

Network Appliance™ FAS3020 and EMC CX500: Comparison of Usability and Performance

Test report prepared under contract from Network Appliance

Executive summary

Network Appliance™ commissioned VeriTest, a division of Lionbridge Technologies Inc., to compare the usability of a variety of features and functionality that users of both the NetApp® FAS3020 and the EMC CLARiiON CX500 encounter in day to day operations to solve typical problems related to storage provisioning. Additionally, we compared the performance of the NetApp FAS3020 and the EMC CX500 storage server in a variety of configurations including while performing a series of snapshot operations over the course of a one hour interval while the products were under load.

The CX500 configuration used in this test contained two storage processors. To ensure a valid comparison during testing, the FAS3020 used in this test was configured with two storage controllers in an active/active network failover configuration.

For the usability testing, we compared the features and functionality that users of these technologies encounter in day to day operations to solve typical problems related to storage provisioning, data backup and restoration and cloning. We looked at the ease with which these features are utilized from an administrative perspective as well as the general usefulness of the related documentation including both online and hard copy. Additionally, we measured the time required to perform a set of typical administrative tasks using the features provided by both the FAS3020 and CX500 products.

All of the performance tests used an OLTP-like workload that online database applications typically encounter. Like many other industry standard OLTP benchmarks, including the majority of the top 10 TPC-C price/performance results, we used an 8KB request size with a mix of random read and write operations for both the FAS3020 and the CX500. For both the FAS3020 and CX500 products, we used the industry-standard, open-source load generator IOMeter, available from Source Forge at

Key findings

- ❑ In our test configurations, we found that using NetApp FlexVol™ technology allowed us to use 26 fewer physical disk drives on the FAS3020 to provision a sample enterprise class database system compared to the CX500.
- ❑ In our test configurations, with the exception of the snapshot creation and restoration, we found it required significantly less time to complete a series of typical provisioning and administrative tasks on the FAS3020 compared to the CX500.
- ❑ In our test configurations, we found that using the snapshot process with the CX500 during performance testing resulted in a sustained, 60 percent drop in overall performance. We found no sustained degradation in overall performance as a result of using the Snapshot™ technology available in the FAS3020.
- ❑ In test configurations using a provisioned 100GB LUN and the FCP protocol, we found that the FAS3020 configured with RAID-DP generated 5.3 times the performance of the CX500 when using RAID 1/0. When the same test was run using the iSCSI protocol on the FAS3020, NetApp's performance advantage was 3.8 times that of the CX500 with FCP.

<http://sourceforge.net/projects/iometer/> to generate the workload for the performance testing. Please refer to the Testing Methodology section of this report for complete details on how we conducted both the usability and performance testing on the FAS3020 and CX500.

Figure 1 below provides a summary of the results from the usability and provisioning testing. For this test we created 20 separate LUNs connected to two servers for use in a corporate environment consisting of multiple departments and multiple databases. We used publicly available best practices documentation from both NetApp and EMC to create a plan to provision both the FAS3020 and the CX500 for use in a simulated corporate environment consisting of multiple departments and multiple databases. Specifically, we used the following documentation from both NetApp and EMC:

- NetApp : Block Management with Data ONTAP™ 7G: FlexVol™, FlexClone™, and Space Guarantees(http://www.netapp.com/tech_library/ftp/3348.pdf)
- EMC : Configuration Planning Guide - P/N 300-001-273 - REV A02

The table shows the elapsed time required to complete a specific task using both the FAS3020 and the CX500. In general, we found that the number of steps required to perform the tests was comparable between the FAS3020 and the CX500. However, with the exception of the snapshot creation and restoration, we found performing these tasks on the FAS3020 required considerably less time compared to the CX500. Additionally, we found that using FlexVol technology on the FAS3020 allowed us to provision an enterprise class database environment using 26 fewer physical disk drives compared to the CX500 using a combination of RAID 1/0 and RAID 5. Additionally, the amount of time listed in Figure 1 to extend the size of the Oracle OLTP Database on the CX500 was calculated using the stripe expansion method.




