Technical Report

# NetApp AFF A700 Performance with Oracle Database

Rodrigo Nascimento and Scott Lane, NetApp
March 2017 | TR-4582

## Abstract

This report provides a performance summary for NetApp® and partner systems engineers who are interested in assessing Oracle database performance with a NetApp AFF A700 storage system.

NetApp All Flash FAS (AFF) systems combine the extreme performance capability of flash media with NetApp ONTAP® 9.1 software to provide performance acceleration, operational agility, best-in-class data protection, and business continuity for database deployments.

NetApp

TABLE OF CONTENTS



LIST OF TABLES



LIST OF FIGURES



 

# 1 Introduction

The NetApp solution for Oracle databases delivers industry-leading storage, superior scalability, continuous data access, and automated data management for immediate responses to business opportunities. NetApp has worked with Oracle for years to develop innovative integrated solutions that reduce IT and business costs and complexity. NetApp leads the way with data storage, offering compelling solutions for Oracle databases.

This document describes the performance of the NetApp AFF A700 storage system with an Oracle database workload.

## 1.1 ONTAP Empowers All Flash FAS Performance

NetApp AFF systems meet 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 on the efficiency, reliability, or flexibility of your IT operations. As 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 for your data center.

Designed specifically for flash, the AFF A series all-flash systems deliver industry-leading performance, capacity, density, scalability, security, and network connectivity in dense form factors. 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 7 million IOPS per cluster with submillisecond 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 to provide both 40Gb Ethernet (40GbE) and 32Gb Fibre Channel (FC) connectivity. Therefore, they eliminate the bandwidth bottlenecks that are increasingly moved to the network from storage as flash becomes faster and faster.

NetApp has taken the lead for the 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 FlashEssentials 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 FlashEssentials 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 that are based on customer deployment have shown that total 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%.
- A data fabric enables live workload migration between flash and hard-disk-drive tiers, on the premises or to the cloud.

 

- Quality-of-service capability safeguards service-level objectives in multiworkload and multitenant environments.

### 1.2 NetApp ONTAP 9.1

NetApp ONTAP 9.1 is a major advance in the industry's leading enterprise data management software. This software can integrate the best of next-generation and traditional technologies, incorporating flash, cloud, and software-defined architectures while building a foundation for the Data Fabric. In addition, new customers and existing NetApp Data ONTAP® 8.3 environments can quickly and easily use the rich data services that ONTAP 9.1 delivers.

An essential feature for Oracle databases that are deployed on shared enterprise storage is the capability to deliver consistent and dependable high performance. High performance must be coupled with nondisruptive operations, high availability, scalability, and storage efficiency. You can depend on ONTAP 9.1 and AFF to provide these essential elements.

Built on the ONTAP unified scale-out architecture, AFF consistently meets or exceeds the high-performance demands of Oracle databases. AFF also provides rich data management capabilities, such as integrated data protection and nondisruptive upgrades and data migration. These features help you eliminate performance silos and seamlessly integrate AFF into a shared infrastructure.

ONTAP 9.1 delivers enhanced inline deduplication, inline compression, and inline data compaction capabilities that significantly reduce the amount of flash storage that is required, with no impact on system performance. ONTAP 9.1 also provides industry-leading ecosystem integration with database applications that makes administration of databases and storage systems far more efficient than with other flash storage solutions on the market.

## 2 Executive Summary

NetApp performed this study to showcase the storage performance and the benefits of the AFF A700 running ONTAP 9.1. NetApp measured the Oracle server read latency and the data throughput and peak sustained IOPS of the AFF A700 storage controller running ONTAP 9.1 with an Oracle OLTP workload. All inline storage efficiency features were enabled.

We ran an OLTP workload called *SLOB2* on a two-node AFF A700 cluster that contained a total of twenty-four 3.8TB SSDs. We tested our cluster at a range of load points that drove the storage to peak utilization. At each load point, we collected information about the storage IOPS and server latency.

In the Oracle Database 12*c* performance test, the AFF A700 generated 16,300 IOPS at 200µs and had a peak performance of 563,900 IOPS at 1392µs. For all load points up to 533,700 IOPS, the test configuration maintained consistent server read latencies that were below 1ms.

