



Technical Report

NetApp AFF8080A EX Storage Efficiency and Performance with Oracle Database

Rodrigo Nascimento, Scott Lane, NetApp
June 2016 | TR-4514

Abstract

This report provides a storage efficiency and performance summary for NetApp and partner systems engineers interested in assessing Oracle database storage efficiency and performance with a NetApp® All Flash FAS (AFF) 8080A EX storage system.

NetApp All Flash FAS systems uniquely combine the extreme performance capability of flash media with NetApp ONTAP® 9 software to provide performance acceleration, operational agility, best-in-class data protection, and business continuance for database deployments.

TABLE OF CONTENTS

1	Introduction	3
1.1	ONTAP FlashEssentials Empowers All Flash FAS Performance.....	3
1.2	NetApp ONTAP 9.....	3
1.3	Storage Efficiency	4
2	Executive Summary.....	5
3	Measuring Storage Efficiency and Performance.....	6
3.1	Test Methodology	6
3.2	Hardware and Software	6
3.3	Network Design.....	8
3.4	Database Layout and Storage Provisioning Design.....	9
3.5	Workload Design.....	10
3.6	Storage Efficiency Test Results	10
3.7	Performance Test Results.....	12
4	Conclusion	13
	Appendix: AWR Report	14
	References.....	15

LIST OF TABLES

Table 1)	Oracle hardware and software components.....	6
Table 2)	NetApp storage system hardware and software.....	7
Table 3)	Database layout.	9

LIST OF FIGURES

Figure 1)	Network design.	8
Figure 2)	ONTAP 9.0 storage efficiencies.	12
Figure 3)	AFF8080A EX Oracle Database 12c performance.....	13

1 Introduction

The NetApp solution for Oracle databases delivers industry-leading storage, unprecedented 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 storage efficiency and performance of the NetApp AFF8080A EX storage system with Oracle database workloads.

1.1 ONTAP FlashEssentials Empowers All Flash FAS Performance

NetApp ONTAP FlashEssentials is the power behind the performance and efficiency of All Flash FAS (AFF). ONTAP is a well-known operating system, but what is not widely known is that Data ONTAP, with its WAFL® (Write Anywhere File Layout) file system, is natively optimized for flash media.

ONTAP FlashEssentials encapsulates key features that optimize solid-state-drive (SSD) performance and endurance, including the following:

- 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 deployment have shown that total data reduction technologies have enabled up to 933 times space savings.
- Inline data compaction provides continued innovation beyond compression and deduplication, further increasing storage efficiency.
- 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.
- New advanced drive partitioning (ADP) increases storage efficiency and further increases usable capacity by almost 20%.
- Data-fabric readiness enables live workload migration between flash and hard-disk-drive tiers, on premise or to the cloud.
- Quality-of-service capability safeguards service-level objectives in multiworkload and multitenant environments.

1.2 NetApp ONTAP 9

NetApp ONTAP 9 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. Plus, new customers and existing Data ONTAP 8.3 environments can quickly and easily use the rich data services delivered by ONTAP 9.

An essential feature for Oracle databases 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. Customers can depend on ONTAP 9 and AFF to provide these essential elements.

Built on the Data 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 allow customers to eliminate performance silos and seamlessly integrate AFF into a shared infrastructure. ONTAP 9 delivers enhanced inline deduplication and completely new inline data compaction capability

that significantly reduces the amount of flash storage required, with no impact on system performance. ONTAP 9 also provides industry-leading ecosystem integration with database applications that makes administration of databases and storage systems far more efficient when compared with other flash storage solutions on the market.

1.3 Storage Efficiency

Simply stated, storage efficiency enables you to store the maximum amount of data within the smallest possible space at the lowest possible cost. The following NetApp storage efficiency technologies can help you realize maximum space savings:

