



PowerFlex



Breakthrough Performance:

Dell EMC PowerFlex Family for
Cloud-Native Applications

DELL Technologies

KIOXIA

Cloud-native technologies and applications empower organizations to build and run scalable applications in modern and dynamic multi-cloud environments.

Increasingly, the enterprises are seeking the high performance and scalability of NoSQL databases for cloud-native applications that support AI, Machine Learning, and IoT. The NoSQL database has become the new norm for all businesses supporting modern applications as they are scalable on X86 platforms due to their versatility in storing, processing, and accessing all kinds of business data. The PowerFlex software-defined storage platform with KIOXIA PM5 Series SAS SSDs is easier to manage than traditional storage and reduces operational complexities to deliver better performance and serve the basic needs of cloud-native Kubernetes stateful applications.

CONSUMPTION OPTIONS

- ✓ PowerFlex Appliance
- ✓ PowerFlex Rack
- ✓ VxFlex Ready Nodes



BASIC REQUIREMENTS OF CLOUD-NATIVE KUBERNETES STATEFUL APPLICATIONS:



Optimized storage platform for cloud-native applications that demand low latency and high throughput.



Linear scaling of container platform with consistent response times.

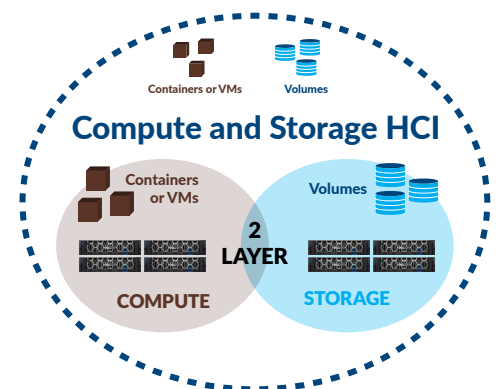


Create, control and manage volumes in your storage platform for Kubernetes stateful applications.

Testing by Cloud Evolutions verified that the PowerFlex platform is a high-performance solution—both resilient and highly scalable.

The PowerFlex family is designed for extensive fault tolerance, scalability and throughput, to deliver six 9's tier 1 availability for business-critical applications. This family offers flexible consumption options in the form of rack, appliance or ready nodes, with three deployment options.

These features are essential for applications such as streaming, cloud-native development, and other demands of today's medium-to-large enterprises. The key differentiator of this platform is that it provides high throughput at low latencies that scales linearly to 1000s of nodes in a self-balancing, fully lifecycle-managed deployment.



DEPLOYMENT OPTIONS

- ✓ 2 Layer (Compute and SAN)
- ✓ Storage Only
- ✓ Single Layer HCI

An Optimized Platform for Microservices

DELL EMC POWERFLEX FAMILY RISES TO THE OCCASION WITH ITS

- ✓ Unmatched Flexibility
- ✓ Simple and Efficient Operation
- ✓ Superior Elasticity



Dell EMC PowerFlex
for **Performance**



Dell EMC PowerFlex
CSI driver for
Persistent Storage



Kubernetes for
**Automation and
Orchestration**

Containers provide an easy way to achieve the goals of cloud-native applications because of their footprint, manageability, portability and packaging. As a way to virtualize operating systems, containers enable rapidly scalable application architectures that reduce time-to-market and eliminate risks associated with monolithic applications, and are quickly becoming the new standard.

As developers and business managers adopt containers for microservices to meet modern market requirements, they strive to enhance the scalability of their applications using Kubernetes for stateful applications. The use of Kubernetes to support data-driven applications helps organizations use containers to modernize legacy applications and develop additional use cases.

Thus the ever-growing requirements of containerized applications for performance and scale make Dell EMC PowerFlex family the choice for stateful applications such as NoSQL databases. Container runtime APIs restrict elasticity and challenge the scalability of distributed databases that can spread across multiple hosts. Containers need to be managed and connected to the outside world for tasks such as scheduling, load balancing and distribution, so this is where a container orchestration tool such as Kubernetes comes in handy. An open-source system for deploying, scaling and managing containerized applications, Kubernetes handles the work of scheduling containers onto a compute cluster and manages the workloads to ensure that the containers run as the user intended.

The PowerFlex family is optimized for Kubernetes for three reasons: persistent storage, its support for storage volumes using the CSI driver, and performance. Storage solutions on the market lack the performance characteristics necessary to support today's data-heavy applications with any reasonable scale or elasticity. PowerFlex can aggregate storage resources and distribute data uniformly across hundreds of nodes, busting an IO bottleneck and providing abundant throughput and scalability.

Stateful application performance and scaling challenges are addressed with storage technologies such as SAS SSDs and NVMe (Non-Volatile Memory Express), so that storage is no longer the bottleneck.

Persistent Storage is supported through PowerFlex leveraging a Container Storage Interface (CSI) plugin for Kubernetes, supporting the broadest set of features for block storage integration, including dynamic provisioning. This CSI plugin is an interface between the logical volumes in the Kubernetes environment called persistent volumes and the storage volumes or LUNs.

