



Technical Report

All-Flash Business Processing SAN and ONTAP 9 Verification Tests Using Microsoft SQL Server Workloads

Business Workloads Group, DFMG, NetApp
September 2016 | TR-4532

Abstract

This document describes the results of a series of performance and functional tests demonstrating the consistent performance and fast failover features of NetApp® ONTAP® 9 operating within the constraints defined in an all-flash business processing (AFBP) configuration. These tests were conducted in a SAN environment using the FC protocol and leveraged an online transaction processing (OLTP)-based workload driven from a Microsoft SQL Server 2014 database.

TABLE OF CONTENTS

| | | |
|----------|---|-----------|
| 1 | Introduction | 4 |
| 2 | Executive Summary | 4 |
| 2.1 | Test Cases..... | 4 |
| 2.2 | Testing Methodology..... | 5 |
| 2.3 | Test Results..... | 6 |
| 3 | Product Overview | 8 |
| 3.1 | AFF8000 Series Systems..... | 8 |
| 3.2 | NetApp ONTAP 9 Software..... | 9 |
| 3.3 | NetApp OnCommand Performance Monitor..... | 9 |
| 3.4 | NetApp SnapCenter..... | 9 |
| 3.5 | Microsoft SQL Server 2014..... | 10 |
| 4 | Network Fabric Topology | 10 |
| 5 | Database Storage Provisioning | 12 |
| 6 | Test Case Overview | 13 |
| 7 | Test Results | 14 |
| 7.1 | Measure Baseline Performance of Configuration Using AFBP Requirements (No Database Backups Performed)..... | 14 |
| 7.2 | Measure Impact of SQL Server Database Backups Initiated from SnapCenter 1.1 on AFBP Baseline Workload..... | 15 |
| 7.3 | Measure Performance Impact of Creating Clones of the SQL Server Database Initiated from SnapCenter 1.1 on AFBP Baseline Workload..... | 16 |
| 7.4 | Measure Performance Impact of Deleting Clones of the SQL Server Database Initiated from SnapCenter 1.1 on AFBP Baseline Workload..... | 16 |
| 7.5 | Performance and Availability Impact of Unplanned Failover Events..... | 17 |
| 7.6 | Performance and Availability Impact of Planned Failover Events..... | 19 |
| 8 | Conclusion | 21 |
| | Appendix A: Detailed Test Cases | 22 |
| | Appendix B: Deployment Details | 28 |

LIST OF TABLES

| | | |
|----------|---|----|
| Table 1) | AFBP test cases..... | 5 |
| Table 2) | Baseline performance, backup, and clone creation/deletion test results..... | 6 |
| Table 3) | Planned and unplanned failover test results..... | 7 |
| Table 4) | SQL Server hardware and software..... | 11 |

| | |
|---|----|
| Table 5) NetApp storage hardware and software. | 12 |
| Table 6) Test case summary. | 13 |
| Table 7) AFBP baseline performance test results. | 14 |
| Table 8) Impact on AFBP baseline performance when conducting backups. | 16 |
| Table 9) Impact on AFBP baseline performance when creating clones of the SQL Server database. | 16 |
| Table 10) Impact of AFBP baseline performance when deleting clones of the SQL Server database. | 17 |
| Table 11) Unplanned storage takeover or giveback and database I/O interruption times. | 18 |
| Table 12) Planned storage takeover and giveback I/O interruption times. | 20 |
| Table 13) Storage configuration and layout. | 28 |
| Table 14) Server network specifications. | 30 |
| Table 15) Storage network specifications. | 30 |
| Table 16) FC back-end switches. | 31 |

LIST OF FIGURES

| | |
|--|----|
| Figure 1) Impact on baseline performance during unplanned storage controller failures. | 8 |
| Figure 2) Network fabric topology. | 11 |
| Figure 3) Storage layout for SQL Server. | 13 |
| Figure 4) Baseline performance results using OPM. | 15 |
| Figure 5) Storage controller IOPS, latency, and performance capacity usage during unplanned storage controller failure testing. | 19 |
| Figure 6) Storage controller IOPS, latency, and performance capacity usage during planned storage controller failure testing. | 21 |

1 Introduction

NetApp storage has been a critical part of SAN-based storage infrastructures delivering enterprise-class performance for business-critical applications. With the advent of all-flash storage systems such as the NetApp AFF8080, customers can now consider leveraging the performance of flash in both enterprise and non-enterprise environments, providing increased flexibility and consistent, low-latency performance to a wider variety of applications. NetApp refers to this expanded focus as all-flash business processing (AFBP).

This document describes the capabilities of NetApp All Flash FAS (AFF) storage platforms used in these business-critical environments. Specifically, it describes a series of tests that demonstrate the performance and availability characteristics of a SQL Server 2014 database that supports OLTP workloads supported by a NetApp AFF8080 storage system.

All the components that were used in the tests were configured in compliance with NetApp best practices for an AFBP SAN environment. The goal of the AFBP SAN configuration is to provide predictable performance for mission-critical and non-mission-critical workloads while taking full advantage of the reliability, availability, and serviceability features of NetApp ONTAP 9 data management software. The NetApp framework for AFBP consists of a set of prescriptive guidelines for the configuration, deployment, and monitoring of AFF controllers in an AFBP scenario.

The goal of the tests was to subject the AFBP-compliant test environment to a transactional-based workload driven from a SQL Server 2014 database that pushed the storage to the limits of what is considered acceptable in an AFBP environment. After establishing this baseline workload, a series of tests were conducted under load to induce both planned and unplanned storage failures into the environment and measure the impact on overall performance.

Additionally, the NetApp SnapCenter® 1.1 application was used to create a series of full, application-consistent backups of the SQL Server 2014 database as well as create and delete clones of the database while under the same transactional workload described earlier. The overall goal of the testing was to demonstrate that if you follow the procedures to configure your MS-SQL database and operate within the AFBP guidelines, the NetApp AFF will provide consistent, predictable performance during normal day-to-day operations as well as fast recovery in the unlikely event of a storage failure.

These AFBP best practices are described in [TR-4515: Best Practices for AFF Business-Processing SAN Workloads ONTAP 9](#). Although these best practices apply to both FAS8000 and AFF8000 systems, this document focuses on the NetApp AFF8080 system.

2 Executive Summary

At a high level, the results of the testing demonstrated that normal, day-to-day activities such as creating backups, NetApp Snapshot® copies, and clones of a database did not significantly impact the overall performance of the SQL Server 2014 environment when configured using the AFBP guidelines. Additionally, with ONTAP 9 and Windows Server 2012R2 client software, there were significant improvements to the overall user experience with regard to both planned and unplanned storage failure events with the SQL Server 2014 database compared to previous versions of ONTAP.

The remainder of this section provides a high-level overview of the test methodology and results depicted in this document. Please refer to the Appendix A for detailed test cases.

2.1 Test Cases

As shown in Table 1, the test cases include establishing a performance baseline in an AFBP-compliant environment using a SQL Server 2014 database and AFF8080 storage system for a typical OLTP workload. The Microsoft TPC-E benchmarking tool was used to establish a performance baseline for the AFBP SAN environment with the following workload simulating an OLTP database environment:

- 90% reads and 10% writes and a 100% random access pattern

To achieve this baseline, we followed the methodology described in [TR-4515: Best Practices for AFF Business-Processing SAN Workloads ONTAP 9](#) for calculating the maximum performance capacity available in the test environment. In this example, we used the TPC-E load generation tool and NetApp OnCommand® Performance Monitor (OPM) to drive and analyze a series of workloads with the characteristics described above that subjected the AFF8080 storage system to increasingly higher levels IOPS over a long period of time. The goal was to plot a curve that mapped IOPS and latency at the different workload levels.

The end goal of this effort was to determine the level of IOPS and latency at which OPM believes the AFF8080 had been pushed to the maximum performance capacity, as defined in [TR-4515: Best Practices for AFF Business-Processing SAN Workloads ONTAP 9](#).

In the case of AFBP configurations, the requirements were to subject the storage to performance levels that did not exceed 50% of the level of the maximum performance capacity as defined in [TR-4515: Best Practices for AFF Business-Processing SAN Workloads ONTAP 9](#).

After establishing the performance baseline, a series of additional tests were executed while the AFF8080 storage system was subjected to the AFBP performance baseline workload generated from the SQL Server database.

These tests included planned and unplanned storage failures and used NetApp SnapCenter 1.1 to create and delete database backups and clones while under load. The impact of each action on the baseline performance was observed and recorded. Table 1 is a summary of the AFBP test cases.

Table 1) AFBP test cases.

| Test Case | Description |
|-----------|--|
| TC-01 | Measure baseline performance of configuration using AFBP requirements (no database backups performed) |
| TC-02 | Measure impact of SQL Server database backups initiated from SnapCenter 1.1 on AFBP baseline workload |
| TC-03 | Measure performance impact of creating clones of the SQL Server database initiated from SnapCenter 1.1 on AFBP baseline workload |
| TC-04 | Measure performance impact of deleting clones of the SQL Server database initiated from SnapCenter 1.1 on AFBP baseline workload |
| TC-05 | Measure performance and availability impact of failover events—unplanned failovers |
| TC-06 | Measure performance and availability impact of failover events—planned failovers |

For more information about how each test was conducted, see the corresponding test results in Appendix A.

2.2 Testing Methodology

At a high level, test cases TC–02 through TC–06 were executed by using the following methodology:

1. Start the specific AFBP baseline workload defined in TC–01 and allow it to run for a specific amount of time to verify steady-state performance was achieved.
2. Introduce the specific feature or failure scenario while under load and allow the test to continue for a specific amount of time.
3. For failure-based tests, correct the issue (for example, perform a storage controller giveback) and allow the test to run to completion.

4. Gather both storage and database performance data to observe the impact to overall baseline performance over the duration of the testing.

2.3 Test Results

This section provides a high-level overview of the test results. These tests demonstrated that configuring and maintaining an environment using the AFBP guidelines allows for normal MS-SQL database operations such as snapshots, backups, and clones to be used with little to no impact on overall performance.

Additionally, in the unlikely event of a more disruptive event such as a storage controller failure, the surviving storage controller was able to provide adequate performance during the failure and return performance back to prefailure levels within seconds of correcting the failure.