Figure 1. Usability and Provisioning Test Results Summary

During the performance tests, we used the following configurations:

- Against both the FAS3020 and CX500 using the single 100GB OLTP database LUN defined in Provisioning Test Case #2 and created during the provisioning tests. Please refer to the Testing Methodology section of this report for specifics of how conducted the provisioning test.
- The FAS3020 configuration populated with 112 physical disk drives each of which was 36GB in size and 15,000 RPM. We tested the FAS3020 using RAID-DP with both FCP and iSCSI host

connections. RAID-DP™ (Double Parity) uses two parity disks per RAID group to decrease the likelihood that a double disk failure will cause data loss.

- The CX500 configuration populated with 112 physical disk drives each of which was 36GB in size and 15,000 RPM. We tested the CX500 using only FCP with RAID 5 and RAID 1/0.

At the time we started testing in early February, we found that 36GB/15,000 RPM Fibre Channel disk drives were the most cost effective and highest performing disk drives for both the FAS3020 and CX500. As a result, we used these drives for all tests not involving Serial ATA (SATA) disk drives.

Figure 2 below shows the results of the performance testing in IOPS when using the provisioned 100GB OLTP database LUN created during the provisioning tests described above for both the FAS3020 and the CX500. For these tests, we used the following configurations:

- NetApp FAS3020 using Fibre Channel host attach with Fibre Channel disk drives
- NetApp FAS3020 using iSCSI host attach with Fibre Channel disk drives
- NetApp FAS3020 using Fibre Channel host attach with SATA disk drives
- EMC CX500 using Fibre Channel and RAID 1/0 with Fibre Channel disk drives

For this test, we used an OLTP workload consisting of 60% random read and 40% random write operations all using an 8KB request size. In this test, we found that the FAS3020 using the Fibre Channel protocol generated 5.3 times the performance (9229 vs. 1728) compared to the CX500 when using the Fibre Channel protocol. Additionally, we found that the FAS3020 using iSCSI generated 3.8 times the performance (6565 vs. 1728) compared to the CX500 when using the Fibre Channel protocol. Finally, we found that the FAS3020 using the Fibre Channel protocol and 7200 RPM SATA drives generated 1.5 times the performance (2686 vs. 1728) compared to the CX500 when using the Fibre Channel protocol.




Figure 2: Test Results for Provisioned Performance Using Provisioned 100GB OLTP LUN

Figure 3 below shows the results of the snapshot performance testing. These results compare the relative performance generated during the snapshot performance test when conducting the snapshot process to the performance generated using the same test script when not conducting the snapshot process. A value of 100 indicates no difference in performance between the two cases while values less than 100 indicate that the performance noted during the test utilizing the snapshot process was less than observed during the baseline test where no snapshot process was performed.

For this test, we used an IOMeter test script that generated an OLTP workload containing a mixture of 60% random reads and 40% random writes using an 8KB request size. We set the run time in the IOMeter test script to one hour and let the test run continuously. For the FAS3020 we then created a series of 30 snapshot copies of the provisioned 100GB OLTP production database LUN at 2 minute intervals during the course of the 60 minute test run. Because the CX500 has a restriction of allowing only a total of 8 snapshot copies for any given LUN, we created a series of 8 snapshot copies of the provisioned 100GB OLTP production database LUN at 5 minute intervals and then allowed the test to run for the last 20 minutes with no additional snapshot copies created.

During the 60 minute duration of the test, we ran the Performance Monitor application from Microsoft on the host system running Windows Server 2003 to measure the read and write activity on the logical volumes being accessed on both the FAS3020 and the CX500 during the testing. By monitoring host activity while taking the snapshot copies, we were able to determine the impact on the overall performance as a result of the snapshot process over the entire one hour test. This was not possible using IOMeter alone as it reports only a single average IOPS metric calculated over the entire test run time.

