



Technical Report

NetApp AFF8080 EX Performance and Server Consolidation with Oracle Database

Workload Engineering, NetApp
June 2015 | TR-4415

Abstract

This report provides a performance and server consolidation summary for NetApp® and partner systems engineers who are interested in assessing Oracle database performance and return on investment (ROI) with a NetApp AFF8080 EX storage system.

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

TABLE OF CONTENTS

1	Introduction	4
1.1	Data ONTAP FlashEssentials Empowers All Flash FAS Performance	4
1.2	NetApp Clustered Data ONTAP 8.3.1	4
1.3	Storage Efficiency	4
1.4	Oracle	5
1.5	Database Server Consolidation	5
2	Executive Summary	6
2.1	Measuring Performance	6
2.2	Improving CPU Utilization	7
2.3	Enabling Database Consolidation	7
3	Measuring Performance	8
3.1	Test Methodology	8
3.2	Hardware and Software	8
3.3	Network Design	9
3.4	Workload Design	11
3.5	AFF8080 EX Oracle Database 12c Test Results	12
3.6	Storage Efficiency Test Results	12
4	Improving Oracle Database CPU Utilization	13
4.1	Test Methodology	13
4.2	Hardware and Software	13
4.3	Database Layout and Storage Provisioning Design	14
4.4	Oracle Utilization Test Results	16
5	Enabling Database Consolidation	17
6	Conclusion	19
	References	19
	Version History	19

LIST OF TABLES

Table 1)	Oracle hardware and software components	8
Table 2)	NetApp storage system hardware and software	9
Table 3)	Database layout	11
Table 4)	Oracle hardware and software	14

Table 5) NetApp AFF storage hardware and software.	14
Table 6) Database layout.	15
Table 7) Cost/benefit analysis.	18

LIST OF FIGURES

Figure 1) AFF8080 EX IOPS versus server read latency.	7
Figure 2) Network design.	10
Figure 3) AFF8080 EX Oracle Database 12c performance.	12
Figure 4) Network diagram.	16
Figure 5) Storage IOPS and database server CPU utilization improvements.	17

1 Introduction

This document describes the performance of the NetApp AFF8080 EX storage system with Oracle database workloads, a consolidation scenario for Oracle database servers, and a return-on-investment (ROI) scenario realized by replacing your legacy storage array with a NetApp AFF8080 EX storage system.

1.1 Data ONTAP FlashEssentials Empowers All Flash FAS Performance

NetApp Data ONTAP FlashEssentials is the power behind the performance and efficiency of All Flash FAS (AFF). Data 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.

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

- Mars[™] operating system innovations are now in clustered Data ONTAP.
- NetApp data-reduction technologies, including inline compression and deduplication, can provide up to 10:1 space savings, on average, for a typical use case. Savings can be further increased by using NetApp Snapshot[®] and NetApp FlexClone[®] technologies.
- Coalesced writes to free blocks maximize performance and flash media longevity.
- Rebuilt I/O path optimizes flash performance.
- Parallelized processing delivers low latency.
- Advanced drive partitioning increases storage efficiency.
- Data-fabric readiness enables live workload migration between flash and hard disk drive (HDD) tiers, on premises or to the cloud.
- Quality-of-service (QoS) capability safeguards service-level objectives in multiworkload and multi-tenant environments.

1.2 NetApp Clustered Data ONTAP 8.3.1

An essential feature for Oracle databases deployed on shared enterprise storage is the ability 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 clustered Data ONTAP 8.3.1 and AFF to provide these essential elements.

Built on clustered Data ONTAP unified scale-out architecture, AFF consistently meets or exceeds the high performance demands of Oracle databases. It 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. Clustered Data ONTAP delivers an enhanced inline compression capability that significantly reduces the amount of flash storage required and carries near-zero effects on system performance. It 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.

NetApp is a global enterprise scale-out storage and data management fabric provider, and clustered Data ONTAP has been an industry-leading operating system since 2012. Onsite ready but cloud connected, clustered Data ONTAP is a complete solution that is future-proof in a rapidly changing technology environment.

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.
- **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%. Always-on deduplication running every minute provides more space savings than postprocessed deduplication.
- **Snapshot technology.** NetApp Snapshot technology provides low-cost, instantaneous, point-in-time copies of the file system (volume) or LUN by preserving Data 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.