In AFBP-compliant environments, the following criteria are required for optimal operations:

- The window of time for both storage node takeover and giveback operations is 2–10 seconds for a planned event.
- The window of time for a storage node takeover operation is 5–15 seconds for an unplanned event, while the subsequent giveback operation is 2–10 seconds.
- Average storage performance capacity under normal operating conditions reported by OPM should be approximately 50%, indicating that the existing workload is consuming approximately 50% of the total available capacity of the AFF8080 storage system.

Table 2 provides a summary of the test results from the perspective of the MS-SQL database and report read, write, and total IOPS observed as well as both read and write latencies. The results recorded in Table 2 demonstrate the following key items of interest and demonstrate that performance was maintained well within the guidelines for an AFBP-compliant environment:

- The AFBP baseline performance shows 137,812 IOPS were achieved with database read latencies well under 1ms while consuming 50% of the total storage performance capacity.
- Using SnapCenter 1.1 to initiate a series of hourly full database backups had no impact on IOPS, latencies, or consumption of storage performance capacity.
- Creating and deleting clones of the MS-SQL database had only a minimal impact on IOPS and latencies, but consumed slightly more storage performance capacity.

Table 2) Baseline performance, backup, and clone creation/deletion test results.

| Workload Mix | Total IOPS | Read IOPS | Write IOPS | Read Latency (ms) | Average Storage Performance Capacity from OPM |
|-----------------------------------|------------|-----------|------------|-------------------|---|
| AFBP baseline | 137,812 | 129,710 | 8,102 | 0.40 | 50% |
| AFBP baseline + hourly backups | 130,638 | 122,935 | 7,702 | 0.39ms | 50% |
| AFBP baseline + DB clone creation | 139,394 | 131,216 | 8,178 | 0.38ms | 57% |
| AFBP baseline + DB clone deletion | 136,943 | 128,933 | 8,011 | 0.38ms | 57% |

For the storage takeover and giveback testing, the AFBP baseline workload was started and ran for a total of 10 hours. During this time a series of 10 planned and unplanned storage failures was introduced and corrected while measuring the impact on the overall AFBP baseline performance. The results below in Table 3 show the following:

- The average time, in seconds (sec), reported by ONTAP to perform the takeover and giveback operations across all 10 planned and unplanned storage controller failures (Takeover Time and

Giveback Time columns, respectively). The takeover and giveback times in Table 3 represent the amount of time that the storage controllers were physically unable to serve I/O using the FC protocol during the takeover and giveback processes, respectively. These values were well within the guidelines defined by the AFBP for storage controller failures of 2–10 seconds for planned and 2–15 seconds for unplanned storage failures

- The average amount of time, in seconds, the SQL Server database was unable to request data from the AFF8080 during all 10 planned and unplanned storage controller failures (Takeover Time and Giveback Time columns, respectively). This data is heavily influenced by the amount of time required for the database host to work through FC path availability issues caused by the loss of the AFF8080 storage controller.

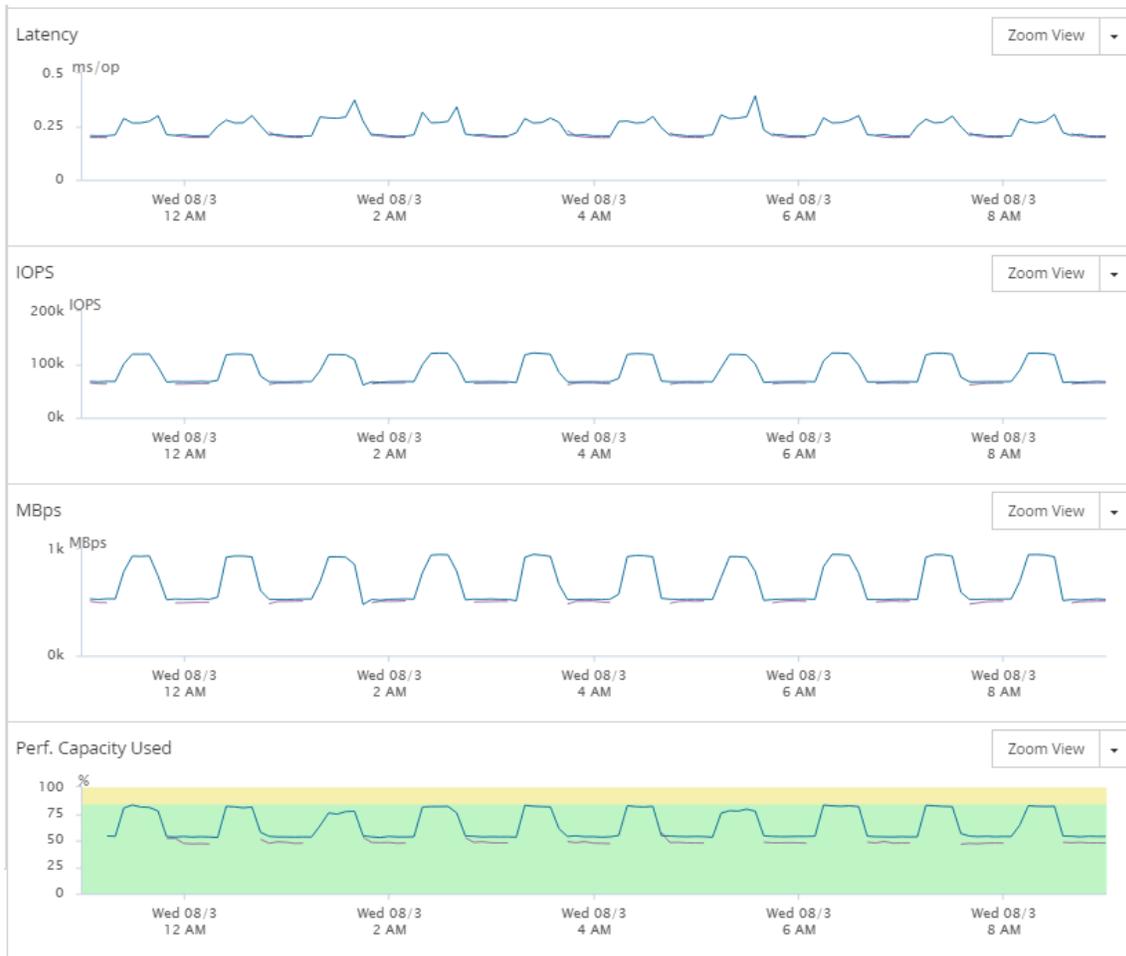
Table 3) Planned and unplanned failover test results.

| | Takeover Time (ms) | Giveback Time (ms) | Database I/O Interruption | |
|--------------------|--------------------|--------------------|---------------------------|----------------|
| | | | Takeover (sec) | Giveback (sec) |
| Unplanned failover | 4.1 | 0.9 | 6 | 7 |
| Planned failover | 1.3 | 1.5 | 9 | 8 |

Finally, Figure 1 shows the impact on baseline performance for each of the 10 planned and unplanned storage controller failures from the perspective of the surviving AFF8080 storage controller. These charts show the following items of interest:

- As expected, IOPS and latency values increase after the failure of one of the storage controllers forces the surviving storage controller to handle the entire workload for the period of time before the failure is corrected.
- In all cases, the average latency reported by both the storage and database remain well below the 2ms AFBP threshold for operating in a degraded environment.
- After correcting the storage failure and restoring the storage system to pre-failure status, overall performance and associated latencies return to pre-failure levels.
- As expected, when a single storage controller is servicing the entire workload, the percentage of total performance capacity consumed increases but never exceeds the total available storage capacity.

Figure 1) Impact on baseline performance during unplanned storage controller failures.



3 Product Overview

This section describes the NetApp and SQL Server products that were used in the solution.

3.1 AFF8000 Series Systems

NetApp AFF8000 series systems combine a unified scale-out architecture with leading data management capabilities. They are designed to adapt quickly to changing business needs while meeting core IT requirements for uptime, scalability, and cost efficiency. These systems offer the following advantages:

- **Speed the completion of business operations.** Leveraging a new high-performance, multicore architecture and self-managing flash acceleration, AFF8000 unified scale-out systems boost throughput and decrease latency. They deliver consistent application performance across a broad range of SAN and NAS workloads.
- **Streamline IT operations.** Simplified management and proven integration with cloud providers let you deploy the AFF8000 in your data center and in a hybrid cloud with confidence. Nondisruptive operations simplify long-term scaling and improve uptime by facilitating hardware repair, tech refreshes, and other updates without planned downtime.
- **Deliver superior total cost of ownership.** Proven storage efficiency and a twofold improvement in price/performance ratio over the previous generation reduce capacity utilization and improve long-term return on investment.

3.2 NetApp ONTAP 9 Software

NetApp ONTAP 9 software delivers a unified storage platform that enables unrestricted, secure data movement across multiple cloud environments and paves the way for software-defined data centers. It offers advanced performance, availability, and efficiency. ONTAP clustering capabilities help you keep your business running nonstop.

ONTAP is an industry-leading data management software. Its single feature-rich platform allows you to scale infrastructure without increasing IT staff. NetApp ONTAP provides the following benefits:

- Nondisruptive operations:
 - Perform storage maintenance, hardware lifecycle operations, and software upgrades without interrupting your business.
 - Eliminate planned and unplanned downtime.
- Proven efficiency:
 - Reduce storage costs by using one of the most comprehensive storage efficiency offerings in the industry.
 - Consolidate and share the same infrastructure for workloads or tenants with different performance, capacity, and security requirements.
- Seamless scalability:
 - Scale capacity, performance, and operations without compromise, regardless of application.
 - Scale SAN and NAS from terabytes to tens of petabytes without reconfiguring running applications.

3.3 NetApp OnCommand Performance Monitor

NetApp® OnCommand® Performance Manager (OPM), integrated with OnCommand Unified Manager, allows users to monitor and manage the performance of their NetApp ONTAP environment. It allows users to see notifications and alerts about their data storage performance in order to troubleshoot issues, isolate potential problems, and find concrete solutions.

OPM provides:

- A dashboard view of key performance indicators
- Access to the performance explorer, a visual navigation feature that provides a fast, easy way to analyze and compare the performance of your physical and logical objects and clusters
- The ability to drill down and perform root-cause analysis on bottlenecks
- The capability to determine and monitor optimal levels of storage performance capacity to make sure of compliance to AFBP or other guidelines.

