating.

Broadcom and Kaminario dedicated hundreds of hours of testing to engineer the SQL Server DWFT solution to provide optimal reliability and performance. These series of tests pushed Emulex Gen 6 HBAs, the Kaminario K2 all-flash array and HPE DL380 Gen 10 servers to peak performance without hardware failure. The reliability and performance experienced during testing is what can be expected in production environments.

The test provided a sustainable throughput of at least 12,800MB/s, which would allow a DBA the flexibility to simultaneously run an OLTP and OLAP workload, move tables around as needed and generate heavy reports quickly and efficiently without impacting production performance.

### **The Broadcom/Kaminario Reference Architectures Provide the Following Benefits:**

- Deliver a tested and Microsoft validated configuration with proven methodology and performance behavior.
- Deliver outstanding performance with cutting-edge Emulex Gen 6 HBAs and Kaminario K2 all-flash storage arrays.
- Achieve a balanced and optimized solution at all levels by following best practices for both hardware and software components.
- Offer high availability at all levels of setup (host, switches, and storage).
- Help customers avoid the pitfalls of an improperly designed and configured system.
- Reduce future support costs by limiting solution re-architect efforts due to scalability challenges.

## References

### **Additional Resources:**

[Emulex Fibre Channel HBA solutions](#)

[Kaminario K2 solutions](#)



## 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](http://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.