1.4 Oracle

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.

1.5 Database Server Consolidation

In today's data centers, HDD-based, legacy, shared storage systems serve up data to large numbers of individual Oracle systems and related applications. As your environment grows and you add more applications with ever-increasing performance requirements, these legacy storage systems struggle to keep up with the demand, requiring cost-prohibitive upgrades to maintain high performance in the future. This struggle can lead to significant underutilization of database server data center assets, because database applications must wait on slower legacy storage systems that no longer deliver the type of performance that users expect from Oracle. The result is high storage latencies, which slow down your applications and cause widespread performance issues that can severely limit your ability to derive the full value of your database servers.

Because of Oracle licensing policies, it can become expensive to maintain large numbers of Oracle servers that cannot take full advantage of server CPU resources. Eliminating storage limitations in your legacy environment by using an AFF8080 EX can unlock this potential and allow you to consolidate your Oracle environment onto fewer physical servers, save money on licenses, reduce the footprint in your data center, and save on power, cooling, and administrative overhead associated with maintaining more Oracle instances than are necessary to handle the workload.

2 Executive Summary

To showcase the benefits of the AFF8080 EX, NetApp conducted the following studies:

- **Measuring AFF8080 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 AFF8080 EX storage controllers running clustered Data ONTAP 8.3.1 with an Oracle online transaction processing (OLTP) workload.
- **Improving Oracle database CPU utilization with AFF8080 EX.** We compared the Oracle Database 11g Release 2 (11gR2) database server CPU utilization of servers attached to an HDD-based legacy storage system to servers attached to an AFF8080 EX storage system. The goal of the testing was to show how the performance capabilities of the underlying storage system directly affect the utilization efficiency of database server assets.
- **Enabling Oracle database consolidation by using AFF8080 EX.** Oracle license and support costs are generally based on CPU cores, which can significantly increase costs for database servers that contain high numbers of CPU cores. Therefore, you can potentially realize significant savings by reducing the number of Oracle servers in your environment. To investigate this possibility, NetApp leveraged the performance improvements provided by the AFF8080 EX to enable an Oracle database consolidation effort that reduced the number of physical database servers by 50%.


2.1 Measuring Performance

NetApp conducted a study to measure the performance of an AFF8080 EX storage system running clustered Data ONTAP 8.3.1 under an Oracle OLTP workload. The goal of the study was to determine the Oracle server read latency and the peak sustained throughput and IOPS of the AFF8080 EX storage system while running the Oracle OLTP workload.

We ran an OLTP workload called SLOB2 on a two-node AFF8080 EX cluster with two 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 throughput, IOPs, and server latency.

As shown in Figure 1, the Oracle Database 12c performance test demonstrated that the cluster increased from 118K IOPS to a peak performance of 316K IOPS. For all load points below the peak, we were able to maintain consistent server read latencies below 1ms.

Figure 1) AFF8080 EX IOPS versus server read latency.



In addition, we measured the overall storage efficiency savings of our Oracle database implementation deployed on Data ONTAP 8.3.1. In this test configuration, a total space savings of 22:1 can be expected with the SLOB2 workload. This total savings is attributed to the overall efficiency effectiveness of RAID DP, Snapshot copies, and inline compression.

2.2 Improving CPU Utilization

By replacing a legacy storage system with an AFF8080 EX, we observed significant reductions in read latencies in the database environment along with corresponding increases in total IOPS and database server CPU utilization as follows:

- An overall reduction in read latencies at the database hosts of over 91%
- A greater than 12x improvement in overall IOPS observed at the storage
- A greater than 7.5x improvement in the CPU utilization observed on the Oracle database servers

2.3 Enabling Database Consolidation

Replacing your legacy storage system with an AFF8080 EX can significantly improve your overall database performance and increase CPU utilization on Oracle servers. Any change to an environment that maximizes CPU utilization means that Oracle CPU license costs are not wasted. By consolidating your Oracle environments to fewer servers with better CPU utilization, you can save money on licensing and operational costs.

We leveraged the performance improvements provided by the AFF8080 EX to enable an Oracle database consolidation exercise that allowed us to reduce the number of physical database servers and CPU cores by 50% while saving more than \$5 million in Oracle license costs over a three-year period.