3.4 NetApp SnapCenter

NetApp SnapCenter software is a unified, scalable platform for application-consistent data protection and clone management. It simplifies backup, restore, and clone lifecycle management with application-integrated workflows. With storage-based data management, SnapCenter enables increased performance and availability and reduced testing and development times.

NetApp SnapCenter includes both the SnapCenter Server and individual lightweight application, database, and operating system plug-ins that are all controlled from a central management console. The management console delivers a consistent user experience across all applications or databases. It incorporates a single GUI to support critical functions such as job monitoring, event notification, logging, dashboard, reporting, scheduling, and policy management for all application or database plug-ins.

SnapCenter Server also includes Snapshot copy catalog management to facilitate easy rollback to point-in-time copies. SnapCenter Server checks application and database and OS interoperability and then

nondisruptively installs and upgrades software plug-ins on application and database hosts. Those plug-ins can then be managed from the central management console.

Administrators can use the SnapCenter plug-ins for applications and databases so that the application or database is consistent at all levels, which promotes maximum recoverability. Plug-ins for SnapCenter allow a variety of restore capabilities. They can roll forward logs and enable application or database administrators to clone or recover to the latest information available or to a specific point in time.

We used SnapCenter to create and manage SQL Server 2014 database backups and clones during our testing.

3.5 Microsoft SQL Server 2014

SQL Server is the foundation of Microsoft's data platform, delivering mission-critical performance with in-memory technologies and faster insights on any data, whether on premises or in the cloud. Microsoft SQL Server 2014 builds on the mission-critical capabilities delivered in prior releases by providing breakthrough performance, availability, and manageability for mission-critical applications. This release of the SQL Server database engine introduces features and enhancements that increase the power and productivity of architects, developers, and administrators who design, develop, and maintain data storage systems.

4 Network Fabric Topology

Figure 2 shows the architecture of the validation testing configuration. According to the diagram, a two-node NetApp AFF8080 cluster was used, along with 4 shelves containing 96 x SSDs of 400GB. In addition, 8 x MS-SQL servers were used. The diagram also shows dual paths between the RAC nodes and a 10GbE network between the database servers, serving as a cluster interconnect.

Figure 2) Network fabric topology.

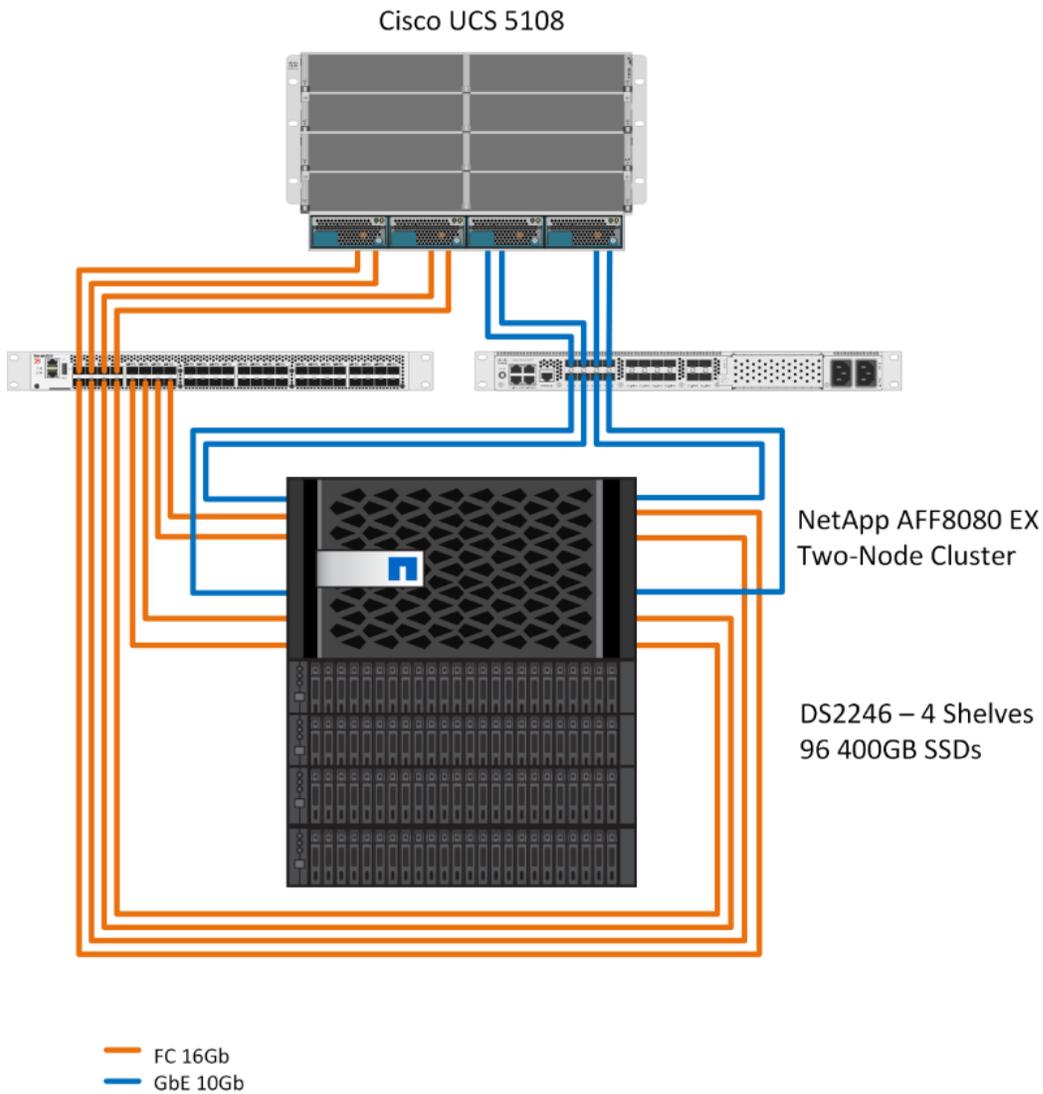


Table 4 and Table 5 show the SQL Server and NetApp hardware and software components used for testing. The storage was configured to make sure of compliance with the AFBP guidelines.

Table 4) SQL Server hardware and software.

| Hardware and Software Components | Details |
|----------------------------------|---|
| SQL Server database servers | Cisco UCS 5108 with 4 x B200 M4 blades (2 x SQL Servers per blade) |
| Server operating system | Windows Server 2012 R2 Datacenter Edition hosted on VMware ESX 5.5.0 |
| SQL Server database version | MS-SQL 2014 Enterprise Edition |
| Processors/server | 2 x socket, 14 cores per server |
| Physical memory/server | 32GB |

| Hardware and Software Components | Details |
|----------------------------------|---|
| FC network | 16Gb FC with multipathing |
| FC HBA | QLogic QLE2672 dual-port PCIe FC HBA |
| Network connections | 4 x Intel 82599ES 10Gbps SFI/SFP+ network connections |
| 16Gb FC switch | Brocade 6510 24-port |
| 10GbE switch | Cisco Nexus 5596 |

Table 5) NetApp storage hardware and software.

| Hardware and Software Components | Details |
|----------------------------------|--|
| Storage controller | AFF8080 EX configured as an HA active-active pair |
| ONTAP software | v9.0 |
| Number/size of SSDs | 96 x 400GB (48 SSDs per cluster node) |
| FC target ports | 8 x 16Gb (4 per node) |
| Ethernet ports | 4 x 10Gb (2 per node) |
| Storage virtual machines (SVMs) | 1 x across both node aggregates |
| Management LIFs (Ethernet) | 4 x 1GbE data (1 per node connected to separate private VLANs) |
| FC LIFs | 8 x 16Gb data |

5 Database Storage Provisioning

Database storage was configured to adhere to AFBP specifications for a NetApp AFF8080 system. A single aggregate was created on each of the two AFF8080 controllers as follows:

- 63 x 400GB SSDs using three RAID groups
- Total aggregate size of 12.07TB

By using this configured storage, a set of volumes was provisioned (each volume contained one FC LUN). The LUNs were used as components of a MS-SQL database. The following LUNs were created on each of the AFF8080 nodes:

- 8 x 3TB LUNs used for MS-SQL database portion
- 8 x 500GB LUNs used for MS-SQL log portion
- 8 x 500GB LUNs used for MS-SQL temp database

Figure 3) Storage layout for SQL Server.

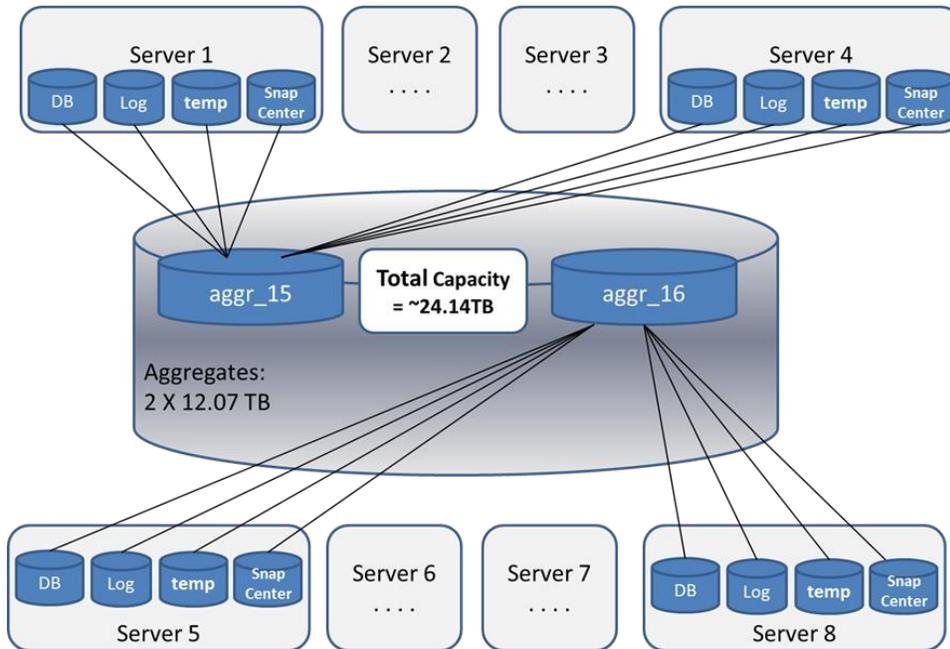


Table 13, located in Appendix A, lists the specifics for each of the AFF8080 storage nodes, including the aggregate, volume, and LUN sizes. Figure 3 shows the storage layout for SQL Server. It shows aggregates aggr_15 and aggr_16, one on each storage node, each using 42 x SSDs, leaving one spare SSD per node. With a RAID group size of 21, the configuration yielded 12.07TB of usable storage space per storage node, for a total of 24.14TB of usable space across both nodes.

