



# FlexPod for Epic Performance Testing

## FlexPod

NetApp  
September 12, 2021

# Table of Contents

- FlexPod for Epic Performance Testing ..... 1
- TR-4784: FlexPod for Epic Performance Testing ..... 1
- Executive Summary ..... 5
- Test methodology ..... 5
- Workload testing ..... 7
- Conclusion ..... 16
- Where to find additional information ..... 16
- Acknowledgements ..... 18

# FlexPod for Epic Performance Testing

## TR-4784: FlexPod for Epic Performance Testing

Brian O'Mahony,  
Ganesh Kamath  
Atul Bhalodia  
Brandon Agee

In partnership with:



### Objective

The objective of this report is to highlight the performance of FlexPod with NetApp All Flash A300 and A700 storage systems with Epic Healthcare workloads.

### Epic Hardware configuration guide

For acceptable end-user performance, Epic production and disaster recovery operational database (ODB) target- read and target- write time requirements are as follows:

- For randomly placed reads to database files measured at the system call level:
  - Average read latencies must be 2ms or less
  - 99% of read latencies must be below 60ms
  - 99.9% of read latencies must be below 200ms
  - 99.99% of read latencies must be below 600ms
- For randomly placed writes to database files measured at the system call level:
  - Average write latencies must be 1ms or less depending on size



These requirements change with time. Epic prepares a customer- specific Epic Hardware Configuration Guide (HCG). Refer to your HCG for details on requirements.

### Overall solution benefits

By running an Epic environment on a FlexPod architectural foundation, healthcare organizations can see an improvement in staff productivity and a decrease in capital and operating expenses. FlexPod Datacenter with Epic delivers several benefits specific to the healthcare industry:

- **Simplified operations and lowered costs.** Eliminate the expense and complexity of legacy proprietary RISC/UNIX platforms by replacing them with a more efficient and scalable shared resource capable of supporting clinicians wherever they are. This solution delivers higher resource utilization for greater ROI.
- **Quicker deployment of infrastructure.** Whether it's in an existing data center or in a remote location, the integrated and tested design of FlexPod Datacenter with Epic enables customers to have new infrastructure up and running in less time with less effort.

- **Scale-out architecture.** Scale SAN and NAS from terabytes to tens of petabytes without reconfiguring running applications.
- **Nondisruptive operations.** Perform storage maintenance, hardware lifecycle operations, and software upgrades without interrupting business operations.
- **Secure multitenancy.** FlexPod supports the needs of shared virtualized server and storage infrastructure, enabling secure multitenancy of facility-specific information, particularly if you are hosting multiple instances of databases and software.
- **Pooled resource optimization.** FlexPod can help reduce physical server and storage controller counts and load- balance workload demands. It can also boost utilization while improving performance.
- **Quality of service (QoS).** FlexPod offers QoS on the entire stack. Industry-leading QoS storage policies enable differentiated service levels in a shared environment. These policies enable optimal performance for workloads and help isolate and control runaway applications.
- **Storage efficiency.** Reduce storage costs with the NetApp 7:1 storage efficiency guarantee.
- **Agility.** The industry-leading workflow automation, orchestration, and management tools offered by FlexPod systems allow IT to be far more responsive to business requests. These business requests can range from Epic backup and provisioning of additional test and training environments to analytics database replications for population health management initiatives.
- **Productivity.** Quickly deploy and scale this solution for optimal clinician end-user experiences.
- **Data Fabric.** The NetApp Data Fabric architecture weaves data together across sites, beyond physical boundaries, and across applications. The NetApp Data Fabric is built for data-driven enterprises in a data-centric world. Data is created and used in multiple locations, and it often must be leveraged and shared with other locations, applications, and infrastructures. Customers want a way to manage data that is consistent and integrated. It provides a way to manage data that puts IT in control and simplifies ever-increasing IT complexity.

## Cisco Unified Computing System, Cisco Nexus and MDS Switching, and ONTAP all-flash storage

The FlexPod for Epic Healthcare delivers the performance, efficiency, manageability, scalability, and data protection that IT organizations need to meet for the most stringent Epic requirements. By accelerating Epic production database performance and by reducing application deployment time from months to weeks, FlexPod helps organizations maximize the potential of their Epic investment.

### Cisco Unified Computing System

As a self-integrating, self-aware system, Cisco Unified Computing System (UCS) consists of a single management domain interconnected with a unified I/O infrastructure. The Cisco UCS for Epic environments has been aligned with Epic infrastructure recommendations and best practices to help make sure that infrastructure can deliver critical patient information with maximum availability.

The foundation of Epic on the Cisco UCS architecture is Cisco UCS technology with its integrated systems management, Intel Xeon processors, and server virtualization. These integrated technologies solve data-center challenges and enable you to meet your goals for data- center design for Epic. Cisco UCS unifies LAN, SAN, and systems management into one simplified link for rack servers, blade servers, and virtual machines. The Cisco UCS is an end-to-end I/O architecture that incorporates Cisco Unified Fabric and Cisco fabric extender (FEX) technology to connect every component in the Cisco UCS with a single network fabric and a single network layer.

The system is designed as a single virtual blade chassis that incorporates and scales across multiple blade chassis. The system implements a radically simplified architecture that eliminates the multiple redundant

devices that populate traditional blade server chassis and result in layers of complexity. Examples include Ethernet switches, Fibre Channel switches, and chassis management modules. The Cisco UCS contains a redundant pair of Cisco fabric interconnects that provide a single point of management and a single point of control for all I/O traffic.