Virtues Of Flash: KIOXIA PM5 Series

For the optimal user experience, SQL and NoSQL database workloads require high transactional performance and consistently low latencies. KIOXIA PM5 Series 12Gb/s enterprise SAS SSDs are perfect for applications demanding the highest performance, reliability, and scale.

These drives offer the rich feature set needed by organizations that rely on modern data loads.



Performance. Up to 82% higher performance levels than the previous generation.



Reliability. Two flash die failure recovery, power loss protections and end-to-end data protection.



Quality of Service. Multi-stream write SSD that improves performance and endurance.



Data Security. Security options that do not impact performance, from no encryption to full FIPS 140-2 certification.

Our Testing Methodology

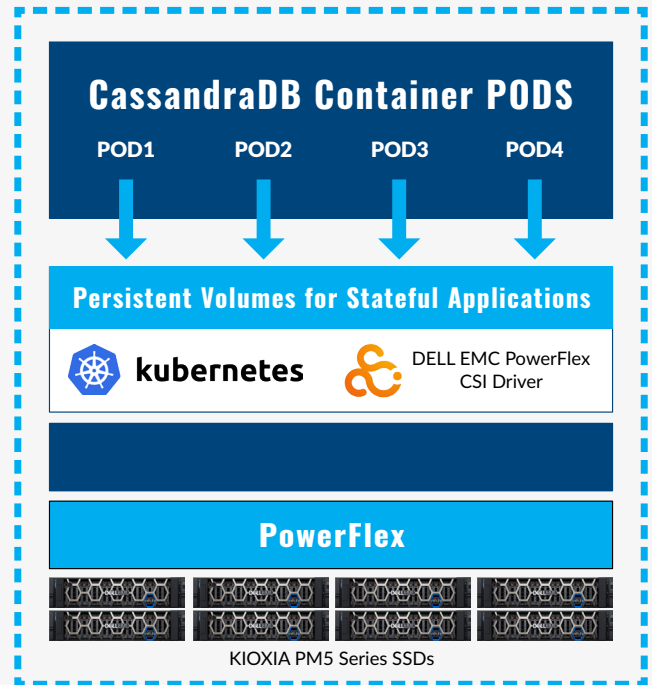
The objective of the testing conducted by Cloud Evolutions was to benchmark performance of the PowerFlex Software Defined Storage platform with KIOXIA PM5 Series SAS SSDs, with Kubernetes orchestrating stateful applications on its platform.

The test configuration of the platform that demonstrated superior performance and scaling capabilities is shown in Appendix A. Key performance indicators included throughput, measured as IOPS (input/output operations per second) and operations per second (ops/sec), and millisecond latency under a variety of workloads. To validate the claims, two tests were designed to measure the PowerFlex platform performance.

The test platform consists of eight nodes, each with an instance of Kubernetes and CassandraDB PODs, in order to assess:

1. **Input/output (IO)**, using the FIO tool to benchmark the platform with Dell EMC PowerFlex and Kubernetes PODs, capturing IOPS and latency.
2. **Platform performance**, using the Cassandra Stress tool when CassandraDB is deployed in containers orchestrated by Kubernetes.

Operations/sec and median latency are considered to measure the performance of Cassandra deployed on the PowerFlex family.



All About the Performance

The results of performance benchmarking executed on Kubernetes PODS on PowerFlex proved exceptional. We observed 1.39 million IOPS with sub-millisecond latency on this PowerFlex family with KIOXIA PM5 Series SAS SSDs optimized for Kubernetes.

**SUSTAINED
1.39M IOPS
IN JUST 8U
FOOTPRINT**



**SUB-
MILLISECOND
LATENCY**




The charts below depict results of a variety of read, write and mixed read-and-write workloads. Thread counts were varied to identify “sweet spots” achieving maximum throughput of operations per second. Across all types of workloads—read, write and mixed—the PowerFlex software-defined storage platform demonstrated high performance, resilience and scalability well beyond the expectations of any traditional architectures. With reads and writes fairly distributed on each node in the cluster, the results showed that the PowerFlex family achieved:

- Total data populated within 32 Cassandra pods: 18 Billion records (~4.5 TB).
- Total records read: 320 Million records.
- Over 302,000 operations per second (ops/sec) for 100% reads.
- Over 271,000 ops/sec for 100% writes.
- Over 260,000 ops/sec for mixed workloads of 50 percent reads and 50 percent writes, with median latency of 1.2 milliseconds at 150 threads.
- Consistent linear scaling of ops/sec as more CassandraDB PODS were added to the cluster, from 8 to 32 PODS, without impacting latency.