Eight MS-SQL databases in all are hosted on this cluster, four on each node. The database was created by using the TPC-E tool. This configuration resulted in the TPC-E test workloads being evenly spread across both storage nodes.

6 Test Case Overview

All the test cases in Table 6 were executed by evaluating the operation of a specific NetApp feature in an otherwise normally configured system running a predefined load driven to the SQL Server cluster by the TPC-E workload generator that conformed to the AFBP prescribed guidelines for storage performance capacity.

The full details of the test methodologies for each test case are provided in Appendix B.

Table 6) Test case summary.

| Test Case | Description |
|-----------|--|
| TC-01 | Measure baseline performance of configuration using AFBP requirements (no database backups performed) |
| TC-02 | Measure performance impact of SQL Server database backups initiated from SnapCenter 1.1 on AFBP baseline workload |
| TC-03 | Measure performance impact of creating clones of the SQL Server database initiated from SnapCenter 1.1 on AFBP baseline workload |

| Test Case | Description |
|-----------|--|
| TC-04 | Measure performance impact of deleting clones of the SQL Server database initiated from SnapCenter 1.1 on AFBP baseline workload |
| TC-05 | Measure performance and availability impact of failover events: unplanned failovers |
| TC-06 | Measure performance and availability impact of failover events: planned failovers |

For more information about how each test was conducted, see the corresponding test results section and Appendix B.

7 Test Results

The following sections summarize the tests that were performed and the procedures and the results of each test.

7.1 Measure Baseline Performance of Configuration Using AFBP Requirements (No Database Backups Performed)

Before any tests were performed, a series of tests were conducted to establish a consistent performance baseline that conformed to the AFBP requirements using the following workloads:

- 90% reads and 10% writes with an 8KB request size and a 100% random access pattern

These baseline performance results became the basis for comparison of all the tests that were executed in the test plan.

To establish a baseline, it was necessary to determine the appropriate workload intensity by using the TPC-E workload generator to drive the load. By using an iterative process, the IOPS levels were driven successively higher until the available performance capacity reported by OPM for the overall cluster was reported to be 50%. This approach is in accordance with AFBP guidelines and is intended to prevent controller failure from resulting in excessive load on the surviving storage controller resulting in significant performance degradation in the AFBP environment.

After establishing the specific TPC-E parameters required to push the performance to 50% of the maximum storage performance capacity as reported by OPM, the test was run for a total of four hours to establish the baseline performance. The TPC-E parameters required to achieve that threshold were recorded and was used for all subsequent testing. This will be referred to as the AFBP baseline workload going forward.

Table 7 shows the resulting IOPS and latency as reported by the SQL Server databases over the duration of the four-hour test. It also shows the average CPU and disk utilization values that were observed over the duration of each of the four-hour baseline runs.

Table 7) AFBP baseline performance test results.

| Workload Mix | Total IOPS | Read IOPS | Write IOPS | Read Latency (ms) | Average Storage Performance Capacity from OPM |
|--------------|------------|-----------|------------|-------------------|---|
| 90/10 | 137,812 | 129,710 | 8,102 | 0.4 | 50% |

The results recorded in Table 7 for the AFBP baseline workload are summarized as follows:

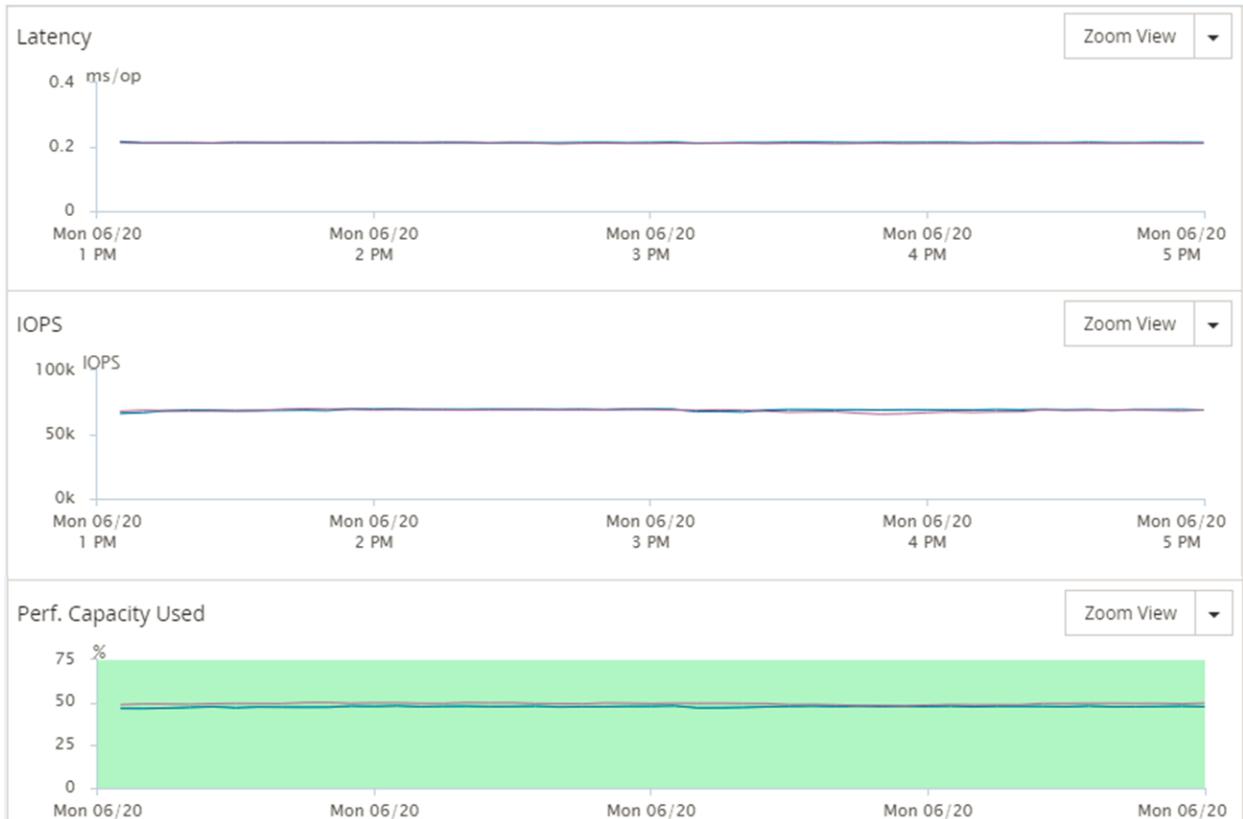
- An average of 137,812 IOPS were sustained over the four-hour test with an average read latency of 0.4ms.

- The AFBP baseline workload consumed an average of 50% of the total available performance capacity of the AFF8080 HA pair, which is the upper limit when operating within the AFBP requirements

It's worth noting that these baseline results do not reflect the maximum performance capabilities of the AFF8080 HA storage but were intentionally limited to conform to the requirements of the AFBP program.

Figure 4 shows the baseline performance from the storage perspective as observed by OPM and includes the total IOPS, average latency, and percentage of performance capacity consumed on the AFF8080. These results line up with results presented above observed at the database level.

Figure 4) Baseline performance results using OPM.



7.2 Measure Impact of SQL Server Database Backups Initiated from SnapCenter 1.1 on AFBP Baseline Workload

The purpose of this test was to measure the performance impact of creating a series of regularly scheduled full, application-consistent backups of the SQL Server database using SnapCenter 1.1. During the backups, the TPC-E workload generator was used to drive the same AFBP baseline workload described in TC-01. The following procedure was used to conduct the test:

1. Using the TPC-E workload utility, initiate the AFBP baseline workload defined in TC-01 for a duration of four hours.
2. Configure SnapCenter to initiate a full, application-consistent backup of the SQL Server database at hourly intervals. Configure the retention interval to four backup versions requiring every fifth backup to remove the oldest backup, causing the deletion of the storage-based Snapshot copy associated with the backup.
3. Record performance monitor data, sysstat, and OPM data over the course of the four-hour run to get a picture of the performance at both the database and storage levels.

Table 8 compares the IOPS, latencies, and storage performance capacity used for the AFBP baseline and the AFBP baseline when conducting hourly full backups of the SQL Server database. These results show essentially no impact on the performance of the database as a result of using SnapCenter to create scheduled database backups.

Table 8) Impact on AFBP baseline performance when conducting backups.

| Workload Mix | Total IOPS | Read IOPS | Write IOPS | Read Latency (ms) | Average Storage Performance Capacity from OPM |
|--------------------------------|------------|-----------|------------|-------------------|---|
| AFBP baseline | 137,812 | 129,710 | 8,102 | 0.4 | 50% |
| AFBP baseline + hourly backups | 130,638 | 122,935 | 7,702 | 0.39 | 50% |

7.3 Measure Performance Impact of Creating Clones of the SQL Server Database Initiated from SnapCenter 1.1 on AFBP Baseline Workload

The purpose of this test was to measure the impact of using SnapCenter 1.1 to create a full clone of the SQL Server database while running the AFBP baseline workload. The following procedure was used to conduct the test.

1. Using the TPC-E workload utility, initiate the AFBP baseline workload defined in TC-01 for a duration of 30 minutes.
2. After running for five minutes, use SnapCenter to create a clone of the SQL Server database using one of the four existing full backups created in previous testing.
3. Record the time required to complete the clone creation process.
4. Record Performance monitor data, sysstat, and OPM data over the course of the 30-minute run to get a picture of the performance at both the database and storage levels.

Table 9 compares the performance using the AFBP baseline to the same workload subjected to the creation of a full clone of the SQL Server database. These results show a very minimal performance impact during the clone creation process with performance over the entire 30-minute duration that is comparable to the AFBP baseline in terms of IOPS and latency.