The Cisco UCS uses service profiles to help ensure that virtual servers in the UCS infrastructure are configured correctly. Service profiles include critical server information about the server identity such as LAN and SAN addressing, I/O configurations, firmware versions, boot order, network VLAN, physical port, and quality-of-service (QoS) policies. Service profiles can be dynamically created and associated with any physical server in the system within minutes rather than within hours or days. The association of service profiles with physical servers is performed as a single, simple operation that enables migration of identities between servers in the environment without any physical configuration changes. It facilitates rapid bare-metal provisioning of replacements for failed servers.

Using service profiles helps to ensure that servers are configured consistently throughout the enterprise. When using multiple Cisco UCS management domains, UCS Central can use global service profiles to synchronize configuration and policy information across domains. If maintenance is required in one domain, the virtual infrastructure can be migrated to another domain. Therefore, applications continue to run with high availability even when a single domain is offline.

Cisco UCS has been extensively tested with Epic over a multiyear period to demonstrate that it meets server configuration requirements. Cisco UCS is a supported server platform, as listed in customers' "Epic Hardware Configuration Guide."

### **Cisco Nexus and Cisco MDS Ethernet and Fibre Channel switching**

Cisco Nexus switches and MDS multilayer directors provide enterprise-class connectivity and SAN consolidation. Cisco multiprotocol storage networking reduces business risk by providing flexibility and options. Supported protocols include Fibre Channel (FC), Fibre Connection (FICON), FC over Ethernet (FCoE), SCSI over IP (iSCSI), and FC over IP (FCIP).

Cisco Nexus switches offer one of the most comprehensive data-center-network feature sets in a single platform. They deliver high performance and density for both the data center and the campus core. They also offer a full feature set for data-center aggregation, end-of-row deployments, and data center interconnect deployments in a highly resilient, modular platform.

The Cisco UCS integrates computing resources with Cisco Nexus switches and a unified I/O fabric that identifies and handles different types of network traffic, including storage I/O, streamed desktop traffic, management, and access to clinical and business applications.

In summary, the Cisco UCS provides the following important advantages for Epic deployments:

- **Infrastructure scalability.** Virtualization, efficient power and cooling, cloud scale with automation, high density, and performance all support efficient data-center growth.
- **Operational continuity.** The design integrates hardware, NX-OS software features, and management to support zero-downtime environments.
- **Transport flexibility.** Incrementally adopt new networking technologies with a cost-effective solution.

Together, Cisco UCS with Cisco Nexus switches and MDS multilayer directors provide a compelling computer, networking, and SAN connectivity solution for Epic.

### **NetApp all-flash storage systems**

NetApp AFF systems address enterprise storage requirements with high performance, superior flexibility, and

best-in-class data management. Built on ONTAP data management software, AFF systems speed up your business without compromising the efficiency, reliability, or flexibility of your IT operations. With enterprise-grade all-flash arrays, AFF systems accelerate, manage, and protect your business-critical data and enable an easy and risk-free transition to flash media for your data center.

Designed specifically for flash, AFF A-series all-flash systems deliver industry-leading performance, capacity, density, scalability, security, and network connectivity in a dense form factor. With the addition of a new entry-level system, the new AFF A-series family extends enterprise-grade flash to midsize businesses. At up to seven million IOPS per cluster with sub-millisecond latency, the AFF A series is the fastest family of all-flash arrays, built on a true unified scale-out architecture.

With the AFF A series, you can complete twice the work at half the latency relative to the previous generation of AFF systems. The members of the AFF A series are the industry's first all-flash arrays that provide both 40Gb Ethernet (40GbE) and 32Gb Fibre Channel (FC) connectivity. Therefore, they eliminate the bandwidth bottlenecks that are increasingly moving from storage to the network as flash becomes faster and faster.

NetApp has taken the lead for all-flash storage innovations with the latest solid-state-drive (SSD) technologies. As the first all-flash array to support 15TB SSDs, AFF systems, with the introduction of the A series, also become the first to use multistream write SSDs. Multistream write capability significantly increases the usable capacity of SSDs.

NetApp ONTAP Flash Essentials is the power behind the performance of All Flash FAS. ONTAP is industry-leading data management software. However, it is not widely known that ONTAP, with its NetApp WAFL (Write Anywhere File Layout) file system, is natively optimized for flash media.

ONTAP Flash Essentials optimizes SSD performance and endurance with the following features, among others:

- NetApp data-reduction technologies, including inline compression, inline deduplication, and inline data compaction, can provide significant space savings. Savings can be further increased by using NetApp Snapshot and NetApp FlexClone technologies. Studies based on customer deployments have shown that these data-reduction technologies have enabled space savings of up to 933 times.
- Coalesced writes to free blocks maximize performance and flash media longevity.
- Flash-specific read-path optimizations provide consistent low latency.
- Parallelized processing handles more requests at once.
- Software-defined access to flash maximizes deployment flexibility.
- Advanced Disk Partitioning (ADP) increases storage efficiency and further increases usable capacity by almost 20%.
- The Data Fabric enables live workload migration between flash and hard-disk-drive tiers on the premises or to the cloud.

QoS capability guarantees minimum service-level objectives in multiworkload and multitenant environments.