In addition to Oracle performance testing on the AFF A700, we also conducted the same testing against a similarly configured AFF A700s. We found that the AFF A700s had performance that was within 3% to 4% of the AFF A700 at each load point.

 

# 3 Measuring Storage Performance

NetApp conducted a study to measure the performance of an AFF A700 storage system running ONTAP 9.1. The following sections describe the methodology and the design considerations that we used to test the AFF A700 running a standard Oracle workload.

## 3.1 Test Methodology

For this study, we used the SLOB2 load-generation tool to simulate an OLTP workload against the Oracle Database 12*c* test configuration. The workload generated a select-update ratio of approximately 75:25 against the Oracle database in the test configuration.

We created an Oracle Real Application Clusters (RAC) environment with one database connected through FC to the AFF A700 system. By using six Oracle database servers and the OLTP load generator, we measured the peak performance of the storage system by generating a workload that was designed to maximize the storage system utilization. We then reran the test while ramping up the workload from 32 users to 440 users. This approach allowed us to gather performance metrics at a range of different load points.

## 3.2 Hardware and Software

We configured six Oracle Database 12*c* servers on six Fujitsu PRIMERGY RX300 S7 servers. We connected the six servers to a single two-node AFF A700 controller through 16Gb FC. The AFF A700 HA pair contained twenty-four 3.8TB internal SSDs.

Table 1 and Table 2 list the hardware and software components that we used for the AFF A700 Oracle performance test configuration.

In addition, we conducted the SLOB performance test against an AFF A700s storage controller. The AFF A700s is a smaller and more compact version of the AFF A700. The AFF A700s has the same specifications that are listed in Table 1 and Table 2. However, we used 24 SSDs with a smaller capacity (960GB SSDs) and an Emulex LPe32000 for the FC host bus adapter (HBA) in our test environment for the AFF A700s. All other test configurations were identical.

**Table 1) Oracle hardware and software components.**

| Hardware and Software Components | Details |
|---|---|
| Oracle Database 12*c* servers | Six Fujitsu PRIMERGY RX300 S7 servers |
| Server operating system | Red Hat Enterprise Linux 6.6 |
| Oracle database version | 12*c* (RAC) |
| Processors per server | Two 6-core Intel Xeon E5-2630 @ 2.30GHz |
| Physical memory per server | 128GB |
| FC network | 8Gb FC with multipathing |
| FC HBA | QLogic QLE2562 dual-port PCIe FC HBA |
| Dedicated public 1GbE ports for cluster management | Two Intel 1350GbE ports |
| 8Gb FC switch | Brocade 6510 24-port switch |
| 10GbE switch | Cisco Nexus 5596 switch |

 

**Table 2) NetApp storage system hardware and software.**

| Hardware and Software Components | Details |
|---|---|
| Storage system | AFF A700 controller, configured as a high-availability (HA) active-active pair |
| ONTAP version | 9.1 |
| Total number of drives | 24 |
| Drive size | 3.8TB |
| Drive type | SSD |
| FC target ports | Eight 16Gb ports (four per node) |
| Ethernet ports | Four 10Gb ports (two per node) |
| Storage virtual machines (SVMs) | One across both node aggregates |
| Ethernet logical interfaces (LIFs) | Four 1Gb management LIFs (two per node connected to separate private VLANs) |
| FC LIFs | Eight 16Gb data LIFs |

## 3.3 Network Design

This section provides the network connectivity details for the tested configurations. The network diagram in Figure 1 shows that the FC SAN was deployed with a Brocade 6510 16Gb FC switch. Each storage node had four ports connected to the FC switch, and each server had two ports connected to the switch. The multiple ports that we used in the FC SAN configurations provided high availability through multipathing and increased bandwidth. At no point in the testing did the network connectivity create a bottleneck.