- **Inline compression.** Data compression reduces the disk space required, regardless of storage protocol, application, or storage tier. Inline compression also reduces the data that must be moved to SSDs, thereby reducing the wear on SSDs. Furthermore, the reduction in the amount of data being written can deliver an increase in overall performance.
- **Inline and always-on deduplication.** Data deduplication cuts storage requirements by reducing redundancies in primary, backup, and archival data. Inline deduplication of zeros speeds up VM provisioning by 20% to 30%. Improvements to inline deduplication in ONTAP 9 provide additional efficiency by extending elimination of duplicate data to blocks in memory and SSDs.¹
- **Inline data compaction.** NetApp inline data compaction in ONTAP 9 provides significant storage savings by compressing and coalescing small I/O into single block writes. Doing so further reduces the disk space required and associated wear on SSDs.
- **Snapshot technology.** NetApp Snapshot technology provides low-cost, instantaneous, point-in-time space-efficient copies of the file system (volume) or LUN by preserving ONTAP architecture and WAFL consistency points without affecting performance. NetApp SnapManager® for Oracle automates and simplifies Oracle database management with backup, recovery, restore, and cloning features with no downtime.
- **Thin provisioning.** Thin provisioning, implemented by NetApp at the NetApp FlexVol® volume and LUN level, defers storage purchases by keeping a common pool of free storage available to all applications.
- **Thin replication.** Thin replication is at the center of the NetApp data protection software portfolio, which includes NetApp SnapMirror® and NetApp SnapVault® software. SnapVault thin replication enables more frequent backups that use less storage capacity because no redundant data is moved or stored. SnapMirror thin replication protects business-critical data while minimizing storage capacity requirements.
- **RAID DP.** NetApp RAID DP® technology protects against double disk failure without sacrificing performance or adding disk-mirroring overhead.
- **FlexClone volumes.** FlexClone virtual cloning reduces the need for storage by enabling multiple, instant, space-efficient, writable copies.
- **Advanced drive partitioning (ADP).** Advanced SSD partitioning with ONTAP 9 increases usable capacity by almost 20%.²

The total effective capacity is the total amount of logical data that can be stored based on savings from compression, deduplication, data compaction, Snapshot copies, and FlexClone copies. The efficiency ratio is calculated by taking the sum of all of the logical data used in each aggregate and

¹ <https://www.netapp.com/us/media/tr-4476.pdf>

² <https://www.netapp.com/us/media/ds-3582.pdf>

comparing it to the sum of physical blocks used to support the logical data. The higher the ratio, the greater the space savings.

2 Executive Summary

To showcase the benefits of the AFF8080A EX running ONTAP 9, NetApp conducted the following studies:

- **Measuring ONTAP 9 storage efficiency.** NetApp measured the storage efficiency of several customer production databases after data reduction by inline compression, inline deduplication, and inline data compaction.
- **Measuring AFF8080A EX Oracle Database 12c performance.** NetApp measured the Oracle server read latency and the data throughput and input/output operations per second (IOPS) of the AFF8080A EX storage controllers running ONTAP 9 with an Oracle online transaction processing (OLTP) workload. All inline storage efficiency features were enabled.

NetApp conducted two studies to measure the storage efficiency and performance of an AFF8080A EX storage system running ONTAP 9. There were two goals. The first was to determine the storage savings on several samples of customer production data from the storage efficiencies delivered by ONTAP 9. The second goal was to determine the Oracle server read latency and the peak sustained IOPS of the AFF8080A EX storage system while running an Oracle OLTP workload.

We measured the overall storage efficiency savings of the customer production Oracle databases deployed on ONTAP 9. In our test configuration, a total space savings of 23:1 was observed. This total savings is attributed to the overall efficiency effectiveness of Snapshot copies, inline compression, and inline data compaction.

Additionally, we ran an OLTP workload called SLOB2 on a 2-node AFF8080A EX cluster with 2 DS2246 shelves containing a total of 48 800GB SSD drives. We tested our cluster at a range of load points that drove the storage to peak utilization. At each load point, we collected information on the storage IOPS and server latency.

The Oracle Database 12c performance test demonstrated that the AFF8080A EX performed from 137K IOPS at 200µs to a peak performance of 332K IOPS at 1600µs. For all load points up to 309K IOPS, the test configuration maintained consistent server read latencies below 1ms.

3 Measuring Storage Efficiency and Performance

NetApp conducted two studies to measure the storage efficiency and performance of an AFF8080A EX storage system running ONTAP 9. The following sections describe the methodology and design considerations used to test the AFF8080A EX while running a standard Oracle workload.

3.1 Test Methodology

In the storage-efficiency study, we measured the storage efficiency of both the data generated by the SLOB2 workload generator as well as three different customer production databases.