The key differentiators with adaptive QOS are as follows:

- Simple self-managing IOPS/TB or throughput MB/TB. Performance grows as data capacity grows.
- Simplified consumption of storage based on service-level performance policies.
- Consolidation of mixed workloads onto a single cluster with guaranteed performance service levels. No more silos are required for critical applications.
- Major cost saving by consolidating nodes and disk.

# Executive Summary

To showcase the storage efficiency and performance of NetApp's All Flash FAS platform, NetApp performed a study to measure Epic EHR performance on AFF A300 and AFF A700 systems. NetApp measured the data throughput, peak IOPS, and average latency of an AFF A300 running ONTAP 9.5 and an AFF A700 storage controller running ONTAP 9.4, each running an Epic EHR workload. In a manner similar to SPC-3 testing, all inline storage efficiency features were enabled.

We ran the Epic GenIO workload generator on an AFF A300 cluster that contained a total of twenty-four 3.8TB SSDs and on an AFF A700 cluster that contained a total of forty-eight 3.8TB SSDs. We tested our cluster at a range of load points that drove the storage to peak CPU utilization. At each load point, we collected information about the storage IOPS and latency.

NetApp has consistently with each software upgrade improved performance in the range of 40-50%. Innovation with performance enhancements has varied based on workload and protocol.

The Epic performance test demonstrated that the AFF A300 cluster IOPS increased from 75,000 IOPS at <1ms to a peak performance of 188,929 IOPS at <1ms. For all load points at or below 200,000 IOPS, we were able to maintain consistent storage latencies of no greater than 1ms. Additionally, the Epic performance test demonstrated that the AFF A700 cluster IOPS increased from 75,000 IOPS at <1ms to a peak performance of 319,000 IOPS at <1ms. For all load points at or below 320,000 IOPS, we were able to maintain consistent storage latencies of no greater than 1ms.

## Test methodology

### Test plan

The GenerationIO tool (GenIO) is used by Epic to validate that storage is production ready. This test focuses on performance by pushing storage to its limits and determining the headroom on storage controllers by ramping up until requirements fail.

The tests performed here are focused on determining headroom as well as using Adaptive Quality of Service (AQOS) to protect critical Epic workloads. For AFF A300 testing, two servers are used with GenIO loaded on both to drive I/O on the storage controllers. Three servers are used with GenIO loaded on all three to drive I/O on the AFF A700 storage controllers. Three servers are used because of server performance limits, and three servers are required for an AFF A700.

### Test environment

#### Hardware and software

For this study, we configured three Red Hat Linux virtual machines (VMs) on VMware ESXi 6.5 running on Cisco UCS B200-M5s. We connected the ESXi hosts to the AFF storage controller nodes with Cisco MDS-series switches by using 16Gb FC on the server side and 16Gb FC on the storage side. The AFF A700 nodes were connected to one DS2446 disk shelf with 3.8TB SSDs by following NetApp cabling best practices.

The three tables below list the hardware and software components that we used for the Epic performance test configuration.

The following table lists Epic Test hardware and software components.

Hardware and software components	Details
Operating system for VM	RHEL 7.4 VMs
Operating system on server blades	VMware ESXi 6.5
Physical server	Cisco UCS B200 M5 x 3
Processors per server	Two 20-core Intel Xeon Gold 6148 2.4Ghz
Physical memory per server	768GB
FC network	16Gb FC with multipathing
FC HBA	FC vHBA on Cisco UCS VIC 1340
Dedicated public 1GbE ports for cluster management	Two Intel 1350GbE ports
16Gb FC switch	Cisco MDS 9148s
40GbE switch	Cisco Nexus 9332 switch

The following table lists NetApp AFF A700 and AFF A300 storage system hardware and software.

Hardware and software components	AFF A700 details	AFF A300 details
Storage system	AFF A700 controller configured as a high-availability (HA) active-active pair	AFF A300 controller configured as a high-availability (HA) active-active pair
ONTAP version	9.4	9.5
Total number of drives	36	24
Drive size	3.8TB	3.8TB
Drive type	SSD	SSD
FC target ports	Eight 16Gb ports (four per node)	Eight 16Gb ports (four per node)
Ethernet ports	Four 10Gb ports (two per node)	Four 10Gb ports (two per node)
Storage virtual machines (SVMs)	One SVM across both node aggregates	One SVM across both node aggregates
Ethernet logical interfaces (LIFs)	Four 1Gb management LIFs (two per node connected to separate private VLANs)	Four 1Gb management LIFs (two per node connected to separate private VLANs)
FC LIFs	Four 16Gb data LIFs	Four 16Gb data LIFs

The following table lists NetApp AFF A700 and AFF A300 storage system layout.

Storage layout	AFF A700 details	AFF A300 details
SVM	Single SVM for Epic application databases	Single SVM for Epic application databases
Aggregates	Two 20TB each	Two 30TB each



Storage layout	AFF A700 details	AFF A300 details
Volumes for production	Sixteen 342GB volumes per RHEL VM	Sixteen 512GB volumes per RHEL VM
LUNs for production	Sixteen 307GB LUNs, one per volume	Sixteen 460GB LUNs, one per volume
Volumes for journal	Two 95Gb volumes per RHEL VM	Two 240Gb volumes per RHEL VM
LUNs for journal	Two 75Gb LUNs, one per volume	Two 190Gb LUNs, one per volume

## Workload testing

### AFF A300 procedure

The AFF A300 HA pair can comfortably run the largest Epic instance in existence. If you have two or more very large Epic instances, you might need to use an AFF A700, based on the outcome of the NetApp SPM tool.