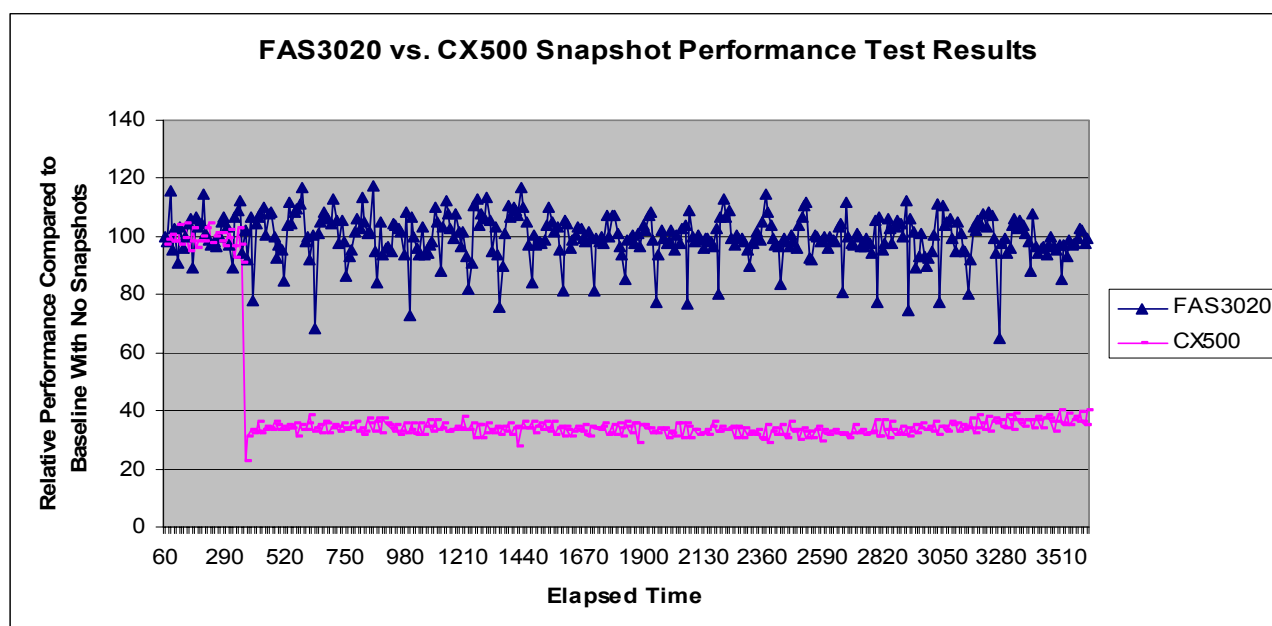


Figure 3: Test Results for Snapshot Performance Tests Using Provisioned 100GB OLTP LUN

In our test configurations, we found that conducting the 30 snapshot copies on the FAS3020 over the course of the test had no sustained impact on the overall performance compared to the test case where no snapshot process was performed. When using the FAS3020, we observed brief periods where performance dropped to approximately 80 percent of that generated without the snapshot process. These periods coincided with the taking of the specific snapshot. However, the overall performance always recovered back to the levels generated using the baseline configuration.

For the CX500, we observed a drop in performance of approximately 60 percent after taking the first of the eight snapshot copies during the testing. This is indicated by the drop in figure 3 at approximately the 300 second point in the test. At this point, the CX500 was generating approximately 40 percent of the pre-snapshot performance. During the remainder of the testing, we conducted an additional seven snapshot copies at 5 minute intervals. During these remaining snapshot copies, we observed no additional drops in the performance of the CX500. However, the overall performance recorded on the CX500 between the snapshot copies did not recover to pre-snapshot levels.

After conducting the last of the eight snapshot copies, we allowed the test to continue against the CX500 for an additional 20 minutes. During this time, the overall performance of the CX500 begins to recover but does not come back to the level of performance observed before we began the snapshot process.

Additionally, we found that using aggregates and flexible volumes on the FAS3020 allowed us to provision an enterprise class database environment using 26 fewer physical disk drives compared to the CX500 using recommended configurations with RAID 1/0 and RAID 5. Figures 5 and 6 below show the specifics of how the various database and log files were laid out on both the FAS3020 and the CX500.

On the FAS3020, we created a single aggregate on each of the two storage controllers each containing a total of 28 physical disks drives. In each aggregate we created two RAID-DP groups of 14 total disk drives for a total of 680 GB of usable space on each of the two storage controllers. We then created a single flexible volume under each of the aggregates that consumed the maximum amount of available space while maintaining a 20 percent space reserve for snapshot copies. The end result was a total of two flexible volumes each with a total of 544GB of usable storage.

After creating the flexible volumes as described above, we created the LUNs associated with the ACME Company database environment on the FAS3020 as shown in figure 5 below. For details of the database environment used by ACME Company, please refer to the Testing Methodology section of this report.