3 Measuring Performance

NetApp studied the performance of an AFF8080 EX storage system to determine its peak sustained throughput, IOPS, and Oracle server read latency. The following sections describe the methodology and design considerations used to test the AFF8080 EX while running a standard Oracle workload.

3.1 Test Methodology

In this 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 AFF8080 EX. We allocated disk space for the Oracle database on the AFF8080 EX storage system. Using the five 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. This allowed us to gather performance metrics at a range of different load points.

In addition to the performance validation testing described earlier, we also measured the inline compression savings of the data written to the Oracle database from the SLOB2 OLTP workload.

3.2 Hardware and Software

For this study, we configured five Oracle Database 12c database servers on five Fujitsu RX300s7 servers. We connected the five servers to a two-node AFF8080 EX through 16Gb FC. We connected each node of the AFF8080 EX to a single DS2246 shelf and populated each shelf with 24 800GB SSD drives.

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	5 Fujitsu RX300s7
Server operating system	Red Hat Enterprise Linux 6.6
Oracle Database version	12c (RAC)
Processors per server	2 6-core Xeon E5-2630 @ 2.30 GHz
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	AFF8080 EX configured as a highly available (HA) active-active pair
Clustered Data ONTAP version	8.3.1
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 2 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 both HA through multipathing and increased bandwidth. At no point in the testing did the network connectivity create a bottleneck.

Figure 2) Network design.



Database Layout and Storage Provisioning Design

Table 3 summarizes the layout for the Oracle database.

We used one Oracle RAC database to host the simulated OLTP environment. Each storage system node housed a single aggregate containing 24 800GB SSD drives that were subdivided into RAID DP groups, plus one 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					Used advanced drive partitioning
	aggr0	root		55	Total aggregate size = 55GB
	aggr				21 data + 2 parity RAID DP + 1 spare Total aggregate size = 12.7TB
Oracle RAC configuration		db1_vol1	200	200	Data files
		db1_vol2	200	200	
		db1_vol3	200	200	
		db1_vol4	200	200	
		db1_vol5	200	200	
		db1_vol6	200	200	
		db1_vol9	200	200	
		db1_vol10	200	200	
		db1_vol11	200	200	
		db1_vol12	200	200	
		db1_vol7	20	100	Redo logs
		db1_vol8	20	100	

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. Clustered Data ONTAP provided Asymmetric Logical Unit Assignment (ALUA) communication to the initiators so that optimal 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, and the other port was zoned for the other two LIFs per node.

3.4 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.

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. 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.

3.5 AFF8080 EX Oracle Database 12c Test Results

As shown in Figure 3, the Oracle performance test demonstrated that the cluster increased from 118K IOPS to a peak performance of 316K IOPS. For all load points below the peak, we were able to maintain consistent server read latencies below 1ms.

Figure 3) AFF8080 EX Oracle Database 12c performance.



3.6 Storage Efficiency Test Results

In addition to the performance validation testing described in this report, we also measured the overall storage efficiency savings of our Oracle database implementation from Data ONTAP 8.3.1. As mentioned in section 3.2 “Hardware and Software,” we used two disk shelves, each containing 24 800GB SSD drives, for a total raw storage of 19.2TB. We used Data ONTAP advanced drive partitioning 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 allocated 2.1TB of storage for a single Oracle database, including data and redo log volumes.

The amount of compression that can be achieved is highly dependent on the specific data that is written and stored in the database. From the 2.1TB of storage, we measured a space savings of 0.9TB from inline compression. This reduction resulted in an effective size of 1.1TB on disk and a storage efficiency ratio of 1.8:1 from inline compression.

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 it 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 24x storage savings over a 24-hour period. Assuming 5% hourly overwrites, the storage savings of each Snapshot copy would be 1TB, for a total of 24TB in savings, resulting in a storage savings ratio of 12:1 from Snapshot efficiencies. The total storage savings ratio achieved through both inline compression and Snapshot copies was 22:1. We discovered that deduplication provided little additional storage savings in our environment, so we did not enable this feature.