#### Data generation

Data inside the LUNs were generated with Epic's Dgen.pl script. The script is designed to create data similar to what would be found inside an Epic database.

The following Dgen command was run from both RHEL VMs, epic-rhel1 and epic-rhel2:

```
./dgen.pl --directory "/epic" --jobs 2 --quiet --pctfull 20
```

-pctfull is optional and defines the percentage of the LUN to fill with data. The default is 95%. The size does not affect performance, but it does affect the time to write the data to the LUNs.

After the dgen process is complete, you can run the GenIO tests for each server.

#### Run GenIO

Two servers were tested. A ramp run from 75,000 to 110,000 IOPS was executed, which represents a very large Epic environment. Both tests were run at the same time.

Run the following GenIO command from the server epic-rhel1:

```
./RampRun.pl -miniops 75000 --maxiops 110000 --background --disable-warmup
--runtime 30 --wijfile /epic/epicjrn/GENIO.WIJ --numruns 10 --system epic-
rhell --comment Ramp 75-110k
```

### GenIO result on the AFF A300

The following table lists GenIO results on the AFF A300

Read IOPs	Write IOPs	Total IOPs	Longest write cycle (sec)	Effective write latency (ms)	Randread average (ms)
142505	46442	188929	44.68	0.115	0.66

## AFF A700 procedure

For larger Epic environments, typically greater than ten million global references, customers can choose the AFF A700.

## Data generation

Data inside the LUNs was generated with Epic's Dgen.pl script. The script is designed to create data similar to what would be found inside an Epic database.

Run the following Dgen command on all three RHEL VMs.

```
./dgen.pl --directory "/epic" --jobs 2 --quiet --pctfull 20
```

`--pctfull` is optional and defines the percentage of the LUN to fill with data. The default is 95%. The size does not affect performance, but it does affect the time to write the data to the LUNs.

After the dgen process is complete you are ready to run the GenIO tests for each server.

## Run GenIO

Three servers were tested. On two servers, a ramp run from 75,000 to 100,000 IOPs was executed, which represents a very large Epic environment. The third server was set up as a bully to ramp run from 75,000 IOPS to 170,000 IOPS. All three tests were run at the same time.

Run the following GenIO command from the server epic-rhel1:

```
./RampRun.pl --miniops 75000 --maxiops 100000 --background --disable-warmup
--runtime 30 --wijfile /epic/epicjrn/GENIO.WIJ --numruns 10 --system epic-
rhell1 --comment Ramp 75-100k
```

## GenIO results on the AFF A700

The following table presents the GenIO results a test of the AFF A700.

Read IOPs	Write IOPs	Total IOPs	Longest write cycle (sec)	Effective write latency (ms)	Randread average (ms)
241,180	78,654	319,837	43.24	0.09	1.05

## Performance SLA with AQOS

NetApp can set floor and ceiling performance values for workloads using AQOS policies. The floor setting guarantees minimum performance. IOPS/TB can be applied to a group of volumes for an application like Epic.

The Epic workload assigned to a QoS policy is protected from other workloads on the same cluster. The minimum requirements are guaranteed while allowing the workload to peak and use available resources on the controller.

In this test, server 1 and server 2 were protected with AQOS, and the third server acted as a bully workload to cause performance degradation within the cluster. AQOS allowed servers 1 and 2 to perform at the specified SLA, while the bully workload showed signs of degradation with longer write cycles.

### Adaptive quality of service defaults

ONTAP comes configured with three default AQOS policies: value, performance, and extreme. The values for each policy can be view with the `qos` command. Use `-instant` at the end of the command to view all AQOS settings.

```
::> qos adaptive-policy-group show
Name           Vserver Wklds  Expected IOPS      Peak IOPS
extreme        fp-g9a  0      6144IOPS/TB 12288IOPS/TB
performance    fp-g9a  0      2048IOPS/TB 4096IOPS/TB
value          fp-g9a  0      128IOPS/TB  512IOPS/TB
```

Here is the syntax to create an AQOS policy:

```
::> qos adaptive-policy-group modify -policy-group aqos- epic-prod1
-expected-iops 5000 -peak-iops 10000 -absolute-min-iops 4000 -peak-iops
-allocation used-space
```

There are a few important settings in an AQOS policy:

- **Expected IOPS.** This adaptive setting is the minimum IOPS/TB value for the policy. Workloads are guaranteed to get at least this level of IOPS/TB. This is the most important setting in this testing. In our example test, the performance AQOS policy was set to 2048IOPS/TB.
- **Peak IOPS.** This adaptive setting is the maximum IOPS/TB value for the policy. In our example test, the performance AQOS policy was set to 4096IOPS/TB.
- **Peak IOPS allocation.** Options are allocated space or used space. Set this parameter to used space, because this value changes as the database grows in the LUNs.
- **Absolute minimum IOPS.** This setting is static and not adaptive. This parameter sets the minimum IOPS regardless of size. This value is only used when size is less than 1TB and has no effect on this testing.

Typically, Epic workloads in production run at about ~1000 IOPS/TB of storage and capacity, and IOPS grows linearly. The default AQOS performance profile is more than adequate for an Epic workload.