Table 9) Impact on AFBP baseline performance when creating clones of the SQL Server database.

| Workload Mix | Total IOPS | Read IOPS | Write IOPS | Read Latency (ms) | Average Storage Performance Capacity from OPM |
|-----------------------------------|------------|-----------|------------|-------------------|---|
| AFBP baseline | 137,812 | 129,710 | 8,102 | 0.4 | 50% |
| AFBP baseline + DB clone creation | 139,394 | 131,216 | 8,178 | 0.38 | 57% |

7.4 Measure Performance Impact of Deleting Clones of the SQL Server Database Initiated from SnapCenter 1.1 on AFBP Baseline Workload

The purpose of this test was to measure the impact of using SnapCenter 1.1 to delete a full clone of the SQL Server database while running the AFBP baseline workload. The following procedure was used to conduct the test.

1. Using the TPC-E workload utility, initiate the AFBP baseline workload defined in TC-01 for a duration of 30 minutes.

2. After running for five minutes, use SnapCenter to delete the clone of the SQL Server database created earlier in the previous clone creation test.
3. Record the time required to complete the clone deletion process.
4. Record performance monitor data, sysstat, and OPM data over the course of the 30-minute run to get a picture of the performance at both the database and storage levels.

Table 10 compares the performance using the AFBP baseline to the same workload subjected to the deletion of a full clone of the SQL Server database using SnapCenter 1.1. These results show a very minimal performance impact during the clone deletion process with performance over the entire 30-minute duration that is comparable to the AFBP baseline in terms of IOPS and latency.

Table 10) Impact of AFBP baseline performance when deleting clones of the SQL Server database.

| Workload Mix | Total IOPS | Read IOPS | Write IOPS | Read Latency (ms) | Average Storage Performance Capacity from OPM |
|-----------------------------------|------------|-----------|------------|-------------------|---|
| AFBP baseline | 137,812 | 129,710 | 8,102 | 0.4ms | 50% |
| AFBP baseline + DB clone deletion | 136,943 | 128,933 | 8,011 | 0.38ms | 57% |

7.5 Performance and Availability Impact of Unplanned Failover Events

The purpose of this test was to measure the impact of unplanned storage node failover on database performance and availability while the database was driving the AFBP baseline workload to the AFF8080 storage system. The following procedure was used to conduct the unplanned failover test:

1. Initiate the baseline AFBP workload consisting of 90% reads and 10% writes with a 100% random access pattern by using the TPC-E tool for a total of 10 hours.
2. At hourly intervals over the course of 10 hours, induce a panic in one of the AFF8080 storage controllers in the HA pair causing a takeover of the failed node by the surviving AFF8080 storage controller. This results in the surviving AFF8080 storage controller handling the entire AFBP baseline workload.
3. Allow the surviving storage controller to service the workload during the time the other storage controller is down.
4. Using the standard timeout values, allow the surviving storage controller to perform a giveback of the failed storage controller to automatically restore the storage to the prefailure state with both controllers in the HA pair equally handling the load.
5. Record performance monitor data, sysstat, and OPM data over the course of the 10-hour run to get a picture of the performance at both the database and storage levels.

Table 11 provides the details of the unplanned storage failure testing as follows:

- The average time, in seconds, reported by ONTAP to perform the takeover and giveback operations across all 10 unplanned storage controller failures (Takeover Time and Giveback Time columns, respectively). The takeover and giveback times in Table 11 represent the amount of time that the storage controllers were physically unable to serve I/O using the FC protocol during the takeover and giveback processes, respectively. These values were well within the guidelines defined by the AFBP for storage controller failures of 2–10 seconds for planned and 2–15 seconds for unplanned storage failures.
- The average amount of time, in seconds, the SQL Server database was unable to request data from the AFF8080 during all 10 unplanned storage controller failures (Takeover Time and Giveback Time columns, respectively). This data is heavily influenced by the amount of time required for the

database host to work through FC path availability issues caused by the loss of the AFF8080 storage controller.

Table 11) Unplanned storage takeover or giveback and database I/O interruption times.

| Parameters Iteration | Takeover Time (sec) | Giveback Time (sec) | Database I/O Interruption | |
|----------------------|---------------------|---------------------|---------------------------|----------------|
| | | | Takeover (sec) | Giveback (sec) |
| 1 | 4.6 | 1.3 | 6 | 7 |
| 2 | 4.4 | 0.7 | 7 | 7 |
| 3 | 4.1 | 0.5 | 4 | 5 |
| 4 | 3.8 | 0.5 | 6 | 6 |
| 5 | 3.9 | 0.6 | 6 | 7 |
| 6 | 4.1 | 1.8 | 5 | 7 |
| 7 | 3.4 | 0.6 | 6 | 6 |
| 8 | 4.9 | 1.6 | 6 | 7 |
| 9 | 4.5 | 0.4 | 5 | 6 |
| 10 | 3.5 | 0.8 | 6 | 8 |

Finally, Figure 5 provides the details from the storage over the course of the 10 unplanned takeover and giveback events from the perspective of the surviving AFF8080 storage controller. The figure provides graphs generated using OPM and shows the following information:

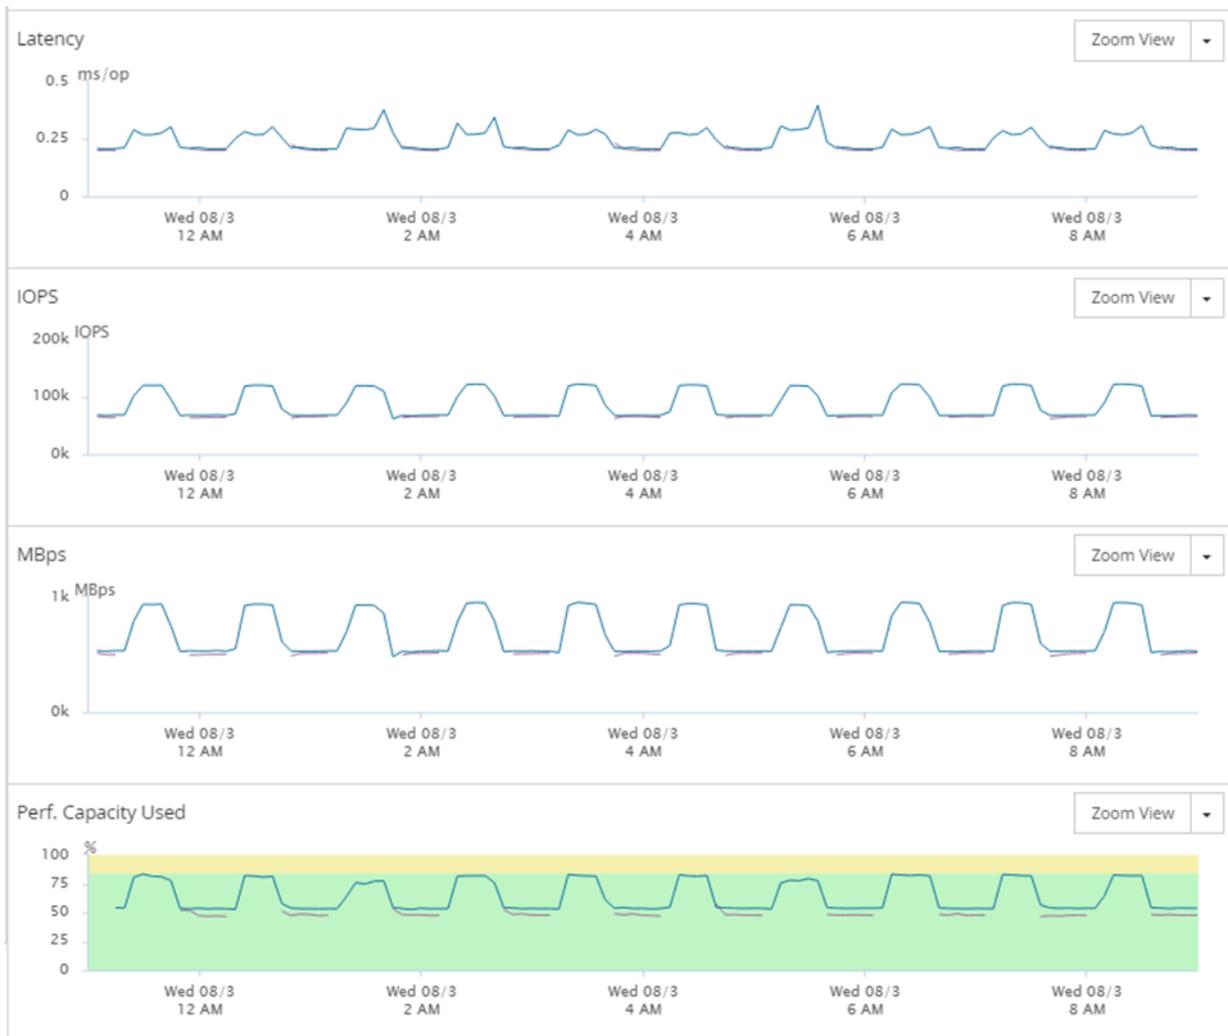
- IOPS and overall average latency.** As expected, overall IOPS increases along with the average latency after the takeover event as the surviving AFF8080 storage controller is handling the entire workload. In this case, the throughput increases from approximately 65,000 IOPS at 0.25ms latency to approximately 110,000 IOPS at 0.4ms overall average latency.

However, after the storage failure is corrected and the giveback completed, overall performance in terms of both IOPS and latency returns to prefailure levels in just a few seconds.

- Percentage of total performance capacity used.** This chart shows the percentage of the overall available storage performance capacity that is currently being consumed. Recall that the AFBP guidelines recommend that total performance capacity be kept at or below 50% of the total available performance capacity for the storage system.

As expected, when the surviving storage controller is handling the entire workload, the amount of performance capacity used on the single controller generally pushes close to or just over 75% and remains there for the duration of the failover. After the storage giveback process is complete and both storage controllers are again sharing the total workload, the performance capacity used drops back to acceptable levels for the AFBP configuration.

Figure 5) Storage controller IOPS, latency, and performance capacity usage during unplanned storage controller failure testing.



7.6 Performance and Availability Impact of Planned Failover Events

The purpose of this test was to measure the impact of a planned storage node failover on database performance and availability while the database was driving the AFBP baseline workload to the AFF8080 storage system. The following procedure was used to conduct the planned failover test:

1. Initiate the baseline AFBP workload consisting of 90% reads and 10% writes with a 100% random access pattern by using the TPC-E tool for a total of 10 hours.
2. At hourly intervals over the course of 10 hours, use the standard storage takeover and giveback commands to initiate a planned takeover of one of the AFF8080 storage controllers in the HA pair, causing the surviving AFF8080 storage controller to handle the entire AFBP baseline workload.
3. Allow the surviving storage controller to service the workload during the time the other storage controller is down.
4. Using the standard timeout values, allow the surviving storage controller to perform a giveback of the failed storage controller to automatically restore the storage to the prefailure state with both controllers in the HA pair equally handling the load.