4 Improving Oracle Database CPU Utilization

HDD-based legacy storage systems can introduce performance issues into your Oracle environment by forcing database servers to wait longer than necessary for I/O responses from the storage. The longer the database servers wait on the storage, the more time it takes to complete transactions, which affects users and causes underutilization of CPU resources on your Oracle database servers.

Replacing a slower legacy storage system with an AFF8080 EX can lead to significantly better performance, lower latencies, and better database server CPU utilization.

4.1 Test Methodology

We conducted a study to measure the performance and database server CPU improvements that were observed as a result of replacing a performance-limited, HDD-based legacy storage system with a NetApp AFF8080 EX storage system running clustered Data ONTAP 8.3.1 under an OLTP workload.

The goal of this test was not to measure the maximum performance of the AFF8080 EX, but rather to use a consistent workload to measure the improvement in overall performance and database server CPU utilization observed after replacing the underperforming legacy storage system with an AFF8080 EX in an existing Oracle environment. For these tests, the legacy storage system was configured with a total of 144 450GB 15K RPM hard drives.

We created a test configuration with 10 database servers connected through FC to both the legacy storage system and the AFF8080 EX. Each of the 10 servers ran Oracle Database 11gR2 and created an Oracle database on both the legacy storage system and the AFF8080 EX storage system.

To generate a load in the test environment, we used the publically available SLOB2 workload generator to drive an OLTP-like workload simultaneously from each of the 10 database servers. Initially, we directed the workload to the Oracle databases on the legacy storage array. We increased the workload until we observed unacceptable performance in the OLTP environment that was using the legacy storage system. For this effort, we defined unacceptable performance as the inability to deliver consistent read latencies under 15ms as measured by observing Automatic Workload Repository (AWR) reports on the Oracle database servers.

We then noted the number of SLOB2 users required and recorded this as our baseline configuration. In addition, we recorded the IOPS observed at the legacy storage system as well as the average CPU utilization and latency on each of the 10 Oracle database servers. We used mpstat to measure average CPU utilization, and we extracted latency from the Oracle AWR reports.

After determining the baseline performance of the legacy storage system, we used the identical database and SLOB2 configurations, directed the same workload to the AFF8080 EX configured with 48 400GB SSDs, and gathered the same metrics.

4.2 Hardware and Software

For this study, we configured 10 Oracle database instances on 10 Fujitsu RX300s6 servers. We connected the 10 servers to a two-node AFF8080 EX storage system through 8Gb FC. Each tablespace had data files on both storage nodes. We connected each node of the AFF8080 EX to a single DS2246 shelf. We populated each shelf with 24 400GB SSD drives.

Table 4 and Table 5 list the hardware and software components used for testing.

Table 4) Oracle hardware and software.

Hardware/Software Components	Details
Oracle Database 11gR2 servers	10 Fujitsu RX300s6
Server operating system	Red Hat Enterprise Linux 6.6
Oracle Database version	11gR2
Processors per server	2 6-core Xeon E5645 @ 2.40GHz
Physical memory per server	48GB
FC network	8Gb FC with multipathing (AFF8080 EX)
FC HBA	QLogic QLE2562 dual-port PCIe FC HBA
Dedicated public 1GbE ports	2 Intel 1350GbE ports
8Gb FC switch	Brocade 6510 24-port
10GbE switch	Cisco Nexus 5596

Table 5) NetApp AFF storage hardware and software.

Hardware/Software Components	Details
Storage system	AFF8080 EX configured as an HA active-active pair
Clustered Data ONTAP version	8.3.1
Total number of drives	48
Drive size	400GB
Drive type	SSD
FC target ports	4 8Gb ports (2 per node)
Ethernet ports	4 10Gb ports (2 per node)
SVMs	1 across both node aggregates
Ethernet LIFs	4 1Gb management LIFs (2 per node connected to separate private VLANs)
FC LIFs	4 8Gb data LIFs

4.3 Database Layout and Storage Provisioning Design

Table 6 summarizes the disk layout for the Oracle database.

The legacy storage array and the AFF8080 EX each contained two storage nodes with their own sets of disk drives. Each node of the legacy storage array contained a total of 72 450GB 15K RPM disk drives. Each node of the AFF8080 EX contained 24 400GB SSD drives that were subdivided into RAID DP groups, plus 1 spare drive. The 10 databases were spread evenly across the 2 nodes of both the legacy storage system and the AFF8080 EX so that each node served data for all 10 databases and included database data, redo logs, temp, and flash recovery areas.