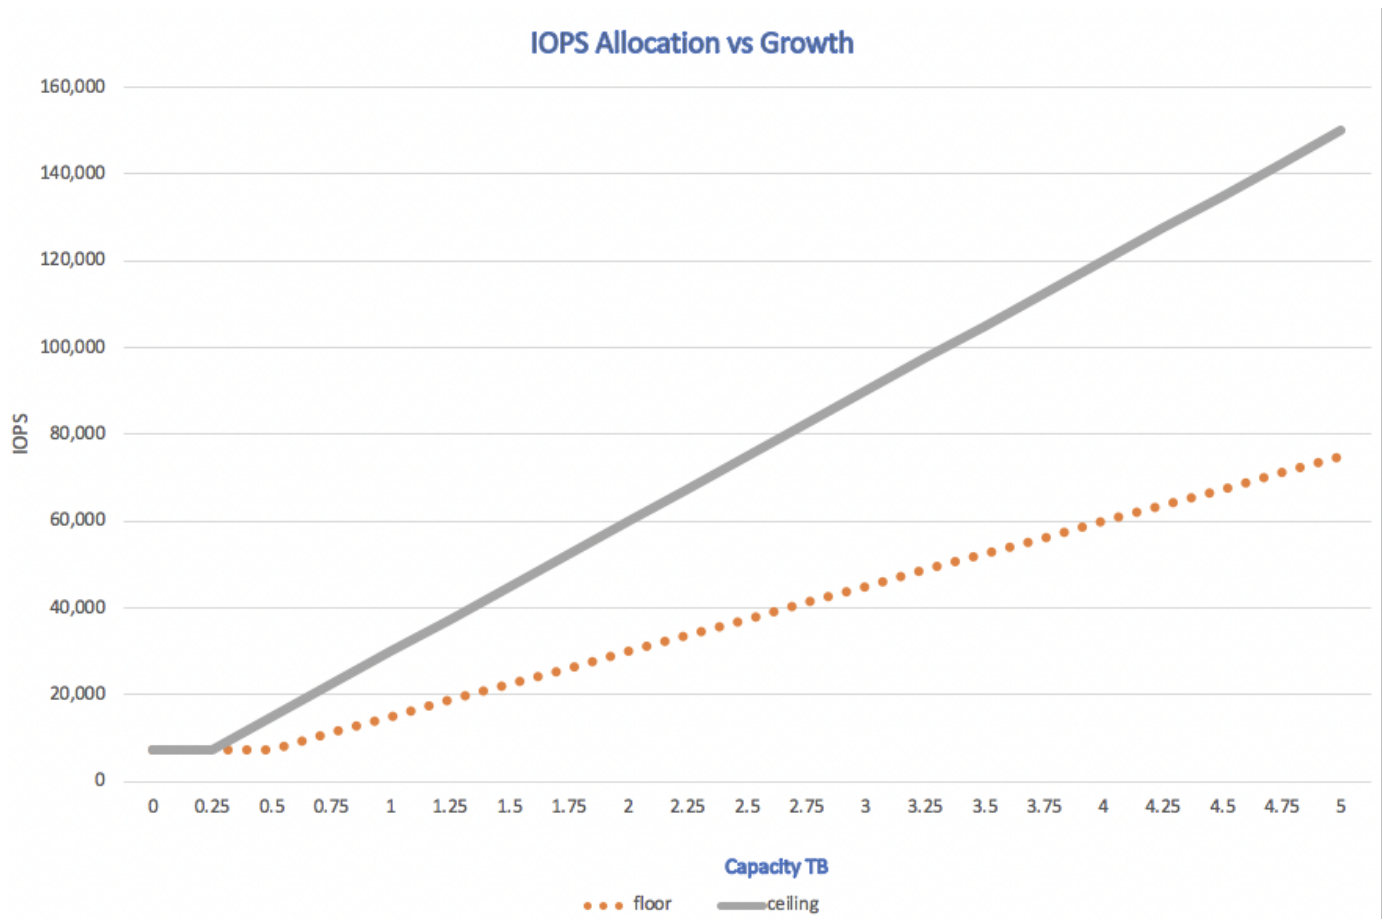
For this testing the lab did not reflect a production size database with a smaller size of 5TB. The goal was to run each test at 75,000 IOPS. The setting for the EpicProd AQOS policy is shown below.

- Expected IOPS/TB = Total IOPS/used space
- 15,000 IOPS/TB = 75,000 IOPS/5TB

The following table presents the settings that were used for the EpicProd AQOS policy.

Setting	Value
Volume size	5TB
Required IOPS	75,000
peak-iops-allocation	Used space
Absolute minimum IOPS	7,500
Expected IOPS/TB	15,000
Peak IOPS/TB	30,000

The following figure shows how floor IOPS and ceiling IOPS are calculated as the used space grows over time.