- Record performance monitor data, sysstat, and OPM data over the course of the 10-hour run to get a picture of the performance at both the database and storage levels.

Table 12 provides the details of the planned storage failure testing as follows:

- The average time, in seconds, reported by ONTAP to perform the takeover and giveback operations across all 10 planned controller failures (Takeover Time and Giveback Time columns, respectively). The takeover and giveback times in Table 12 represent the amount of time that the storage controllers were physically unable to serve I/O using the FC protocol during the takeover and giveback processes, respectively. These values were well within the guidelines defined by the AFBP for storage controller failures of 2–10 seconds for planned and 2–15 seconds for unplanned storage failures.
- The average amount of time, in seconds, the SQL Server database was unable to request data from the AFF8080 during all 10 planned storage controller failures (Takeover and Giveback columns, respectively). This data is heavily influenced by the amount of time required for the database host to work through FC path availability issues caused by the loss of the AFF8080 storage controller.

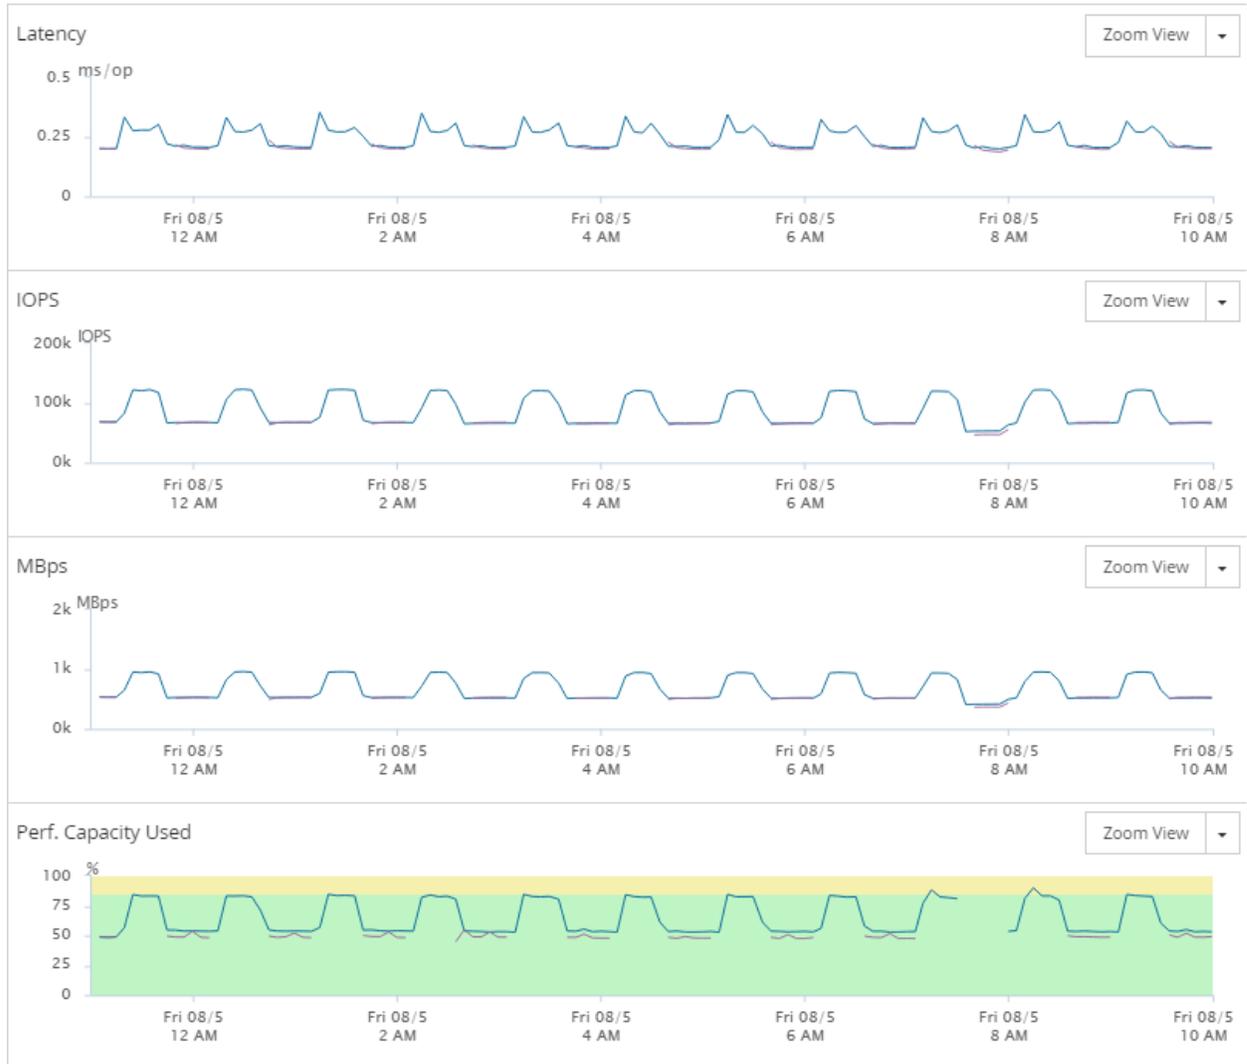
Table 12) Planned storage takeover and giveback I/O interruption times.

| Parameters Iteration | Takeover Time (sec) | Giveback Time (sec) | Database I/O Interruption | |
|----------------------|---------------------|---------------------|---------------------------|----------------|
| | | | Takeover (sec) | Giveback (sec) |
| 1 | 0.5 | 1.4 | 10 | 6 |
| 2 | 0.5 | 1.1 | 12 | 7 |
| 3 | 1.0 | 1.7 | 10 | 7 |
| 4 | 2.7 | 2.7 | 9 | 8 |
| 5 | 1.4 | 1.6 | 9 | 8 |
| 6 | 1.7 | 1.8 | 9 | 8 |
| 7 | 0.5 | 0.7 | 8 | 8 |
| 8 | 1.1 | 1.5 | 9 | 7 |
| 9 | 2.2 | 1.5 | 9 | 5 |
| 10 | 1.5 | 0.8 | 8 | 13 |

Finally, Figure 6 provides the details from the storage over the course of the 10 unplanned takeover and giveback events from the perspective of the surviving AFF8080 storage controller. The figures provide graphs generated using OPM and shows the following information:

- IOPS and overall average latency.** As expected, overall IOPS increases along with the average latency after the takeover event as the surviving AFF8080 storage controller is handling the entire workload. In this case, the throughput increases from approximately 65,000 IOPS at 0.25 ms latency to approximately 110,000 IOPS at 0.4 ms overall average latency.
However, after the storage failure is corrected and the giveback completed, overall performance in terms of both IOPS and latency returns to prefailure levels in just a few seconds.
- Percentage of total performance capacity used.** This chart shows the percentage of the overall available storage performance capacity that is currently being consumed. Recall that the AFBP guidelines recommend that total performance capacity be kept at or below 50% of the total available performance capacity for the storage system.

Figure 6) Storage controller IOPS, latency, and performance capacity usage during planned storage controller failure testing.



As expected, when the surviving storage controller is handling the entire workload, the amount of performance capacity used on the single controller generally pushes close to or just over 75% and remains there for the duration of the failover. Once the storage giveback process is complete and both storage controllers are again sharing the total workload, the performance capacity used drops back to acceptable levels for the AFBP configuration.

8 Conclusion

The AFBP SAN-compliant design used for these tests met the performance, availability, and predictability targets associated with the program when using a SQL Server database. Furthermore, the specific tests that were chosen for this document provide the necessary data for assessment of high-value data management, data protection, and data availability features of NetApp ONTAP 9.

Typical SQL Server database management functions including backups, Snapshot copies, and cloning were shown to have a negligible impact on the overall performance capabilities of the NetApp AFF8080 storage system.

As expected, events such as storage controller failures had a greater impact on overall performance in the AFBP-compliant environment. Even so, the overall environment maintained latency and storage performance capacity levels that continued to be within the acceptable limits of the AFBP guidelines, and performance returned to pre-failure levels within seconds of fixing the initial failure.

Results of the planned failover test indicate that planned failover should not be initiated during periods of peak usage, but should instead be planned for off-peak periods.

Finally, at no point during the testing did we observe that the SQL Server database crashed or had other operational issues as a result of a storage failure. The storage failure did briefly affect storage availability, but the OS did not report any I/O errors and the SQL Server database remained online and available. Generally, the storage was unavailable to the database for five seconds or less.

Attempts were made to use realistic workloads in the execution of all tests, but application workloads and OS configurations can vary greatly between production environments. Customers might see a variation in results. Adherence to documented AFBP SAN best practices promotes optimal storage performance and availability and facilitates business continuity in the event of storage failover events.

Appendix A: Detailed Test Cases

This appendix provides detailed information about the test cases that are described in this document.

TC-01: Measure Baseline Performance of Configuration using AFBP Requirements (No Database Backups Performed)

| Test Case | Details |
|--|--|
| Test number | TC-01 |
| Test case description | Using a test environment configured to be in compliance with the AFBP requirements, measure the overall performance of the SQL Server database when delivering a typical OLTP workload. This test did not include creation/deletion of any database backups during the test period. Total duration of the test was four hours to allow steady-state performance to be reached. |
| Test assumptions | <ul style="list-style-type: none"> • A completely operational NetApp cluster running clustered NetApp ONTAP 9 has been installed and configured within the AFBP guidelines. • A completely operational SQL Server environment has been installed and configured. • The TPC-E utility has been installed and configured to generate the following workloads: <ul style="list-style-type: none"> – 90% reads and 10% writes with a 100% random access pattern |
| Test data or metrics to capture | <ul style="list-style-type: none"> • Microsoft performance monitor data for SQL Server • Sysstat output from all NetApp AFF8080 controllers in the cluster • OPM performance capacity data for the duration of the test |
| Expected results | Baseline performance in this AFBP environment will be measured as IOPS and associated latency at the database level measured by Performance monitor over the course of the four-hour test. Additionally, the reported IOPS will be limited by making sure that total performance capacity as measured by OPM does not exceed 50% as required by AFBP guidelines. |

Test Methodology

1. Using the TPC-E workload utility and an OLTP workload consisting of a mix of 90% reads and 10% writes with a 100% random access pattern, initiate a series of 15-minute test iterations, successively increasing the overall workload directed by the SQL Server database to the AFF8080 storage controllers.
2. Continue test iterations until observing performance levels that push the storage controllers to storage performance capacity levels of 50% (measured by OPM), which is the maximum allowable to remain within the AFBP requirements.
3. Record the TPC-E parameters required to generate this level of load and execute an additional test at this workload level for a duration of four hours to establish steady-state performance.
4. Save the sysstat output and analyze it for the impact on storage availability during storage node takeover and for the impact on storage performance after the takeover.
5. Record performance monitor data, sysstat, and OPM data over the course of the four-hour run to get a picture of the performance at both the database and storage levels.