We found that by using the FlexVol™ technology of the FAS3020, we were able to provision and deploy the database and log files for this test using a total of 56 physical disk drives compared to the CX500 which required 82 physical drives to provision and deploy the identical database layout.

LUN Description	Storage Controller #	RAID Level	LUN Size
Exchange Logs	1	DP	25
Exchange DB	2	DP	100
HR QA DB	2	DP	50
HR QA Logs	1	DP	10
Oracle Prod DB	2	DP	100
Oracle Prod Logs	1	DP	25
Payroll Development DB	1	DP	50
Payroll Development Log	2	DP	10
Oracle Development DB	2	DP	100
Oracle Development Logs	1	DP	25
Oracle QA DB	1	DP	100
Oracle QA Logs	2	DP	25
Payroll Prod DB	2	DP	50
Payroll Prod Log	1	DP	10
HR Development DB	1	DP	50
HR Development Logs	2	DP	10
HR Prod DB	1	DP	50
HR Prod Logs	2	DP	10
Payroll QA DB	1	DP	50
Payroll QA Log	2	DP	10
Totals Size of DB and Log LUNs			860GB

Figure 5. FAS3020 Database and Log File LUN Layout

When conducting this exercise on the CX500, we utilized available EMC documentation and best practices guides to ensure an optimal layout. These documents contained detailed worksheets to help with planning, provisioning and sizing a deployment similar to the one we completed for this testing. Figure 6 below shows how we configured the CX500 to handle the ACME database layout for all database and log volumes while maintaining a 20 percent space reserve for snapshot copies.

As figure 6 shows, we found it required a total of 82 physical disk drives to deploy the ACME Company database environment on the CX500 compared to the 56 physical disk drives required for the same deployment performed on the FAS3020.

Applications	SP	LUN Size in GB	Total Used Space (GB)	RAID Group	RAID Level	BUS	# of Physical Disks	Approx Raw Space (GB)
Exchange Logs	A	25	31	RAID Group 0	5	0	5	165
HR QA Logs	A	10	12.5					
Oracle Development Logs	A	25	31					
Oracle QA Logs	A	25	31					
Payroll QA Log	A	10	12.5					
HR Development Logs	B	10	12.5	RAID Group 1	5	1	5	165
HR Prod Logs	B	10	12.5					
Oracle Prod Logs	B	25	31					
Payroll Development Log	B	10	12.5					
Payroll Prod Log	B	10	12.5					
Exchange DB	B	100	200	RAID Group 2	1/0	1	8	264
Oracle Prod DB	A	100	200	RAID Group 3	1/0	0	8	264
Oracle QA DB	A	100	200	RAID Group 4	1/0	0	8	264
Oracle Development DB	B	100	200	RAID Group 5	1/0	1	8	264
HR QA DB	A	50	100	RAID Group 6	1/0	0	10	330
Payroll Development DB	A	50	100					
Payroll Prod DB	A	50	100					
HR Development DB	B	50	100	RAID Group 7	1/0	1	10	330
HR Prod DB	B	50	100					
Payroll QA DB	B	50	100					
Snapshot Reserve	A	N/A	N/A	RAID Group 8	1/0	0	10	330
Snapshot Reserve	B	N/A	N/A	RAID Group 9	1/0	1	10	330
Totals		860	1599				82	2706

Figure 6: CX500 Database Provisioning Details

Figure 7 below summarizes the data shown in figures 5 and 6. These results show that to deploy the 860GB of database and log LUNs required by the ACME Company required a total of 56 physical disk drives and on the FAS3020 compared to 82 physical disk drives on the CX500. The CX500 therefore required just over 46% more disk drives for a best-practices configuration matching this common enterprise workload.