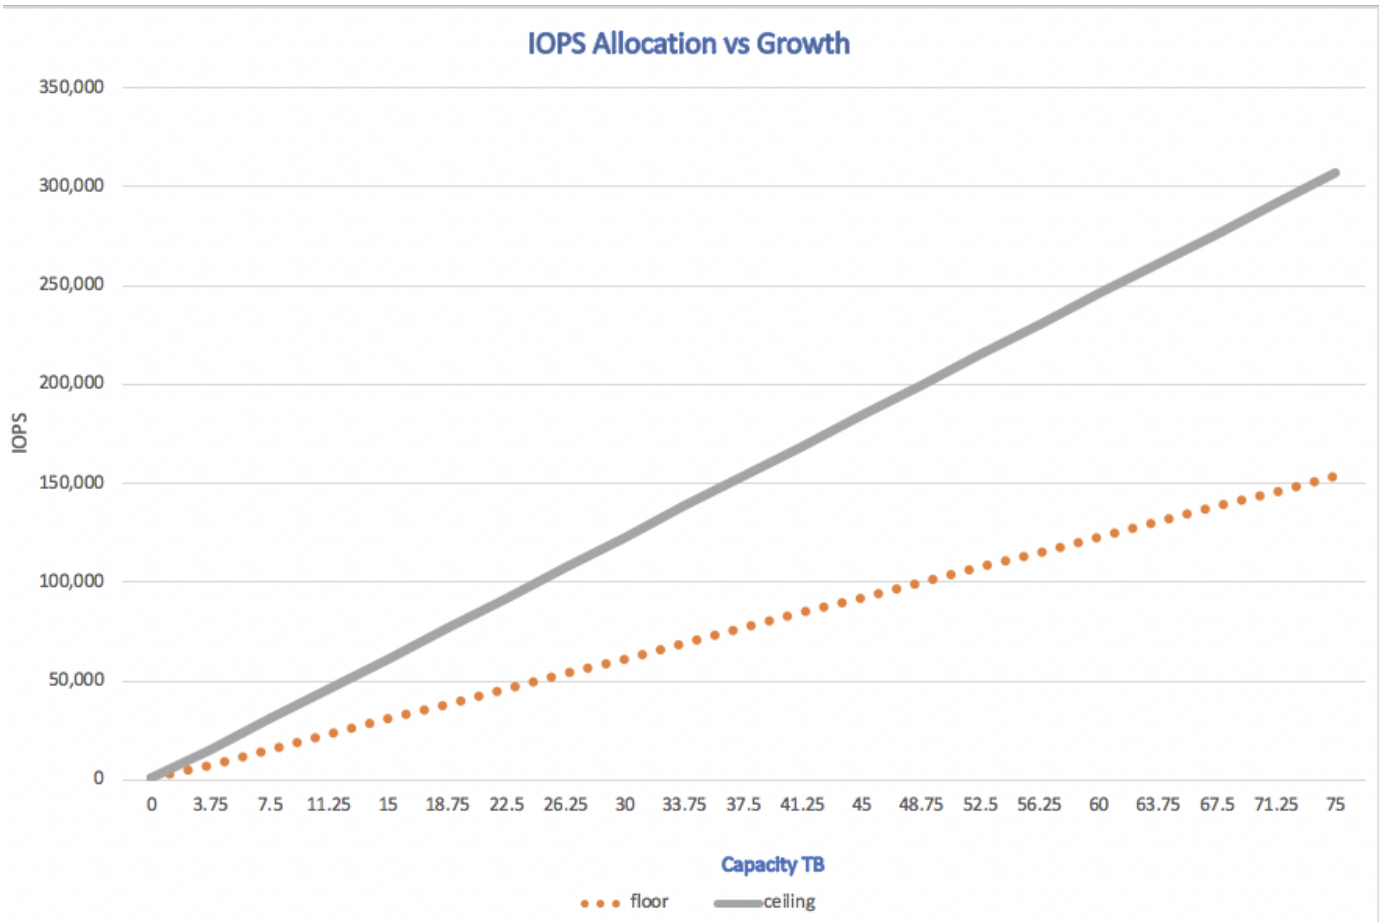


For a production-sized database, you can either create a custom AQOS profile like the one used in the last example, or you can use the default performance AQOS policy. The settings for the performance AQOS policy are show in the table below.

Setting	Value
Volume size	75TB
Required IOPS	75,000
peak-iops-allocation	Used space
Absolute minimum IOPS	500

Setting	Value
Expected IOPS/TB	1,000
Peak IOPS/TB	2,000

The following figure shows how floor and ceiling IOPS are calculated as the used space grows over time for the default performance AQOS policy.



## Parameters

- The following parameter specifies the name of the adaptive policy group:

```
-policy-group <text> - Name
```

Adaptive policy group names must be unique and are restricted to 127 alphanumeric characters including underscores "\_" and hyphens "-". Adaptive policy group names must start with an alphanumeric character. Use the `qos adaptive-policy-group rename` command to change the adaptive policy group name.

- The following parameter specifies the data SVM (called vserver in the command line) to which this adaptive policy group belongs.

```
-vserver <vserver name> - Vserver
```

You can apply this adaptive policy group to only the storage objects contained in the specified SVM. If the system has only one SVM, then the command uses that SVM by default.

- The following parameter specifies the minimum expected IOPS/TB or IOPS/GB allocated based on the storage object allocated size.

```
-expected-iops {<integer>[IOPS[/{GB|TB}]] (default: TB)} - Expected IOPS
```

- The following parameter specifies the maximum possible IOPS/TB or IOPS/GB allocated based on the storage object allocated size or the storage object used size.

```
-peak-iops {<integer>[IOPS[/{GB|TB}]] (default: TB)} - Peak IOPS
```

- The following parameter specifies the absolute minimum IOPS that is used as an override when the expected IOPS is less than this value.

```
[-absolute-min-iops <qos_tput>] - Absolute Minimum IOPS
```

The default value is computed as follows:

```
qos adaptive-policy-group modify -policy-group aqos- epic-prod1  
-expected-iops 5000 -peak-iops 10000 -absolute-min-iops 4000 -peak-iops  
-allocation used-space
```

```
qos adaptive-policy-group modify -policy-group aqos- epic-prod2  
-expected-iops 6000 -peak-iops 20000 -absolute-min-iops 5000 -peak-iops  
-allocation used-space
```

```
qos adaptive-policy-group modify -policy-group aqos- epic-bully  
-expected-iops 3000 -peak-iops 2000 -absolute-min-iops 2000 -peak-iops  
-allocation used-space
```

## Data generation

Data inside the LUNs was generated with the Epic `Dgen.pl` script. The script is designed to create data similar to what would be found inside an Epic database.