Cassandra 100% Reads

For 100% reads from the Cassandra stress tool, the platform achieved 302k Ops/sec, by far highest in its class and configuration, while keeping its latency under 1 millisecond




Cassandra 50% Read 50% Write

For mixed workload tests using the Cassandra stress tool, the platform achieved 264k Ops/sec, by far highest in its class and configuration, while maintaining a latency of just over 1 millisecond

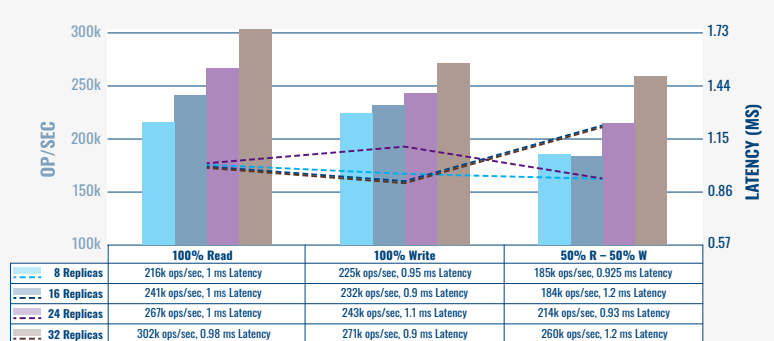


Cassandra 100% Writes

For 100% write workload tests using the Cassandra stress tool, the platform achieved 271k Ops/sec, by far highest in its class and configuration, while keeping its latency under 1 millisecond



Cassandra Performance on Multiple Kubernetes PODS



Results Matter


The PowerFlex platform tested demonstrated rock-steady performance, scalability and data availability for the Kubernetes stateful application, CassandraDB, using the Dell EMC PowerFlex CSI driver. The Dell EMC PowerFlex CSI driver provided persistent volumes for dynamic and static container provisioning, empowering users with efficient platform scaling and durability. With data stored on numerous nodes and in numerous places, PowerFlex distributed data efficiently for CassandraDB, allowing virtually linear scalability, reading and writing data fast with near-constant data availability.

The PowerFlex family with KIOXIA PM5 Series SAS SSDs handled heavy read and write workloads very efficiently, empowering the admins to plan data distribution across multiple servers, clusters and data centers.

Our testing supports the idea that if your real-time data is growing exponentially, this optimized platform—Dell EMC PowerFlex family, KIOXIA SSDs, Kubernetes, and CassandraDB—is the preferred and most robust solution for NoSQL distributed databases. Testing highlighted that CassandraDB's performance on the PowerFlex family, orchestrated by Kubernetes, can offer new levels of reliability, flexibility, and performance.

Dell Technologies' PowerFlex family delivers superior performance, scalability, and reliability in a basic 8-node, multiple Cassandra POD configuration. Under a range of dynamic loads—even those outstripping traditional ideas of IOPS under certain reads and writes—the PowerFlex family proves itself both resilient and self-managing. For medium to large enterprises seeking rapid deployment on a more efficient footprint, the PowerFlex family delivers.

KEY RESULTS



Consistent Scaling and Uniform Distribution of Workload



**Exceptional Performance
8 Node, 32 POD CassandraDB Configuration**



Proven Sub-Millisecond Latency for all Types of Workloads: Reads, Writes, and Mixed



Persistent Volume Support Using PowerFlex CSI Driver to Manage Volumes Consumed by Kubernetes

Appendix A

PowerFlex Node	8 x PowerFlex Nodes (R640 Servers)
	PowerFlex Version: R 3.0.1
	CPU: 2 x Intel® Xeon® Gold 6132 (18C, 2.60 GHz)
	Memory: 384-GB RAM((24 x 16 - GB DIMMs)
	Storage: KIOXIA PM5 Series 12Gb/s SAS SSDs 10 x 1.92 TB
	Persistent Volume: 1 TB for cassandra POD
Network	2 x 25 GB Mellanox
Server BIOS Version	3.21.23.22
SOFTWARE	
Operating System	CentOS 7.6
Kubernetes	v1.13
CSI Plugin	v1.0
Cassandra Image	3.11.6

DISCLAIMER OF WARRANTIES & LIMITATION OF LIABILITY Cloud Evolutions, Inc. has made reasonable efforts to ensure the accuracy and validity of its testing. However, Cloud Evolutions, Inc. specifically disclaims any warranty, expressed or implied, relating to the test results and analysis, their accuracy, completeness or quality, including any implied warranty of fitness for any particular purpose. All persons or entities relying on the results of any testing do so at their own risk, and agree Cloud Evolutions, Inc. its employees and its subcontractors shall have no liability whatsoever from any claim of loss or damage on account of any alleged error or defect in any testing procedure or result. In no event shall Cloud Evolutions, Inc. be liable for indirect, special, incidental, or consequential damages in connection with its testing, even if advised of the possibility of such damages. In no event shall Cloud Evolutions, Inc.'s liability, including for direct damages, exceed the amounts paid in connection with Cloud Evolutions, Inc.'s testing. Customer's sole and exclusive remedies are as set forth herein. Testing is conducted under laboratory conditions using synthetic benchmark tools. Real-world performance may vary.

© 2020 Cloud Evolutions, Inc. All trademarks are the property of their respective owners.