**Figure 1) Network design.**

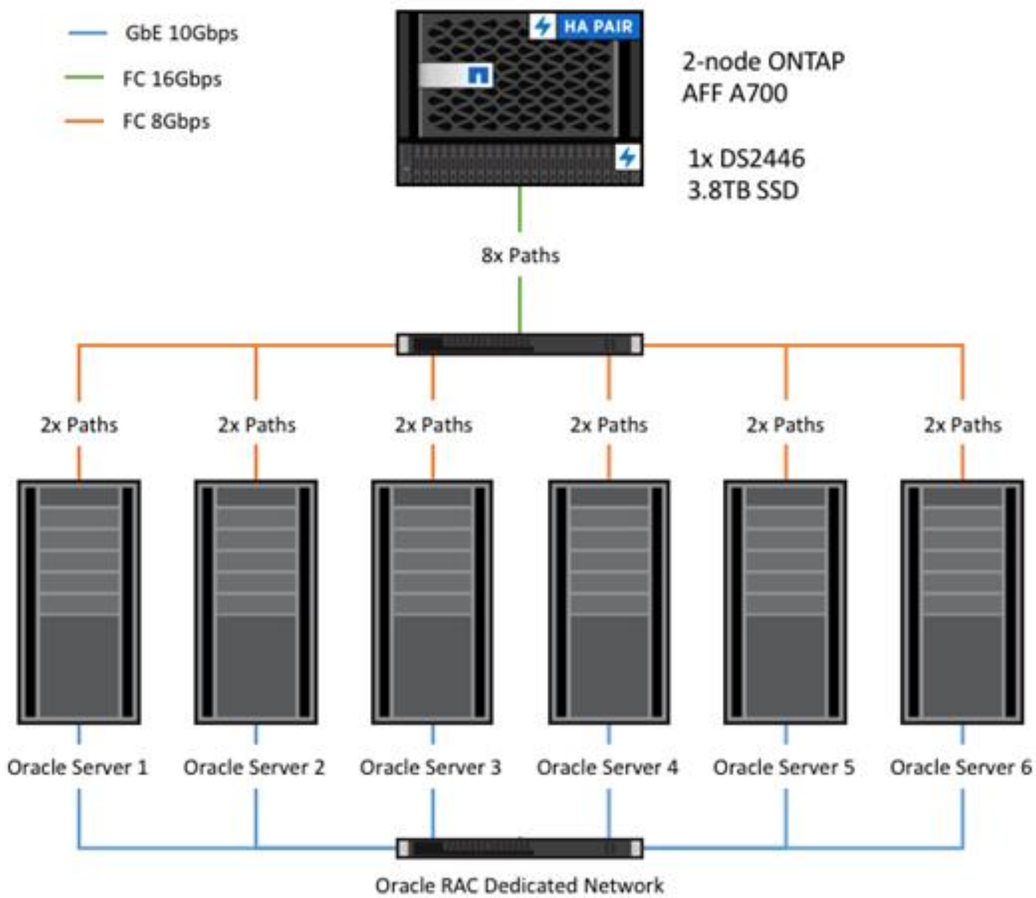


Note: Oracle Server 1 is also the workload generator.

## 3.4 Database Layout and Storage Provisioning Design

Table 3 summarizes the layout for the Oracle database.

In the performance test, we used one Oracle RAC database to host the simulated OLTP environment. The storage system had twenty-four 3.8TB internal SSDs attached in an external shelf. Of those 24 drives, 23 drives were used to create one single NetApp RAID DP® aggregate and 1 was left as a spare drive. We configured the two data aggregates into a single SVM and created a single database by using Oracle RAC.