The following `Dgen` command was run on all three RHEL VMs:

```
./dgen.pl --directory "/epic" --jobs 2 --quiet --pctfull 20
```

## Run GenIO

Three servers were tested. Two ran at a constant 75,000 IOPS, which represents a very large Epic environment. The third server was setup as a bully to ramp run from 75,000 IOPS to 150,000 IOPS. All three tests were run at the same time.

### Server epic\_rhel1 GenIO test

The following command was run to assign EpicProd AQOS settings to each volume:

```
::> vol modify -vserver epic -volume epic_rhell_* -qos-adaptive-policy  
-group AqosEpicProd
```

The following GenIO command was run from the server epic-rhel1:

```
./RampRun.pl -miniops 75000 --maxiops 75000 --background --disable-warmup  
--runtime 30 --wijfile /epic/GENIO.WIJ --numruns 10 --system epic-rhel1  
--comment Ramp constant 75k
```

### Server epic\_rhel2 GenIO test

The following command was run to assign EpicProd AQOS settings to each volume:

```
::> vol modify -vserver epic -volume epic_rhel2_* -qos-adaptive-policy  
-group AqosEpicProd
```

The following GenIO command was run from the server epic-rhel2:

```
./RampRun.pl --miniops 75000 --maxiops 75000 --background --disable-warmup  
--runtime 30 --wijfile /epic/GENIO.WIJ --numruns 10 --system epic-rhel2  
--comment Ramp constant 75k
```

### Server epic\_rhel3 GenIO test (bully)

The following command assigns no AQOS policy to each volume:

```
::> vol modify -vserver epic -volume epic_rhel3_* -qos-adaptive-policy  
-group non
```

The following GenIO command was run from the server epic-rhel3:

```
./RampRun.pl --miniops 75000 --maxiops 150000 --background --disable
-warmup --runtime 30 --wijfile /epic/GENIO.WIJ --numruns 10 --system epic-
rhel3 --comment Ramp 75-150k
```

## AQOS test results

The tables in the following sections contain the output from the summary.csv files from each concurrent GenIO test. To pass the test, the longest write cycle must have been below 45 seconds. The effective write latency must have been below 1 millisecond.

### Server epic\_rhel1 GenIO results

The following table illustrates GenIO results for AQOS server epic\_rhel1.

Run	Read IOPS	Write IOPS	Total IOPS	Longest write cycle (sec)	Effective write latency (ms)
10	55655	18176	73832	32.66	0.12
11	55653	18114	73768	34.66	0.1
12	55623	18099	73722	35.17	0.1
13	55646	18093	73740	35.16	0.1
14	55643	18082	73726	35.66	0.1
15	55634	18156	73791	32.54	0.1
16	55629	18138	73767	34.74	0.11
17	55646	18131	73777	35.81	0.11
18	55639	18136	73775	35.48	0.11
19	55597	18141	73739	35.42	0.11

### Server epic\_rhel2 GenIO results

The following table illustrates GenIO results for AQOS server epic\_rhel2.

Run	Read IOPS	Write IOPS	Total IOPS	Longest write cycle (sec)	Effective write latency (ms)
10	55629	18081	73711	33.96	0.1
11	55635	18152	73788	28.59	0.09
12	55606	18154	73761	30.44	0.09
13	55639	18148	73787	30.37	0.09
14	55629	18145	73774	30.13	0.09
15	55619	18125	73745	30.03	0.09
16	55640	18156	73796	33.48	0.09



Run	Read IOPS	Write IOPS	Total IOPS	Longest write cycle (sec)	Effective write latency (ms)
17	55613	18177	73790	33.32	0.09
18	55605	18173	73779	32.11	0.09
19	55606	18178	73785	33.19	0.09

### Server epic\_rhel3 GenIO results (bully)

The following table illustrates GenIO results for AQOS server epic\_rhel3.

Run	Write IOPS	Total IOPS	Longest WIJ Time (sec)	Longest Write Cycle (sec)	Effective Write Latency (ms)
10	19980	81207	21.48	40.05	0.1
11	21835	88610	17.57	46.32	0.12
12	23657	95955	19.77	53.03	0.12
13	25493	103387	21.93	57.53	0.12
14	27331	110766	23.17	60.57	0.12
15	28893	117906	26.93	56.56	0.1
16	30704	125233	28.05	60.5	0.12
17	32521	132585	28.43	64.38	0.12
18	34335	139881	30	70.38	0.12
19	36361	147633	22.78	73.66	0.13

### AQOS test results analysis

The results from the previous section demonstrate that the performance of the servers epic\_rhel1 and epic\_rhel2 are not affected by the bully workload on epic\_rhel3. epic\_rhel3 ramps up to 150,000 IOPS and starts to fail the GenIO test as it hits the limits of the controllers. The write cycle and latency on epic\_rhel1 and epic\_rhel2 stay constant while the bully server spirals out of control.

