

# AEROSPIKE CERTIFICATION TOOL (ACT) BENCHMARK ON OPENFLEX F3200



## Contents

1	Executive Summary.....	2
2	Solution Highlights.....	2
3	Technology Overview.....	2
3.1	OpenFlex™ F3200 and E3000 Overview.....	2
3.2	Aerospike Overview.....	2
4	OpenFlex™ F3200 with Aerospike.....	3
4.1	Aerospike Certification Tool (ACT) Benchmark on OpenFlex™ F3200 .....	3
4.2	ACT Tests & Configuration Details .....	3
4.3	ACT Benchmark results on OpenFlex™ F3200 .....	3
5	Summary .....	3
6	Resources & Additional links .....	3

# 1 Executive Summary

Faster data access and scalability has become a top priority for every business. To keep pace with constantly growing capacity demands, Organizations needs to keep storage costs as low as possible. But operating a vast storage infrastructure, which spans across multiple data centers around the globe, was becoming increasingly complex and expensive to manage. Relying on storage solutions from multiple vendors, each of which used separate interfaces and support organizations, exacerbated the problem. Additionally, organizations are needed to continue delivering rock-solid reliability as it added capacity over time. Companies whose business are built on providing reliable cybersecurity and data protection, the organizations could not afford to sacrifice quality to lower costs.

For firms struggling to manage high and rapidly growing volumes of operational data in real-time, it is worth exploring how recent advances in database management technology provided by Aerospike with OpenFlex™ Solution can help achieve these goals.

OpenFlex™ architecture and products allow storage to be disaggregated from compute, enabling applications to share a common pool of storage capacity. Data can easily be shared between applications or needed capacity can be allocated to an application regardless of location. By decoupling software from the underlying platform enterprises can build solutions with the greatest flexibility spanning the portfolio of industry standard OpenFlex™ servers hardware offerings. This provides a decisive step forward in reducing the cost of ownership for data centers deployments to come and it addresses all the problems which is caused by the Volume, Variety, Velocity, Access Speed and Consistency of data. Aerospike provides a database system that processes large volumes of operational data in real time while delivering exceptional runtime performance, high availability, and cost efficiency while still keeping your data safe.

The open source Aerospike Certification Tool (ACT) for Solid State Drive benchmarks helps to measure latency under large workloads and increasing throughput. The OpenFlex™ Composable Platform and Aerospike together provides capabilities to IT operators that enable them to connect disaggregated resources intelligently and manage, modify, and scale these components over time. With the combined solution IT can achieve Greater productivity, agility, performance and faster time-to-market Solution.

## 2 Solution Highlights

The purpose of this document is to showcase the ACT benchmark on OpenFlex™ F3200. The open source ACT helps in evaluating the latencies the device is capable of under the large workloads and increase in throughput. The document reports the ACT benchmark results achieved with OpenFlex™ F3200 device and the configuration used to achieve the same. With the ACT testing, we have validated that high performance is achievable with little or no loss in speed and latency when compared to Direct Attached Storage (DAS). OpenFlex™ F3200 is a fabric device that leverage this OCI approach in the form of disaggregated data storage using NVMe-over-Fabrics (NVMe-oF). The Aerospike database platform is modelled on the classic shared-nothing database architecture. The database cluster consists of server nodes, each of which has CPUs, DRAMs, and rotational disks (HDDs) and optional flash storage units (SSDs).

## 3 Technology Overview

### 3.1 OpenFlex™ F3200 and E3000 Overview


With the exponential growth in data, along with the increasing diversity of workflows and demands on IT infrastructure, businesses need to increase speed, agility, and time-to-value for their customers. Emerging as a solution for this, composable infrastructure is a new architectural approach that—using **NVMe™-over-Fabrics**—will vastly improve compute and storage utilization, performance, and agility in the data center. OpenFlex is based on scale-out performance and open composability.



OpenFlex is Western Digital’s architecture that supports **Open Composable Infrastructure (OCI)**, OpenFlex F3200 and E3000 are fabric devices that leverage this OCI approach in the form of disaggregated data storage using **NVMe-over-Fabrics (NVMe-oF)**. NVMe-oF is a networked storage protocol that allows storage to be disaggregated from compute to make that storage widely available to multiple applications and servers.