To measure the storage efficiency of data generated by the SLOB2 workload generator, we used SLOB2 to populate an Oracle database. We created one aggregate on each of the storage nodes of the AFF8080A EX. After installing the client operating systems and Oracle database software and configuring all software, we used SLOB2 to populate the Oracle database as it would be before a SLOB2 performance test run. After the database population was completed, we measured the storage efficiency of the workload files using the `storage aggregate show-efficiency ONTAP` command on the cluster.

For our storage-efficiency testing, we also used three different customer production databases of sizes 26GB, 70GB, and 4.7TB. To measure the storage efficiency of the customer production databases, we first created two aggregates, one on each storage node. Each aggregate contained one volume. We copied the database into a volume on one node and then used the Linux command `dd` to copy the data into the volume on the second node. Using the `dd` command, we used a block size of 8KB to simulate the native Oracle I/O block size. Copying these database files allowed the ONTAP inline efficiencies to operate on the data. We used the ONTAP `storage aggregate show-efficiency` command to collect the storage-efficiency measurements.

For the performance study, we used the SLOB2 load-generation tool to simulate an OLTP workload against the Oracle Database 12c 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 Fibre Channel (FC) to the AFF8080A EX. We allocated disk space for the Oracle database on the AFF8080A EX storage system. Using the six Oracle database servers and the OLTP load generator, we measured the peak performance of the storage system by generating a workload designed to maximize the storage system utilization. We then reran the test while ramping up the workload from 32 users to 440 users. Doing so allowed us to gather performance metrics at a range of different load points.

3.2 Hardware and Software

We configured six Oracle Database 12c database servers on six Fujitsu RX300s7 servers. We connected the six servers to a two-node AFF8080A EX through 16Gb FC. Each AFF node was connected to one DS2246 disk shelf with 800GB SSD drives following NetApp cabling best practices. The shelves were connected to the nodes, each in its own stack, in a multipath-HA configuration.

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

Table 1) Oracle hardware and software components.

Hardware/Software Components	Details
Oracle Database 12c servers	6 Fujitsu RX300s7 servers
Server operating system	Red Hat Enterprise Linux 6.6

Hardware/Software Components	Details
Oracle database version	12c (RAC)
Processors per server	2 6-core Xeon E5-2630 @ 2.30GHz
Physical memory per server	128GB
FC network	8Gb FC with multipathing
FC host bus adapter (HBA)	QLogic QLE2562 dual-port PCIe FC HBA
Dedicated public 1GbE ports for cluster management	2 Intel 1350GbE ports
8Gb FC switch	Brocade 6510 24-port
10GbE switch	Cisco Nexus 5596

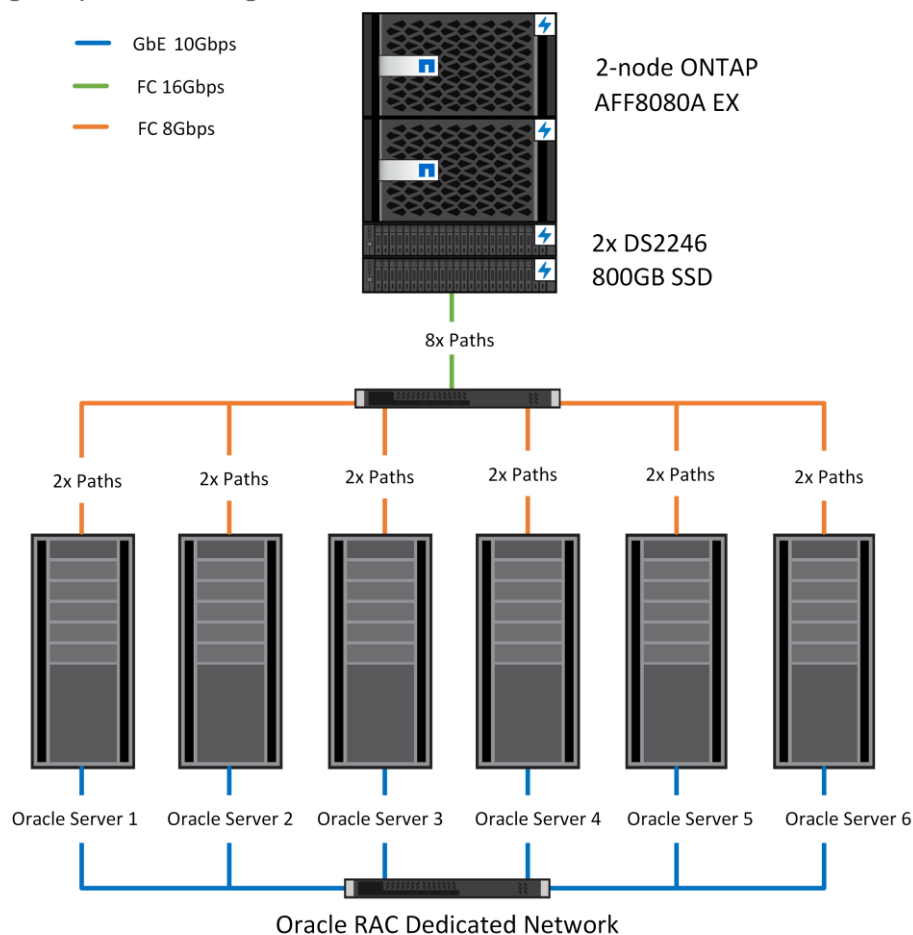
Table 2) NetApp storage system hardware and software.