This illustrates how an AQOS minimum policy can effectively isolate workloads from bullies and guarantee a minimum level of performance.

AQOS has a number of benefits:

- It allows for a more flexible and simplified architecture. Critical workloads no longer need to be siloed and can coexist with noncritical workloads. All capacity and performance can be managed and allocated with software rather than by using physical separation.
- It saves on the amount of disk and controllers required for Epic running on an ONTAP cluster.
- It simplifies the provisioning of workloads to performance policies that guarantee consistent performance.
- Optionally, you can also implement of NetApp Service Level Manager to perform the following tasks:
  - Create a catalog of services to simplify provisioning of storage.
  - Deliver predictable service levels so that you can consistently meet utilization goals.

- Define service-level objectives.

## Conclusion

By 2020, all Epic customers must be on flash storage. NetApp ONTAP was the first all- flash array to get a high-comfort rating from Epic, and it is listed under Enterprise Storage Arrays. All NetApp platforms that run a GA version of ONTAP are high comfort.

Epic requires that critical workloads like Production, Report, and Clarity are physically separated on storage allocations called pools. NetApp provides multiple pools of storage in a single cluster with each node and offers a simplified single cluster and single OS for the entire Epic solution. ONTAP supports all protocols for NAS and SAN, with mixed tiers of storage for SSD, HDD, and cloud.

The introduction of Adaptive QoS in ONTAP 9.3, with significant enhancements in ONTAP 9.4, allows for the creation of storage pools with software without the need for physical separation. This capability greatly simplifies architecture development, permits the consolidation of nodes and disks, and improves performance for critical workloads like production by spreading across nodes. It also eliminates storage performance issues caused by bullies and guarantees consistent performance for the life of the workload.

## Where to find additional information

To learn more about the information that is described in this document, see the following documents or websites:

### FlexPod Design Zone

- NetApp FlexPod Design Zone

<https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/flexpod-design-guides.html>

- FlexPod DC with FC Storage (MDS Switches) Using NetApp AFF, vSphere 6.5U1, and Cisco UCS Manager

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_esxi65u1\\_n9fc.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html)

- Cisco Best Practices with Epic on Cisco UCS

[https://www.cisco.com/c/dam/en\\_us/solutions/industries/healthcare/Epic\\_on\\_UCS\\_tech\\_brief\\_FNL.pdf](https://www.cisco.com/c/dam/en_us/solutions/industries/healthcare/Epic_on_UCS_tech_brief_FNL.pdf)

### NetApp technical reports

- TR-4693: FlexPod Datacenter for Epic EHR Deployment Guide

<https://www.netapp.com/us/media/tr-4693.pdf>

- TR-4707: FlexPod for Epic Directional Sizing Guide

<https://www.netapp.com/us/media/tr-4707.pdf>

- TR-3929: Reallocate Best Practices Guide

<https://www.netapp.com/us/media/tr-3929.pdf>

- TR-3987: Snap Creator Framework Plug-In for InterSystems Caché

<https://www.netapp.com/us/media/tr-3987.pdf>

- TR-3928: NetApp Best Practices for Epic

<https://www.netapp.com/us/media/tr-3928.pdf>

- TR-4017: FC SAN Best Practices

<https://www.netapp.com/us/media/tr-4017.pdf>

- TR-3446: SnapMirror Async Overview and Best Practices Guide

<https://www.netapp.com/us/media/tr-3446.pdf>

## ONTAP documentation

- NetApp product documentation

<https://www.netapp.com/us/documentation/index.aspx>

- Virtual Storage Console (VSC) for vSphere documentation

<https://mysupport.netapp.com/documentation/productlibrary/index.html?productID=30048>

- ONTAP 9 Documentation Center

<http://docs.netapp.com/ontap-9/index.jsp>

## Cisco Nexus, MDS, Cisco UCS, and Cisco UCS Manager guides

- Cisco UCS Servers Overview

<https://www.cisco.com/c/en/us/products/servers-unified-computing/index.html>

- Cisco UCS Blade Servers Overview

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>

- Cisco UCS B200 M5 Datasheet

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>

- Cisco UCS Manager Overview

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-manager/index.html>

- Cisco UCS Manager 3.2(3a) Infrastructure Bundle (requires Cisco.com authorization)

<https://software.cisco.com/download/home/283612660/type/283655658/release/3.2%25283a%2529>

- Cisco Nexus 9300 Platform Switches

<https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78->

[736967.html](#)

- Cisco MDS 9148S FC Switch

<https://www.cisco.com/c/en/us/products/storage-networking/mds-9148s-16g-multilayer-fabric-switch/index.html>

## Acknowledgements

The following people contributed to the creation of this guide:

- Ganesh Kamath, Technical Marketing Engineer, NetApp
- Atul Bhalodia, Technical Marketing Engineer, NetApp
- Brandon Agee, Technical Marketing Engineer, NetApp
- Brian O'Mahony, Solution Architect – Healthcare, NetApp
- Ketan Mota, Product Manager, NetApp
- Jon Ebmeier, Technical Solutions Architect, Cisco Systems, Inc
- Mike Brennan, Product Manager, Cisco Systems, Inc

## Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system-without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.