**Table 3) Database layout.**

| Storage | Aggregate Name | Volume Name | LUN Size (GB) | Volume Size (GB) | Description |
|---|---|---|---|---|---|
| Per node | n1_aggr0 | – | – | – | Used ADP<br>Total aggregate size = 970GB |
| | n2_aggr0 | – | – | – | Used ADP<br>Total aggregate size = 970GB |
| | n1_aggr1 | – | – | – | 21 data + 2 parity RAID DP + 1 spare<br>Total aggregate size = 31TB |
| | n2_aggr1 | – | – | – | 21 data + 2 parity RAID DP + 1 spare<br>Total aggregate size = 31TB |
| Oracle RAC configuration | n1_aggr1 | dbFCP_vol1 | 128 | 128 | Data files |
| | n2_aggr1 | dbFCP_vol2 | 128 | 128 | |
| | n1_aggr1 | dbFCP_vol3 | 128 | 128 | |
| | n2_aggr1 | dbFCP_vol4 | 128 | 128 | |
| | n1_aggr1 | dbFCP_vol5 | 128 | 128 | |
| | n2_aggr1 | dbFCP_vol6 | 128 | 128 | |
| | n1_aggr1 | dbFCP_vol7 | 128 | 128 | |
| | n2_aggr1 | dbFCP_vol8 | 128 | 128 | |
| | n1_aggr1 | dbFCP_vol9 | 128 | 128 | |
| | n2_aggr1 | dbFCP_vol10 | 128 | 128 | |
| | n1_aggr1 | dbFCP_vol11 | 128 | 128 | |
| | n2_aggr1 | dbFCP_vol12 | 128 | 128 | |
| | n1_aggr1 | dbFCP_vol13 | 128 | 128 | |
| | n2_aggr1 | dbFCP_vol14 | 128 | 128 | |
| | n1_aggr1 | dbFCP_vol15 | 128 | 128 | |
| | n2_aggr1 | dbFCP_vol16 | 128 | 128 | |
| | n1_aggr1 | dbFCP_vol17 | 100 | 100 | Redo logs |
| | n2_aggr1 | dbFCP_vol18 | 100 | 100 | – |
| | n1_aggr1 | dbFCP_vol19 | 20 | 20 | Grid: CRS and Voting |

 

| Storage | Aggregate Name | Volume Name | LUN Size (GB) | Volume Size (GB) | Description |
|---|---|---|---|---|---|
| | `n2_aggr1` | `dbFCP_vol20` | 20 | 20 | |

We used one igroup per server to contain the FC initiators. We then created disk groups with an allocation unit size of 64MB by using the Oracle Automatic Storage Management (ASM) volume manager. Those ASM disk groups provided the storage that was required to create the tablespaces. The FC SAN was configured on the Brocade switch. ONTAP provided asymmetric logical unit access (ALUA) communication to the initiators so that direct paths were used for host I/O access according to the multipathing I/O load-balancing policies on the host.

We deployed zoning in our configuration to balance the FC connections by using eight paths per LUN. We used two HBA ports per server and four LIFs per node. One server port was zoned for two LIFs per node; the other port was zoned for the other two LIFs per node.

## 3.5 Workload Design

We used SLOB2 to generate our OLTP workload. Each database server applied the workload to Oracle database, log, and temp files. We configured the workload to be 75% selects and 25% updates with a block size of 8KB. We configured the SLOB2 scale parameter with 1,013,760 rows per user schema, which is equivalent to 7.7GB of SLOB2 data per user. To create our 2.1TB SLOB2 database, we used 256 users.

To collect performance results, we increased the number of Oracle users in SLOB2 from a minimum of 32 users to a maximum of 440 users by using the SLOB2 `thread_per_schema` parameter. At each load point, we verified that the storage system and the Oracle servers could maintain steady-state behavior without failure. We also made sure that there were no bottlenecks across servers or networking systems.

**Note:** We took care in these test steps to simulate real database and customer workloads. However, we acknowledge that workloads can vary across databases. In addition, these test results were obtained in a closed lab environment with no competing workloads on the same infrastructure. In a typical shared-storage infrastructure, other workloads share resources. Therefore, your results might vary from the results that are found in this report.

## 3.6 Performance Test Results

We measured the performance of our Oracle database implementation running on ONTAP 9.1. Figure 2 depicts the performance results that we collected from testing the AFF8080A EX running clustered Data ONTAP 8.3.1, the AFF8080 EX running ONTAP 9.1, and the AFF A700 running ONTAP 9.1.

In Figure 2, Oracle performance test results on the AFF A700 running ONTAP 9.1 demonstrate significantly improved performance over the other test configurations. We found that the AFF A700 controller generated 16,300 IOPS at 200µs and had a peak performance of 563,900 IOPS at 1392µs. For all load points up to 533,700 IOPS, server read latencies stayed consistently below 1ms. As shown in Figure 2, ONTAP 9.1 provides a 73% increase in IOPS under a millisecond over the results for clustered Data ONTAP 8.3.1.