For the AFF8080 EX, we used clustered Data ONTAP advanced drive partitioning to create two partitions on each shelf: one for the root aggregate and one for the data aggregate.

Table 6) Database layout.

Storage	Aggregate Name	Name	LUN Size (GB)	Vol Size (GB)	Description
Per controller					Used advanced drive partitioning
	aggr0	root		55	Total aggregate size = 55GB
	aggr				21 data + 2 parity RAID DP + 1 spare Total aggregate size = 6.3TB
Per Oracle DB	One ASM data disk group	datavol1_1	512	512	Data files
		datavol1_2	512	512	
		datavol1_3	512	512	
	One ASM redo log disk group	logvol1_1	40	40	Redo logs
		logvol1_2	40	40	
	One ASM temp disk group	tempvol1_1	64	64	Temp
		tempvol1_2	64	64	
	One ASM flash recovery area disk group	fravol1_1	64	64	Flash recovery area
		fravol1_2	64	64	


We used three data LUNs to create one Oracle ASM data disk group per database. Within each data disk group, we created a 1TB tablespace for data. Additionally, we used the two redo log LUNs to create one redo log disk group. Within each redo log disk group, we created 2 30GB redo log members. Similarly, we created one temp disk group by using the two temp LUNs and one flash recovery area disk group by using the two flash recovery area LUNs. Each of the 10 databases used the per-Oracle-DB layout described in Table 6.

As mentioned earlier, the legacy storage system was connected to the 10 database servers through 8Gb FC connections, using 1 igroup per server. We verified connectivity and performance to the legacy storage system before testing.

The resulting ASM disks were used to create the Oracle ASM disk groups. Those disk groups were created with an allocation unit size of 64MB and external redundancy. The FC SAN was configured on the Brocade switch. Clustered Data ONTAP provided ALUA communication to the initiators so that optimal paths were used for host I/O access according to the MPIO load-balancing policies on the host.

Figure 4 shows a high-level picture of the environment.

Figure 4) Network diagram.



4.4 Oracle Utilization Test Results


In our test environment, we found that the legacy storage system eventually became the limiting factor, resulting in consistent read latencies of approximately 15ms.

After replacing the legacy storage system with the AFF8080 EX, we observed the following performance-related improvements:

- An overall 91% reduction in read latencies at the database hosts
- A greater than 12x improvement in the overall IOPS observed at the storage
- A greater than 7.5x improvement in the CPU utilization observed on the Oracle database servers

Figure 5 shows the results of the tests, which demonstrated that eliminating performance bottlenecks caused by slower HDD-based legacy storage systems can significantly improve the overall performance of your Oracle Database 11gR2 environment. Improved IOPS, lower latencies, and effective use of CPU resources allow your Oracle systems to process more transactions faster, which improves business efficiency and keeps both internal and external users satisfied.

Figure 5) Storage IOPS and database server CPU utilization improvements.



5 Enabling Database Consolidation

Through testing, we demonstrated that replacing your performance-limited, HDD-based legacy storage system with an AFF8080 EX can significantly improve your overall performance and allow you to make more efficient use of the Oracle database server CPU resources in your existing data center.

Using AFF8080 EX, storage bottlenecks are almost entirely eliminated. You can now confidently plan to consolidate your Oracle Database 11gR2 instances onto fewer servers and potentially realize significant savings on Oracle Database 11gR2 licenses, support, power, floor space, and administrative costs.

Using the results of the previous testing, we performed an analysis of the existing environment before and after replacing the legacy storage system with an AFF8080 EX. Specifically, we consolidated the 10 Oracle Database 11gR2 instances and related databases that were currently running on 10 physical servers onto fewer physical servers.

The goal of this exercise was to compare the total cost of ownership (TCO) of the existing Oracle Database 11gR2 environment in contrasting scenarios over a three-year period:

- In the first scenario, you continue to use the existing legacy storage system for the Oracle environment by renewing the service agreements with your existing storage and database partners. Because of performance limitations with the legacy storage, you are unable to justify the risk of consolidating your Oracle environment at this time.
- In the second scenario, you immediately replace the legacy storage system with a new NetApp AFF8080 EX and embark on an effort to consolidate your current database environment from 10 physical database servers down to a total of 5, while maintaining the same number of databases and applications in your current environment.