Hardware/Software Components	Details
Storage system	AFF8080A EX configured as a high-availability (HA) active-active pair
ONTAP version	9.0
Total number of drives	48
Drive size	800GB
Drive type	SSD
FC target ports	8 16Gb ports (4 per node)
Ethernet ports	4 10Gb ports (2 per node)
Storage virtual machines (SVMs)	1 across both node aggregates
Ethernet logical interfaces (LIFs)	4 1Gb management LIFs (2 per node connected to separate private VLANs)
FC LIFs	8 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. Each server had two ports connected to the switch. The multiple ports used in the FC SAN configurations provided HA 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. Each storage system node housed a single shelf containing 24 800GB SSD drives. Of those 24 drives, 23 drives were used to create one single 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 using Oracle RAC.

Table 3) Database layout.

Storage	Aggregate Name	Volume Name	LUN Size (GB)	Vol Size (GB)	Description
Per node	n1_aggr0				Used ADP Total aggregate size = 970GB
	n2_aggr0				Used ADP Total aggregate size = 970GB
	n1_aggr1				21 data + 2 parity RAID DP + 1 spare Total aggregate size = 12.7TB
	n2_aggr1				21 data + 2 parity RAID DP + 1 spare Total aggregate size = 12.7TB
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

Storage	Aggregate Name	Volume Name	LUN Size (GB)	Vol Size (GB)	Description
	n2_aggr1	dbFCP_vol18	100	100	
	n1_aggr1	dbFCP_vol19	20	20	Grid: CRS and Voting
	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 using the Oracle Automatic Storage Management (ASM) volume manager. Those ASM disk groups provided the storage required to create the tablespaces. The FC SAN was configured on the Brocade switch. ONTAP provided asymmetric logical unit assignment (ALUA) communication to the initiators so that direct paths were used for host I/O access according to the multipathing input/output (MPIO) load-balancing policies on the host.

We deployed zoning in our configuration to balance the FC connections 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 our performance results, we tested the environment by increasing the number of Oracle users in SLOB2 from a minimum of 32 users up to a maximum of 440 users 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, but we acknowledge that workloads 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. Your results might vary from those found in this report.

3.6 Storage Efficiency Test Results

We measured the overall storage efficiency savings of four databases on ONTAP 9. The four databases that we used were:

- 2.1TB of data generated from the SLOB2 workload generator
- 26GB database files from a customer production environment
- 70GB database files from a customer production environment
- 4.7TB database files from a customer production environment

For the SLOB2 workload generator storage-efficiency testing, we used two disk shelves, each containing 24 800GB SSD drives, for a total raw storage of 38.4TB, as mentioned in section 3.2 “Hardware and Software.” We used ONTAP ADP to partition the drives, install the root partition, and create the data partition by using RAID DP to provide data redundancy, which resulted in total usable storage of 25.5TB. Of this usable space, we used the SLOB2 workload generator to create 2.1TB of storage for a single Oracle database, including data, redo log, and grid volumes. The amount of compression that can be

achieved is highly dependent on the specific contents of the data that is written and stored in the database.