 

Figure 2) All Flash FAS Oracle Database 12*c* performance.

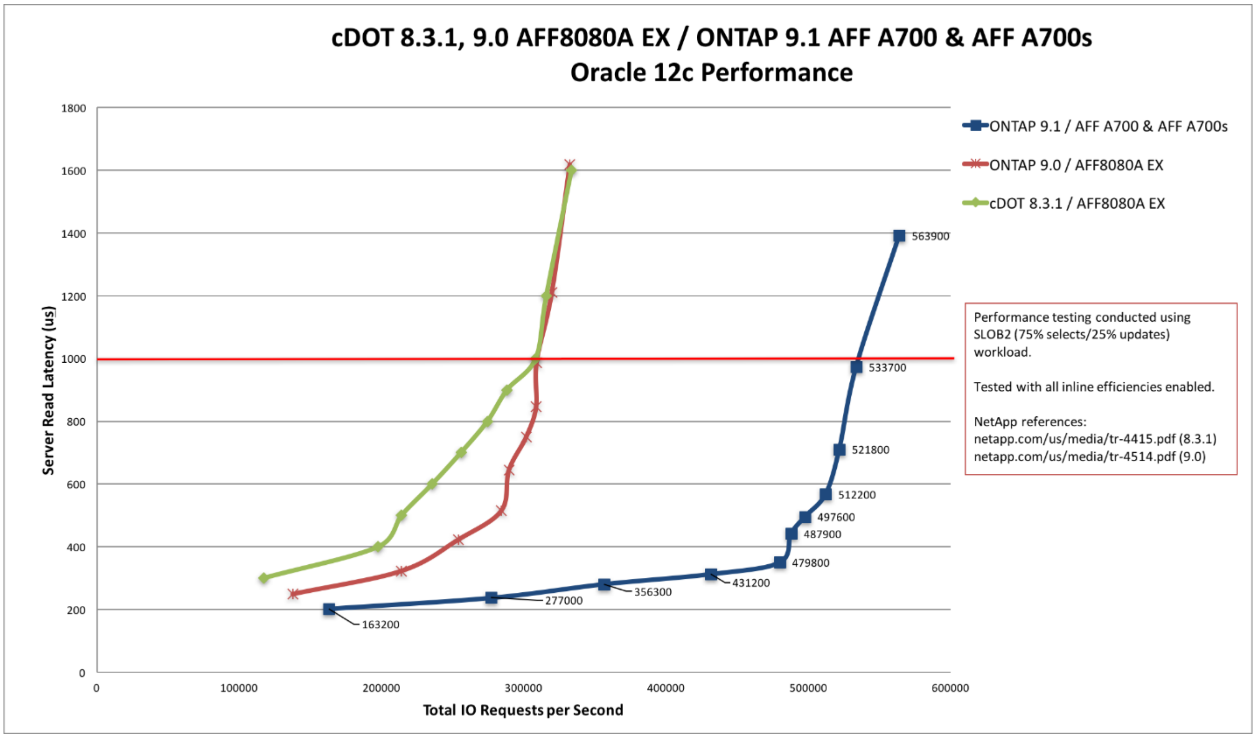


We collected an Oracle Automatic Workload Repository (AWR) report at the 512K IOPS performance point. The "Database Summary" section of the AWR shows the elapsed time and the database time for that specific performance point. The "Top Timed Events" section shows the top 10 events and their respective latencies. The "System Statistics - Per Second" section shows the number of physical reads and writes per second as well as the number of redo log operations in kilobytes per second. For details from the AWR report, see the Appendix.

In addition, we conducted a performance test with the same test configuration against the similarly configured AFF A700s. We found that the A700s had similar performance characteristics, performing within 3% to 4% of the AFF A700 at each load point.

For more information about the performance test methodology and results for the AFF8080A EX running ONTAP 9.1, see NetApp TR-4514: NetApp AFF8080A EX Storage Efficiency and Performance with Oracle Database.