By enabling applications to share a common pool of storage capacity, data can be easily shared between applications, or needed capacity can be allocated to an application regardless of location. Exploiting NVMe device-level performance, NVMe-oF promises to deliver the lowest end-to-end latency from application to shared storage. NVMe-oF enables composable infrastructures to deliver the data locality benefits of NVMe DAS (low latency, high performance) while providing the agility and flexibility of sharing storage and compute.

#### OpenFlex F3200 Specification:




Specification	Value
Max Raw Data Storage Capacity per device	61.4 TB
Data Ingest Capability	2x 50G Ethernet
Data Transfer Rates	12 GBps *
Number per enclosure	Up to 10
Hot Swappable	Yes
Service Window	5 minutes
Dynamic Provisioning support	Yes

The **maximum data storage capacity is 614TB \*** when leveraging a full set of 10 F3200 fabric devices. **Each E3000 is capable of hosting 10 quantity of F3200.** F3200 is capable of scaling up to **2Million IOps** and cumulatively we can scale for each E3000 up to **20 Million IOps** in a 3U solution.

#### F3200 Data Plane:



**System Data Ingest Architecture:** The system main data ingest architecture uses two separate 50G Ethernet connections each on a dual QSFP28 connector on the rear I/O of the chassis.



This completes the connection from the device that is inserted into a chassis slot, through the backplane into the QSFP connectors. The architecture supports the hot swap nature of the devices and do not require any sort of shut down or disconnection before servicing. Each 100G Ethernet connection is split in half at the QSFP28 connectors resulting in 50G per connector allowing for dual port functionality with the device.

F3200 uses NVMe over Fabrics technology with NVMe-RoCEv2 protocol for all network transmission.

Read more @ [Western Digital OpenFlex Composable Infrastructure](#)

### 3.2 Aerospike Overview

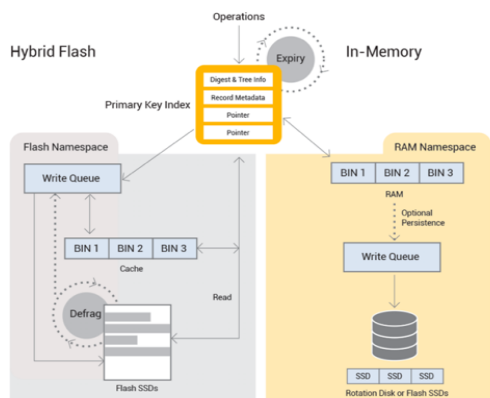
Aerospike is a distributed NoSQL system that provides extremely fast – and predictable – read/write access to operational data sets that span billions of records in databases of 10s – 100s TB. Its patented Hybrid Memory Architecture™ (HMA) delivers exceptional performance using a much smaller server footprint than competing solutions. Aerospike uses dynamic random-access memory (DRAM) for index and user data. Optionally, applications can commit each write directly to fast, non-volatile memory (solid state disks or SSDs), which Aerospike treats as raw devices for Aerospike, which is a Distributed, Scalable database.

Aerospike enterprises overcome seemingly impossible data bottlenecks to compete and win with a fraction of the infrastructure complexity and cost of legacy NoSQL databases. Aerospike’s patented Hybrid Memory Architecture™ delivers an unbreakable competitive advantage by unlocking the full potential of modern hardware, delivering previously unimaginable value from vast amounts of data at the edge, to the core and in the cloud. Aerospike empowers customers to instantly fight fraud; dramatically increase shopping cart size; deploy global digital payment networks; and deliver instant, one-to-one personalization for millions of customers.

The Aerospike Architecture mainly has the below key objectives:

- ✓ Create a flexible, scalable platform for web-scale applications.
- ✓ Provide the robustness and reliability expected from traditional databases.

✓ Provide operational efficiency with minimal manual involvement.  
Aerospike Storage Architecture:



To meet the demanding needs of real-time enterprise applications, Aerospike namespaces (databases) can be configured to operate with strong consistency (SC), which prevents stale reads, dirty reads, and data loss. Firms can also configure Aerospike to relax its strong consistency support and operate in available / partition tolerant (AP) mode, which offers backward compatibility with early Aerospike releases. However, given that performance is comparable in both modes and that strong consistency simplifies application logic, Aerospike expects most firms to prefer to configure their namespaces to operate with strong consistency. It's important to note that Aerospike's support for consistency is immediate regardless of its operational mode

Read more @ [Aerospike](#)

## 4 OpenFlex™ F3200 with Aerospike

Enterprises are adopting NoSQL databases, such as Aerospike, and replacing their relational database deployments. It allows apps to get high-performance, quick processing of information at a massive scale. Composable Infrastructure is created for these environments where elasticity, fast deployment, and rapid scale are paramount. The Solution design is based on OpenFlex™ F3200 with Aerospike using the Aerospike Hybrid storage.