TC-02: Measure Impact of SQL Server Database Backups Initiated from SnapCenter 1.1 on AFBP Baseline Workload

| Test Case | Details |
|--|--|
| Test number | TC-02 |
| Test case description | <p>While subjected to the AFBP baseline workload computed as a result of TC-01 for a total of four hours, configure NetApp SnapCenter 1.1 to create a full backup of the SQL Server database every hour for a total of four full, application-consistent database backups.</p> <p>Observe the impact on overall AFBP baseline performance during the backups to make sure that conducting normal SQL Server database backups did not significantly impact database performance.</p> |
| Test assumptions | <ul style="list-style-type: none"> • A completely operational NetApp cluster running clustered NetApp ONTAP 9 has been installed and configured within the AFBP guidelines. • A completely operational SQL Server environment has been installed and configured. • The TPC-E utility has been installed and configured to generate the following workloads: <ul style="list-style-type: none"> – 90% reads and 10% writes with a 100% random access pattern • SnapCenter 1.1 was installed into the environment and configured to successfully create application-consistent backups of the SQL Server database using the AFF8080 storage configuration. |
| Test data or metrics to capture | <ul style="list-style-type: none"> • Microsoft performance monitor data for SQL Server • Sysstat output from all NetApp AFF8080 controllers in the cluster • OPM performance capacity data for the duration of the test |
| Expected results | Baseline performance in this AFBP environment is minimally affected by the SnapCenter backup process, which includes creating and deleting Snapshot copies of the volumes containing the database files. |

Test Methodology

1. Using the TPC-E workload utility, initiate the AFBP baseline workload defined in TC-01 for a duration of four hours.

2. Configure SnapCenter to initiate a full, application-consistent backup of the SQL Server database at hourly intervals. Configure the retention interval to four backup versions requiring every fifth backup to remove the oldest backup, causing the deletion of the storage-based Snapshot copy associated with the backup.
3. Record performance monitor data, sysstat, and OPM data over the course of the four-hour run to get a picture of the performance at both the database and storage levels.

TC-03: Measure Performance Impact of Creating Clones of the SQL Server Database Initiated from SnapCenter 1.1 on AFBP Baseline Workload

| Test Case | Details |
|--|--|
| Test number | TC-03 |
| Test case description | <p>While subjected to the AFBP baseline workload computed as a result of TC-01 for a total of 30 minutes, configure NetApp SnapCenter 1.1 to create a clone of the SQL Server database using one of the existing application consistent backup copies created in TC-02.</p> <p>Observe the impact on overall AFBP baseline performance during the clone creation to make sure that conducting normal cloning operations using SQL Server database backups did not significantly impact database performance.</p> |
| Test assumptions | <ul style="list-style-type: none"> • A completely operational NetApp cluster running clustered NetApp ONTAP 9 has been installed and configured within the AFBP guidelines. • A completely operational SQL Server environment has been installed and configured. • The TPC-E utility has been installed and configured to generate the following workloads: <ul style="list-style-type: none"> – 90% reads and 10% writes with a 100% random access pattern • SnapCenter 1.1 was installed into the environment and configured to successfully create application-consistent backups of the SQL Server database using the AFF8080 storage configuration. |
| Test data or metrics to capture | <ul style="list-style-type: none"> • Microsoft performance monitor data for SQL Server • Sysstat output from all NetApp AFF8080 controllers in the cluster • OPM performance capacity data for the duration of the test |
| Expected results | Baseline performance in this AFBP environment is minimally affected by the SnapCenter clone creation process, which includes creating a series of FlexClone® volumes supporting the SQL Server database. |

Test Methodology

1. Using the TPC-E workload utility, initiate the AFBP baseline workload defined in TC-01 for a duration of 30 minutes.
2. After running for five minutes, use SnapCenter to create a clone of the SQL Server database using one of the four existing full backups created in previous testing.
3. Record the time required to complete the clone creation process.
4. Record performance monitor data, sysstat, and OPM data over the course of the 30-minute run to get a picture of the performance at both the database and storage levels.

TC-04: Measure Performance Impact of Deleting Clones of the SQL Server Database Initiated from SnapCenter 1.1 on AFBP Baseline Workload

| Test Case | Details |
|--|--|
| Test number | TC-04 |
| Test case description | While subjected to the AFBP baseline workload computed as a result of TC-01 for a total of 30 minutes, configure NetApp SnapCenter 1.1 to delete an existing clone of the SQL Server database. The database clone selected was created in TC-05. Observe the impact on overall AFBP baseline performance during the clone deletion to make sure that conducting normal cloning operations using SQL Server database backups did not significantly impact database performance. |
| Test assumptions | <ul style="list-style-type: none"> • A completely operational NetApp cluster running clustered NetApp ONTAP 9 has been installed and configured within the AFBP guidelines. • A completely operational SQL Server environment has been installed and configured. • The TPC-E utility has been installed and configured to generate the following workloads: <ul style="list-style-type: none"> – 90% reads and 10% writes with a 100% random access pattern • SnapCenter 1.1 was installed into the environment and configured to successfully create application-consistent backups of the SQL Server database using the AFF8080 storage configuration. |
| Test data or metrics to capture | <ul style="list-style-type: none"> • Microsoft performance monitor data for SQL Server • Sysstat output from all NetApp AFF8080 controllers in the cluster • OPM performance capacity data for the duration of the test |
| Expected results | Baseline performance in this AFBP environment is minimally impacted by the SnapCenter clone deletion process |

Test Methodology

1. Using the TPC-E workload utility, initiate the AFBP baseline workload defined in TC-01 for a duration of 30 minutes.
2. After running for five minutes, use SnapCenter to delete the clone of the SQL Server database that was initially created in TC-05.
3. Record the time required to complete the clone deletion process.
4. Record performance monitor data, sysstat, and OPM data over the course of the 30-minute run to get a picture of the performance at both the database and storage levels.

TC-05: Performance and Availability Impact of Unplanned Failover Events: Storage Controller Panic

| Test Case | Details |
|------------------------------|--|
| Test number | TC-05 |
| Test case description | <p>Measure the impact on database availability and performance when one storage node in the HA pair is forced into a panic situation resulting in storage node failover to the remaining storage controller in the HA pair.</p> <p>Additionally, measure the impact on performance after correcting the issues that caused the storage failover after a successful giveback has occurred restoring the overall environment to the pre-failure condition.</p> <p>This test conducted a series of 10 identical storage failovers and givebacks over the 10-hour duration of the testing. The SQL Server database continued to be fully backed up as described in TC-02 every hour.</p> |

| Test Case | Details |
|--|---|
| Test assumptions | <ul style="list-style-type: none"> • A completely operational NetApp cluster running clustered NetApp ONTAP 9 has been installed and configured within the AFBP guidelines. • A completely operational SQL Server environment has been installed and configured. • SnapCenter 1.1 was installed into the environment and configured to successfully create application consistent backups of the SQL Server database using the AFF8080 storage configuration each hour. • The TPC-E utility has been installed and configured to generate the following workloads: <ul style="list-style-type: none"> – 90% reads and 10% writes with a 100% random access pattern |
| Test data or metrics to capture | <ul style="list-style-type: none"> • Microsoft performance monitor data for SQL Server • Sysstat output from all NetApp AFF8080 controllers in the cluster • OPM performance capacity data for the duration of the test |
| Expected results | <p>An unplanned storage failover event that results from a storage controller panic will behave as follows:</p> <ul style="list-style-type: none"> • Have minimal impact on database availability. • Performance during the failover event is minimally affected and continues to be observed operating below the maximum storage performance capacity threshold as reported by OPM. • Performance returns to prefailover levels after restoring the configuration to its prefailure condition with observed levels of performance capacity back with the AFBP guidelines. • No impact on the amount of time required for SQL Server database backups over the course of the testing. |

Test Methodology

1. Initiate the baseline AFBP workload consisting of 90% reads and 10% writes with a 100% random access pattern by using the TPC-E tool for a total of 10 hours.
2. At hourly intervals over the course of 10 hours, induce a panic in one of the AFF8080 storage controllers in the HA pair, causing a takeover of the failed node by the surviving AFF8080 storage controller. This results in the surviving AFF8080 storage controller handling the entire AFBP baseline workload.
3. Allow the surviving storage controller to service the workload during the time the other storage controller is down.
4. Using the standard timeout values, allow the storage to perform a giveback of the failed storage controller to automatically restore the storage to the prefailure state with both controllers in the HA pair equally handling the load.
5. Record performance monitor data, sysstat, and OPM data over the course of the 10-hour run to get a picture of the performance at both the database and storage levels.