# 4 Conclusion

A NetApp AFF A700 running ONTAP 9.1 software provides very high IOPS at consistent low latencies when serving an Oracle Database 12*c* OLTP workload. Our tests showed that the AFF A700 cluster produced 16,300 IOPS at 200µs and had a peak performance of 563,900 IOPS at 1392µs. For all load points up to 533,700 IOPS, we observed consistent server read latencies below 1ms on customer production data. Performance testing that used an identical test methodology with an AFF A700s resulted in performance within 3% to 4% of the AFF A700 at each load point.

 

## Appendix: AWR Report

The following three screenshots show the AWR report that we collected at the 512K IOPS point of the NetApp AFF A700s performance test.

### System Statistics - Per Second

| I# | Logical Reads/s | Physical Reads/s | Physical Writes/s | Redo Size (k)/s | Block Changes/s | User Calls/s | Execs/s | Parses/s | Logons/s | Txns/s |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | 72,292.91 | 67,518.76 | 17,519.62 | 14,213.02 | 34,743.64 | 4.50 | 1,070.41 | 5.34 | 0.58 | 265.58 |
| 2 | 72,489.39 | 67,732.63 | 17,574.33 | 14,256.62 | 34,842.91 | 4.50 | 1,074.37 | 5.73 | 0.58 | 266.38 |
| 3 | 73,287.83 | 68,479.26 | 17,763.71 | 14,412.82 | 35,223.23 | 4.50 | 1,086.26 | 5.83 | 0.58 | 269.29 |
| 4 | 73,316.86 | 68,446.73 | 17,746.89 | 14,408.89 | 35,222.36 | 5.00 | 1,086.65 | 6.22 | 0.58 | 269.19 |
| 5 | 71,896.22 | 67,101.57 | 17,406.79 | 14,127.26 | 34,538.77 | 4.50 | 1,065.49 | 5.93 | 0.58 | 263.94 |
| 6 | 71,822.74 | 67,104.65 | 17,410.29 | 14,123.16 | 34,518.94 | 4.49 | 1,064.96 | 5.80 | 0.58 | 263.91 |
| Sum | 435,105.95 | 406,383.61 | 105,421.63 | 85,541.77 | 209,089.85 | 27.50 | 6,448.14 | 34.84 | 3.47 | 1,598.30 |
| Avg | 72,517.66 | 67,730.60 | 17,570.27 | 14,256.96 | 34,848.31 | 4.58 | 1,074.69 | 5.81 | 0.58 | 266.38 |
| Std | 656.09 | 617.37 | 157.17 | 129.64 | 314.88 | 0.21 | 9.74 | 0.28 | 0.00 | 2.41 |

## WORKLOAD REPOSITORY REPORT (RAC)

**Database Summary**

| Database | | | | Snapshot Ids | | Number of Instances | | Number of Hosts | | Report Total (minutes) | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Id | Name | RAC | Block Size | Begin | End | In Report | Total | In Report | Total | DB time | Elapsed time |
| 499148728 | LWNFCP | YES | 8192 | 40 | 41 | 6 | 6 | 6 | 6 | 13,212.83 | 30.10 |

**Database Instances Included In Report**

- Listed in order of instance number, I#

| I# | Instance | Host | Startup | Begin Snap Time | End Snap Time | Release | Elapsed Time(min) | DB time(min) | Up Time(hrs) | Avg Active Sessions | Platform |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | LWNFCP1 | s144.netapp.lab.local | 14-Dec-16 12:03 | 14-Dec-16 19:43 | 14-Dec-16 20:13 | 12.1.0.2.0 | 30.10 | 2,222.12 | 8.17 | 73.82 | Linux x86 64-bit |
| 2 | LWNFCP2 | s149.netapp.lab.local | 14-Dec-16 12:03 | 14-Dec-16 19:43 | 14-Dec-16 20:13 | 12.1.0.2.0 | 30.10 | 2,222.12 | 8.17 | 73.82 | Linux x86 64-bit |
| 3 | LWNFCP3 | s147.netapp.lab.local | 14-Dec-16 12:03 | 14-Dec-16 19:43 | 14-Dec-16 20:13 | 12.1.0.2.0 | 30.10 | 2,222.18 | 8.17 | 73.83 | Linux x86 64-bit |
| 4 | LWNFCP4 | s148.netapp.lab.local | 14-Dec-16 12:03 | 14-Dec-16 19:43 | 14-Dec-16 20:13 | 12.1.0.2.0 | 30.10 | 2,222.15 | 8.17 | 73.83 | Linux x86 64-bit |
| 5 | LWNFCP5 | s146.netapp.lab.local | 14-Dec-16 12:03 | 14-Dec-16 19:43 | 14-Dec-16 20:13 | 12.1.0.2.0 | 30.10 | 2,162.17 | 8.17 | 71.83 | Linux x86 64-bit |
| 6 | LWNFCP6 | s145.netapp.lab.local | 14-Dec-16 12:03 | 14-Dec-16 19:43 | 14-Dec-16 20:13 | 12.1.0.2.0 | 30.10 | 2,162.09 | 8.17 | 71.83 | Linux x86 64-bit |