For this analysis, we used the NetApp Realize investment-analysis software application. The NetApp Realize tool has been examined and validated by International Data Corporation (IDC), which concluded that results produced by NetApp Realize are based on sound methodologies. In addition, IDC found that NetApp Realize uses industry-standard default input values that fall within accepted industry ranges. For additional information about the Realize tool, contact your NetApp account representative.

The analysis was conducted based on cash flows, which is the method normally employed for evaluating investment decisions. As mentioned earlier, the storage solution options evaluated included the existing

legacy storage system and the AFF8080 EX. Return on investment (ROI) and payback period were calculated for the AFF8080 EX over a three-year period to understand the costs and benefits of consolidating the database environment by replacing the existing legacy storage.

The Realize tool used the existing legacy storage system and the 10 currently allocated database servers running Oracle as a base for estimating the potential cost savings that could be realized by using the AFF8080 EX. The value of the potential savings was used in the ROI calculations.

The relevant portions of the existing legacy data storage hardware and software environments, such as support renewal, space, power, and cooling for the legacy environment, were included in the analysis. We also assumed that the legacy environment would be phased out immediately in favor of the AFF8080 EX.

For this exercise, we used list prices for Oracle licenses and software assurance gathered from the Oracle website.

Table 7 provides a summary of the analysis, which revealed the following key points:

- Replacing the legacy storage system with an AFF8080 EX combined with the overall database consolidation would yield an ROI of 90%.
- Consolidating the Oracle environment from 10 physical servers to 5 physical servers and reducing Oracle license costs by 50% would achieve more than \$5 million in savings over a three-year period.
- Reducing the footprint of the overall environment saves power, space, and labor costs.
- Replacing the legacy storage system with an AFF8080 EX running clustered Data ONTAP 8.3.1 enables additional savings by using NetApp nondisruptive operations (NDO) features, such as reductions in planned and unplanned down times.

Table 7) Cost/benefit analysis.

Value	Analysis Results
ROI	90%
Net present value (NPV)	\$4,800,000
Payback period (months)	6 months
Cost reduction	More than \$5 million saved over a 3-year analysis period compared to the legacy storage system
Savings on power, space, and administration	\$40,000
Additional savings from NDO benefits (not included in ROI)	\$90,000

These findings represent just one possible scenario for consolidating Oracle servers with NetApp storage. Contact your local NetApp account representative to perform a server consolidation ROI analysis tailored to your specific requirements.

Finally, we repeated the test, using SLOB2 and the database environment after consolidating from 10 servers to 5 servers. For this effort, each of the five remaining Oracle servers drove an OLTP workload to two of the original Oracle databases. We observed the following key points:

- As expected, we observed a 2x increase in database server CPU utilization after reducing the number of Oracle servers.
- Overall performance was in line with what we observed in the original environment that used 10 database servers.

6 Conclusion

The NetApp AFF8080 EX solution provides extremely high IOPS at consistently low latencies while serving an Oracle Database 12c OLTP workload. Our testing showed that the AFF8080 EX throughput increased from 118K IOPS to a peak performance of 316K IOPS as we increased the OLTP workload. For all load points below the peak, we were able to maintain consistent server read latencies below 1ms and provide space-efficiency savings of up to 22:1.

We demonstrated that significant improvements in performance and database server CPU utilization are possible when you replace your older legacy storage system with an AFF8080 EX. Because it delivers excellent performance with consistently low latencies, the AFF8080 EX is a great option to enable additional savings during Oracle Database 11gR2 database consolidation efforts.

The AFF8080 EX solution delivers both financial and technical benefits. AFF8080 EX not only provides an attractive investment opportunity but also positions your IT environment for the future by providing the foundation for an agile data infrastructure that includes integrated data protection, nondisruptive operations, seamless scalability, intelligent data management, and embedded security.

Many of the financial and business advantages of AFF8080 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 and 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.

References

The following references are used in this TR:

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

Version History

Version	Date	Document Version History
Version 1.0	May 2015	Initial release

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 © 1994–2015 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, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, 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-4415-0615