TC-06: Measure Performance and Availability Impact of Planned Failover Events

| Test Case | Details |
|------------------------------|---|
| Test number | TC-06 |
| Test case description | Measure the impact on database availability and performance when one storage node in the HA pair is removed from the configuration in a planned maintenance |

| Test Case | Details |
|--|--|
| | <p>window resulting in storage node failover to the remaining storage controller in the HA pair.</p> <p>Additionally, measure the impact on performance after correcting the issues that caused the storage failover after a successful giveback has occurred restoring the overall environment to the prefailure condition.</p> <p>This test conducted a series of 10 identical storage failovers and givebacks over the 10-hour duration of the testing. The SQL Server database continued to be fully backed up as described in TC-02 every hour.</p> |
| Test assumptions | <ul style="list-style-type: none"> • A completely operational NetApp cluster running clustered NetApp ONTAP 9 has been installed and configured within the AFBP guidelines. • A completely operational SQL Server environment has been installed and configured. • SnapCenter 1.1 was installed into the environment and configured to successfully create application-consistent backups of the SQL Server database using the AFF8080 storage configuration each hour. • The TPC-E utility has been installed and configured to generate the following workloads: <ul style="list-style-type: none"> – 90% reads and 10% writes with a 100% random access pattern |
| Test data or metrics to capture | <ul style="list-style-type: none"> • Microsoft performance monitor data for SQL Server • Sysstat output from all NetApp AFF8080 controllers in the cluster • OPM performance capacity data for the duration of the test |
| Expected results | <p>A planned storage failover event that results from a storage controller panic behaves as follows:</p> <ul style="list-style-type: none"> • Have minimal impact on database availability. • Performance during the failover event is minimally affected and continue to be observed operating below the maximum storage performance capacity threshold as reported by OPM. • Performance returns to prefailover levels after restoring the configuration to its prefailure condition with observed levels of performance capacity back with the AFBP guidelines. • No impact on the amount of time required for SQL Server database backups over the course of the testing. |

Test Methodology

1. Initiate the baseline AFBP workload consisting of 90% reads and 10% writes with a 100% random access pattern by using the TPC-E tool for a total of 10 hours.
2. At hourly intervals over the course of 10 hours, use the standard storage takeover and giveback commands to initiate a planned failure of one of the AFF8080 storage controllers in the HA pair, causing a the surviving AFF8080 storage controller to handle the entire AFBP baseline workload.
3. Allow the surviving storage controller to service the workload during the time the other storage controller is down.
4. Using the standard timeout values, allow the storage to perform a giveback of the failed storage controller to automatically restore the storage to the prefailure state with both controllers in the HA pair equally handling the load.
5. Record performance monitor data, sysstat, and OPM data over the course of the 10-hour run to get a picture of the performance at both the database and storage levels.

Appendix B: Deployment Details

This appendix provides detailed information, in Table 13 through Table 16, about the deployment details of the architecture described in this document.

Table 13 lists the storage configuration and layout, including aggregate, volume, LUN sizes, and mappings.

Table 13) Storage configuration and layout.

| Storage Node | Aggregate Name | Volume Name | Volume Size | LUN Size | Description |
|----------------|------------------------|----------------------|-------------|---------------------------------------|---|
| nst-fas8080-15 | | | | | Used advanced drive partitioning |
| | aggr0_nst_fas8080_15_0 | vol0 | 348.6GB | | Total aggr size = 368.42GB |
| | aggr_15 | | | | 42 x 400GB SSDs (~12.07TB TB) |
| | | blade1vm1dbvol | 3TB | 2.6TB | Blade 1 VM 1 MS-SQL TPCE database LUN |
| | | blade1vm1logvol | 500GB | 480GB | Blade 1 VM 1 MS-SQL TPCE log LUN |
| | | blade1vm1snapinfovol | 2TB | 1.8TB | Blade 1 VM 1 LUN for SnapCenter log directory |
| | | blade1vm1tempvol | 500GB | 480GB | Blade 1 VM 1 MS-SQL tempDB LUN |
| | | blade2vm1dbvol | 3TB | 2.6TB | Blade 2 VM 1 MS-SQL TPCE database LUN |
| | | blade2vm1logvol | 500GB | 480GB | Blade 2 VM 1 MS-SQL TPCE Log LUN |
| | | blade2vm1snapinfovol | 2TB | 1.8TB | Blade 2 VM 1 LUN for SnapCenter log directory |
| | | blade2vm1tempvol | 500GB | 480GB | Blade 2 VM 1 MS-SQL tempDB LUN |
| | | blade3vm1dbvol | 3TB | 2.6TB | Blade 3 VM 1 MS-SQL TPCE database LUN |
| | | blade3vm1logvol | 500GB | 480GB | Blade 3 VM 1 MS-SQL TPCE Log LUN |
| | | blade3vm1snapinfovol | 2TB | 1.8TB | Blade 3 VM 1 LUN for SnapCenter log directory |
| | | blade3vm1tempvol | 500GB | 480GB | Blade 3 VM 1 MS-SQL tempDB LUN |
| | blade4vm1dbvol | 3TB | 2.6TB | Blade 1 VM 1 MS-SQL TPCE database LUN | |

| Storage Node | Aggregate Name | Volume Name | Volume Size | LUN Size | Description |
|----------------|----------------------|-----------------------|-------------|--------------------------------|---|
| | | blade4vm1 logvol | 500GB | 480GB | Blade 1 VM 1 MS-SQL TPCE log LUN |
| | | blade4vm1 snapinfovol | 2TB | 1.8TB | Blade 1 VM 1 LUN for SnapCenter log directory |
| | | blade4vm1 tempvol | 500GB | 480GB | Blade 1 VM 1 MS-SQL tempDB LUN |
| | | sdwroot | 1GB | NA | SVM root volume |
| | nst_fas8080_15_aggr1 | | | | Total aggr size = 946.9GB |
| nst-fas8080-16 | | | | | Used advanced drive partitioning |
| | | vol0 | 348.6GB | | Total aggr size = 368.42GB |
| | aggr_16 | | | | 42 x 400GB SSDs (~12.07TB TB) |
| | | Blade1vm2 dbvol | 3TB | 2.6TB | Blade 1 VM 2 MS-SQL TPCE database LUN |
| | | Blade1vm2 logvol | 500GB | 480GB | Blade 1 VM 2 MS-SQL TPCE log LUN |
| | | Blade1vm2 snapinfovol | 2TB | 1.8TB | Blade 1 VM 2 LUN for SnapCenter log directory |
| | | Blade1vm2 tempvol | 500GB | 480GB | Blade 1 VM 2 MS-SQL tempDB LUN |
| | | blade2vm2 dbvol | 3TB | 2.6TB | Blade 2 VM 2 MS-SQL TPCE database LUN |
| | | blade2vm2 logvol | 500GB | 480GB | Blade 2 VM 2 MS-SQL TPCE log LUN |
| | | blade2vm2 snapinfovol | 2TB | 1.8TB | Blade 2 VM 2 LUN for SnapCenter log directory |
| | | blade2vm2 tempvol | 500GB | 480GB | Blade 2 VM 2 MS-SQL tempDB LUN |
| | | Blade3vm2 dbvol | 3TB | 2.6TB | Blade 3 VM 2 MS-SQL TPCE database LUN |
| | | Blade3vm2 logvol | 500GB | 480GB | Blade 3 VM 2 MS-SQL TPCE log LUN |
| | | Blade3vm2 snapinfovol | 2TB | 1.8TB | Blade 3 VM 2 LUN for SnapCenter Log Directory |
| | Blade3vm2 tempvol | 500GB | 480GB | Blade 3 VM 2 MS-SQL tempDB LUN | |

| Storage Node | Aggregate Name | Volume Name | Volume Size | LUN Size | Description |
|--------------|----------------|-----------------------|-------------|----------|---|
| | | Blade4vm2 dbvol | 3TB | 2.6TB | Blade 4 VM 2 MS-SQL TPCE database LUN |
| | | Blade4vm2 logvol | 500GB | 480GB | Blade 4 VM 2 MS-SQL TPCE log LUN |
| | | Blade4vm2 snapinfovol | 2TB | 1.8TB | Blade 4 VM 2 LUN for SnapCenter log directory |
| | | Blade4vm2 tempvol | 500GB | 480GB | Blade 4 VM 2 MS-SQL tempDB LUN |

Tables 14 through 16 list the network details.

Table 14) Server network specifications.

| Host Name | Interface | IP Address | Speed | MTU | Purpose |
|-----------|-----------|----------------|--------|------|---------|
| sqlblade1 | eth0 | 10.231.149.151 | 10Gbps | 1500 | Public |
| sqlblade2 | eth0 | 10.231.149.152 | 10Gbps | 1500 | Public |
| sqlblade3 | eth0 | 10.231.149.153 | 10Gbps | 1500 | Public |
| sqlblade4 | eth0 | 10.231.149.154 | 10Gbps | 1500 | Public |
| sqlblade5 | eth0 | 10.231.149.155 | 10Gbps | 1500 | Public |
| sqlblade6 | eth0 | 10.231.149.156 | 10Gbps | 1500 | Public |
| sqlblade7 | eth0 | 10.231.149.157 | 10Gbps | 1500 | Public |
| sqlblade8 | eth0 | 10.231.149.158 | 10Gbps | 1500 | Public |

Table 15) Storage network specifications.

| SVM | LIF | Port | IP Address | Speed | MTU | Role |
|-------------------|----------------------|--------------------|----------------|-------|-------|--------------------|
| Cluster | nst-fas8080-15_clus1 | nst-fas8080-15:e0a | 10.226.227.63 | 10Gb | 9,000 | Cluster |
| Cluster | nst-fas8080-15_clus2 | nst-fas8080-15:e0c | 10.226.227.64 | 10Gb | 9,000 | Cluster |
| Cluster | nst-fas8080-16_clus1 | nst-fas8080-16:e0a | 10.226.227.65 | 10Gb | 9,000 | Cluster |
| Cluster | nst-fas8080-16_clus2 | nst-fas8080-16:e0c | 10.226.227.66 | 10Gb | 9,000 | Cluster |
| nst-fas8080-15-16 | cluster_mgmt | nst-fas8080-15:e0i | 10.233.181.29 | 1Gb | 1,500 | Cluster management |
| nst-fas8080-15-16 | nst-fas8080-15_mgmt1 | nst-fas8080-15:e0M | 10.233.180.208 | 1Gb | 1,500 | Node management |

| SVM | LIF | Port | IP Address | Speed | MTU | Role |
|-------------------|----------------------|--------------------|----------------|-------|-------|-----------------|
| nst-fas8080-15-16 | nst-fas8080-16_mgmt1 | nst-fas8080-16:e0M | 10.233.180.208 | 1Gb | 1,500 | Node management |
| SQL_SVM | SNAPCENTERMANAGEMENT | nst-fas8080-15:e0d | 10.233.181.26 | 10Gb | 1,500 | Data |

Table 16) FC back-end switches.

| Host Name | IP Address |
|---------------|--------------|
| ste-nb-6510-1 | 10.226.26.29 |

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-4532-0916