As shown in Figure 2, from the 2.1TB of SLOB2 storage, we measured a space savings of 976GB from inline compression. This resulted in a storage efficiency ratio of approximately 2:1 from inline compression. In this dataset, we found significant additional storage savings from the new inline data compaction feature. When evaluating the storage efficiency of both inline compression and inline data compaction, we found savings of approximately 17:1 from this dataset.

Similarly, we measured the storage efficiency savings of data of files from three customer production databases with dataset sizes of 26GB, 70GB, and 4.7TB. For the 26GB dataset, we measured a storage efficiency ratio of 1.7:1 from inline compression and a total savings of 3:1 when considering both inline compression and inline data compaction.


For the 70GB dataset, we measured a storage efficiency ratio of 1.3:1 from inline compression and a total savings of 1.5:1 when considering both inline compression and inline data compaction.

For the 4.7TB dataset, we measured a storage efficiency ratio of 1.3:1 from inline compression and a total savings of 1.5:1 when considering both inline compression and inline data compaction.

Overall, the three customer production datasets provided a total weighed average savings from both inline compression and inline data compaction that resulted in a storage efficiency ratio of 1.5:1.

As shown in Figure 2, both inline compression and inline data compaction represent significant storage efficiency in all of our test database configurations. The solid green section of each bar represents the amount of space saved by inline compression. The hashed green section represents the amount of space saved by inline data compaction and the blue section is the amount of data actually written to disk.

Figure 2) ONTAP 9.0 storage efficiencies.




Space-efficient NetApp Snapshot copies can provide additional storage efficiency benefits. There is no performance penalty for creating Snapshot copies because data is never moved as it is with other copy-out technologies. The cost for Snapshot copies is only at the rate of block-level changes, not at 100% for each backup, as is the case with mirror copies. Snapshot technology can help you save on storage costs for backups and restores, and the technology opens up a number of efficient data management possibilities.

In a typical real-world Oracle customer deployment, hourly database volume Snapshot copies can account for up to an additional 24 times storage savings over a 24-hour period. Assuming 5% hourly overwrites and given our customer production 3 database environment, the storage savings of each Snapshot copy would be 2.8TB, for a total savings of 66.5TB. The total storage savings ratio achieved through inline compression, inline data compaction, and Snapshot copies was 23:1. Due to the unique header and tail information that Oracle adds to each block, deduplication provided little additional storage efficiency savings; therefore, we did not include calculations from this feature.

3.7 Performance Test Results

In addition to the storage efficiency testing, we measured the performance of our Oracle database implementation running on ONTAP 9. As shown in Figure 3, the Oracle performance test results on ONTAP 9 demonstrated that the cluster generated from 137K IOPS at 200μs to a peak performance of 332K IOPS at 1600μs. For all load points up to 309K IOPS, we were able to maintain consistent server read latencies below 1ms. As shown in Figure 3, ONTAP 9 provides up to 28% increase in IOPS under a millisecond over clustered Data ONTAP 8.3.1.

Figure 3) AFF8080A EX Oracle Database 12c performance.



We collected an Oracle Automated Workload Repository (AWR) report at the 309K IOPS performance point. The “Database Summary” section of the AWR shows the elapsed time and 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 amount of redo log operations in KB/s. For details of the AWR report, see the Appendix.

4 Conclusion

NetApp AFF8080A EX running ONTAP 9 software provides extremely high IOPS at consistently low latencies while serving an Oracle Database 12c OLTP workload. Our tests showed that the AFF8080A EX cluster performed from 137K IOPS at 200μs to a peak performance of 332K IOPS at 1600μs. For all load points up to 309K IOPS, we were able to maintain consistent server read latencies below 1ms and provide total space-efficiency savings of up to 23:1 on customer production data.

Many of the financial and business advantages of AFF8080A EX derive from the extensive storage-efficiency portfolio available with NetApp storage solutions. NetApp storage-efficiency technologies work together on a single unified architecture, and they can be enabled or disabled to serve any requirement, application, or environment. NetApp leads the way in bringing value to its customers. The company has built its reputation based on storage efficiency, helping customers achieve what they previously thought impossible and partnering with customers to get the most value out of their IT environments.

Appendix: AWR Report

The following screenshots show the AWR report that was collected at the 309K IOPS point of the performance test.