### 4.1 Aerospike Certification Tool (ACT) Benchmark on OpenFlex™ F3200

The Aerospike Certification Tool (ACT) is an open source tool that Aerospike provides for stress-testing SSD drives and measuring latencies under stress. Instead of relying on third-party benchmarks, you can use this tool to generate



performance metrics in your environment. You can determine the load and measure latencies on the device. Testing in your environment allows you to eliminate all other variables and isolate the SSDs as the sole variable. Aerospike provides ACT so that you can stress test any SSDs in your local environment. When you perform local stress testing, you can document precise latencies for the drives/devices in your environment. You can also easily compare the latencies of drives from disparate manufacturers without having to accommodate for environmental differences.

As part of Aerospike Certification Test on OpenFlex™ F3200 device, we have executed Aerospike ACT Storage Performance test on the OpenFlex™ F3200 device to test the maximum load the device can handle as per ACT Storage performance test criteria.

### 4.2 ACT Tests & Configuration Details

The ACT tests are intended to evaluate the performance of SSDs for their real-time data needs.

Below configuration has been setup for ACT test on OpenFlex™ F3200

Product	OpenFlex™ F3200 Fabric Device
Interface	Dual QSFP28 (2x50Gb)
Host Server	Dual Socket server with 20 Core CPU each.80 logical cores in total with HT enabled
Host OS	Red Hat Enterprise Linux Server release 7.6 (Maipo)
Kernel	3.10.0-957.el7.x86_64
Physical Core	80
Host NIC	CX5 - MCX516A-CCAT
NIC Package Version	4.6-1.0.1.1
ACT version	5.3
No. of Volumes	10

A profile has three fundamental characteristics:

- Read / write ratio
- Object size
- Latency requirements of read operations

The test is run by increasing throughput and determining whether the latency requirement is met --and increasing throughput again, if the device is within latency requirements, until the latency SLA is no longer met.

The purpose of this certification is:

- Determine if an SSD device(s) will stand up to the demands of a high-speed real-time database
- Evaluate the upper limits of the throughput you can expect from a device(s).

### 4.3 ACT Benchmark results on OpenFlex™ F3200

OpenFlex™ F3200 device is tested with ACT at the specified speed with a 67% read, 33% write ratio of 1.5 KB objects over 24 hours.

ACT was able to achieve a very high performance of **1.158M tps (1158K tps)** Below is the detail of the latencies & tps achieved with ACT 5.3

Refer to below snippet from Aerospike website where the OpenFlex results are published. [ACT SSD Certification](#)

Storage Solution	Speed (tps)	>1ms	>8ms	>64ms	Device	Protocol	ACT	Source
Western Digital	1,158,000	4.49%	0.00%	0.00%	OpenFlex F-Series	NVMe-oF	5.3	Western Digital

Below are the recommendations to achieve better performance

- ✓ Based on available budget, customer can opt for higher CPU/Memory configuration to achieve maximum results.
- ✓ Similarly, depending on the data requirements, customers can either go for lower capacity or higher capacity Fabric device.

## 5 Summary

The ACT certification results proves that high performance is achievable with little or no loss in speed and latency when compared to Direct Attached Storage (DAS) with OpenFlex™ F3200 device. Customers can also save up to 70% in costs by eliminating overprovisioned storage and compute resources with this solution. With the Low Latencies, High Performance levels, Consistencies & Scalability of OpenFlex, it will provide a new range of operational capabilities for Aerospike. OpenFlex & Aerospike can be used in E-commerce, AL-ML, Edge Databases, Spark, Kafka, Query and Reporting, Transactional Database Deployments.

## 6 Resources & Additional links

ACT test results: [ACT SSD Certification](#)

OpenFlex: [Western Digital OpenFlex Composable Infrastructure](#)