### Top Timed Events

- Instance '*' - cluster wide summary
- '*' Waits, %Timeouts, Wait Time Total(s) : Cluster-wide total for the wait event
- '*' 'Wait Time Avg (ms)' : Cluster-wide average computed as (Wait Time Total / Event Waits) in ms
- '*' Summary 'Avg Wait Time (ms)' : Per-instance 'Wait Time Avg (ms)' used to compute the following statistics
- '*' [Avg/Min/Max/Std Dev] : average/minimum/maximum/standard deviation of per-instance 'Wait Time Avg(ms)'
- '*' Cnt : count of instances with wait times for the event

| | Wait | | Event | | Wait Time | | | Summary Avg Wait Time (ms) | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| I# | Class | Event | Waits | %Timeouts | Total(s) | Avg(ms) | %DB time | Avg | Min | Max | Std Dev | Cnt |
| * | User I/O | db file sequential read | 733,971,758 | 0.00 | 753,637.73 | 1.03 | 95.06 | 1.03 | 1.02 | 1.04 | 0.01 | 6 |
| | | DB CPU | | | 74,488.00 | | 9.40 | | | | | 6 |
| | System I/O | log file parallel write | 2,792,446 | 0.00 | 3,018.71 | 1.08 | 0.38 | 1.08 | 1.04 | 1.15 | 0.05 | 6 |
| | System I/O | db file parallel write | 14,238,559 | 0.00 | 1,863.18 | 0.13 | 0.24 | 0.13 | 0.13 | 0.13 | 0.00 | 6 |
| | Cluster | gc cr grant 2-way | 4,918,825 | 0.00 | 1,123.34 | 0.23 | 0.14 | 0.23 | 0.22 | 0.24 | 0.01 | 6 |
| | Concurrency | cursor: pin S wait on X | 5,506 | 0.00 | 59.51 | 10.81 | 0.01 | 10.77 | 9.36 | 12.61 | 1.09 | 6 |
| | System I/O | control file sequential read | 38,272 | 0.00 | 36.04 | 0.94 | 0.00 | 0.94 | 0.91 | 0.96 | 0.02 | 6 |
| | Concurrency | library cache: mutex X | 13,088 | 0.00 | 25.38 | 1.94 | 0.00 | 1.94 | 1.58 | 2.34 | 0.29 | 6 |
| | Other | LGWR wait for redo copy | 388,341 | 0.00 | 19.19 | 0.05 | 0.00 | 0.05 | 0.05 | 0.05 | 0.00 | 6 |
| | User I/O | read by other session | 15,590 | 0.00 | 15.89 | 1.02 | 0.00 | 1.02 | 1.00 | 1.04 | 0.01 | 6 |

 

# References

The following references were used in this technical report:


- The Silly Little Oracle Benchmark v2.2 (SLOB2)
  http://kevinclosson.net/2012/02/06/introducing-slob-the-silly-little-oracle-benchmark/
- NetApp TR-4514: NetApp AFF8080A EX Storage Efficiency and Performance with Oracle Database
  http://www.netapp.com/us/media/tr-4514.pdf
- NetApp AFF A700 and A700s product webpage
  http://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx


 

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.



NetApp