System Statistics - Per Second

#	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	45,341.98	42,175.56	10,975.65	8,877.12	21,774.57	4.07	666.91	3.77	0.52	166.16
2	45,269.86	42,098.64	10,966.33	8,860.23	21,731.36	4.07	664.72	2.90	0.52	165.83
3	43,218.79	40,213.07	10,478.66	8,468.31	20,759.57	4.55	635.58	3.06	0.52	158.53
4	43,140.19	40,185.37	10,465.68	8,457.23	20,729.05	4.05	635.47	3.48	0.51	158.37
5	43,159.67	40,215.20	10,461.74	8,464.04	20,740.23	4.05	635.20	2.82	0.51	158.48
6	43,265.82	40,315.15	10,492.76	8,481.06	20,788.50	4.06	637.24	3.52	0.52	158.80
Sum	263,396.31	245,202.99	63,840.82	51,607.99	126,523.28	24.85	3,875.12	19.56	3.10	966.17
Avg	43,899.38	40,867.16	10,640.14	8,601.33	21,087.21	4.14	645.85	3.26	0.52	161.03
Std	1,090.64	984.98	256.53	207.30	516.26	0.20	15.49	0.38	0.00	3.85

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
478409441	LWNFCP	YES	8192	471	472	6	6	6	6	7,690.20	30.20

Database Instances Included In Report

- Listed in order of instance number, #

#	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	11-Apr-16 10:18	12-Apr-16 14:02	12-Apr-16 14:32	12.1.0.2.0	30.22	1,321.73	28.23	43.74	Linux x86 64-bit
2	LWNFCP2	s147.netapp.lab.local	11-Apr-16 10:18	12-Apr-16 14:02	12-Apr-16 14:32	12.1.0.2.0	30.22	1,321.70	28.23	43.74	Linux x86 64-bit
3	LWNFCP3	s146.netapp.lab.local	11-Apr-16 10:18	12-Apr-16 14:02	12-Apr-16 14:32	12.1.0.2.0	30.22	1,261.72	28.23	41.76	Linux x86 64-bit
4	LWNFCP4	s149.netapp.lab.local	11-Apr-16 10:18	12-Apr-16 14:02	12-Apr-16 14:32	12.1.0.2.0	30.22	1,261.67	28.23	41.75	Linux x86 64-bit
5	LWNFCP5	s148.netapp.lab.local	11-Apr-16 10:18	12-Apr-16 14:02	12-Apr-16 14:32	12.1.0.2.0	30.22	1,261.68	28.23	41.75	Linux x86 64-bit
6	LWNFCP6	s145.netapp.lab.local	11-Apr-16 10:18	12-Apr-16 14:02	12-Apr-16 14:32	12.1.0.2.0	30.22	1,261.70	28.23	41.76	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)				
	Class	Event	Waits	%Timeouts	Total(s)	Avg(ms)	%DB time	Avg	Min	Max	Std Dev	Cnt
*	User I/O	db file sequential read	444,350,605	0.00	438,802.25	0.99	95.10	0.99	0.99	0.99	0.00	6
		DB CPU			38,883.33		8.43					6
	System I/O	log file parallel write	1,788,334	0.00	1,844.78	1.03	0.40	1.03	1.02	1.05	0.01	6
	System I/O	db file parallel write	5,400,970	0.00	930.57	0.17	0.20	0.17	0.17	0.17	0.00	6
	Cluster	gc cr grant 2-way	3,468,905	0.00	645.01	0.19	0.14	0.19	0.18	0.20	0.01	6
	System I/O	control file sequential read	40,473	0.00	24.02	0.59	0.01	0.59	0.58	0.62	0.02	6
	Other	reliable message	33,280	0.00	15.32	0.46	0.00	0.49	0.38	0.64	0.10	6
	User I/O	read by other session	11,690	0.00	12.33	1.05	0.00	1.05	1.01	1.08	0.03	6
	Configuration	undo segment extension	507	96.65	10.67	21.05	0.00	21.15	19.63	22.44	1.15	6
	Other	oracle thread bootstrap	580	0.00	9.32	16.07	0.00	16.07	15.46	16.88	0.63	6

References

The following references were used in this TR:

- The Silly Little Oracle Benchmark v2.2 (SLOB2)
<http://kevinclosson.net/2012/02/06/introducing-slob-the-silly-little-oracle-benchmark/>

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.

Copyright Information

Copyright © 1994–2016 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, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Fitness, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, SnapCopy, Snap Creator, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. TR-4514-0616