Test Configuration	Total Space Required Space for LUNs	Total Disk Drives Required
FAS3020	860 GB	56
CX500	860 GB	82

Figure 7. FAS3020 vs. CX500 ACME Company Database Provisioning Space Summary

Performance Testing Results

Performance Test Case #1: Performance Using the OLTP Production Database LUN

This section provides the results of the performance testing conducted against the OLTP database LUN as provisioned during the ACME Company database provisioning exercise described in this test report on both the FAS3020 and the CX500. For these tests, we used the following configurations:

- NetApp FAS3020 using Fibre Channel host attach with Fibre Channel disk drives
- NetApp FAS3020 using iSCSI host attach with Fibre Channel disk drives
- NetApp FAS3020 using Fibre Channel host attach with SATA disk drives
- EMC CX500 using Fibre Channel host attach and RAID 1/0 with Fibre Channel disk drives

Figure 8 below shows the results of the performance testing in IOPS for both the FAS3020 and the CX500 when using the provisioned 100GB OLTP database LUN. For this test, we used an OLTP workload consisting of 60% random read and 40% random write operations all using an 8KB request size. In this test, we found that the FAS3020 using the Fibre Channel protocol generated 5.3 times the performance (9229 vs. 1728) compared to the CX500 when using the Fibre Channel protocol. Additionally, we found that the FAS3020 using iSCSI generated 3.8 times the performance (6565 vs. 1728) compared to the CX500 when using the Fibre Channel protocol. Finally, we found that the FAS3020 using the Fibre Channel protocol and 7200 RPM SATA drives generated 1.5 times the performance (2686 vs. 1728) compared to the CX500 when using the Fibre Channel protocol.

In addition to performance in IOPS, figure 8 also shows the average latency in milliseconds reported by IOMeter for each of the test configurations for both the FAS3020 and the CX500. In this test configuration, we found that the FAS3020 using the Fibre Channel protocol reported a 77 percent lower average latency (28 vs. 120 milliseconds) compared to the CX500 when using the Fibre Channel protocol. Additionally, we found that the FAS3020 using iSCSI reported a 67 percent lower average latency (39 vs. 120 milliseconds) compared to the CX500 when using the Fibre Channel protocol. Finally, we found that the FAS3020 using the Fibre Channel protocol and 7200 RPM SATA drives reported a 4 percent lower average latency (115 vs. 120 milliseconds) compared to the CX500 using the Fibre Channel protocol.




Figure 8. Average IOPS/Latency Results for Performance Testing Using the Oracle OLTP Database LUN

Performance Test Case #2: Measure Relative Performance Impact When Taking Successive Snapshot Copies of the Provisioned OLTP Production LUN

This section provides the details of the snapshot performance testing conducted on both the FAS3020 and the CX500. For this test, we used an IOMeter test script that generated an OLTP-like workload containing a mixture of 60% random reads and 40% random writes using an 8KB request size size. We measured the relative performance impact of creating a series of snapshot copies of the 100GB LUN representing the OLTP production database created during the provisioning tests. For this test, we set the run time in the IOMeter test script to one hour and let the test run continuously. For the FAS3020 we created a series of 30 snapshot copies at 2 minute intervals during the course of the 60 minute test run. Because the CX500 is limited to a total of 8 snapshot copies for any given LUN, we created a series of 8 snapshot copies at 5 minute intervals and then allowed the test to run for the last 20 minutes with no additional snapshot copies created.

During the 60 minute duration of the test, we also ran the Performance Monitor application from Microsoft on the host system running Windows Server 2003 to measure the read and write activity on the logical volumes being accessed on both the FAS3020 and the CX500 during the testing. We configured Performance Monitor to capture information related to the following logical disk counters at 10 second intervals during the testing:

- Number of disk reads per second
- Number of disk writes per second
- Average read latency in seconds per operation
- Average write latency in seconds per operation

By monitoring these counters while taking the snapshot copies, we were able to determine the impact on the overall performance as a result of the snapshot process over the entire one hour test. This was not possible using IOMeter alone as it reports only a single average IOPS metric calculated over the entire test run time.

To compute the results presented for this test, we recorded the total IOPS values generated at each of the 10 second intervals using Performance Monitor over the course of the test when not performing the snapshot process and used these values as our baseline. This baseline represented the overall performance of the test configuration when not performing the snapshot process.

We then recorded the total IOPS values generated at each of the 10 second intervals using Performance Monitor over the course of the test when performing the snapshot process. This data represented the overall performance of the test configuration when the snapshot process was performed. At each data point, we computed the difference in the number of IOPS between the baseline configuration using no snapshot process and the configuration where we conducted the snapshot process. We calculated the difference as a percentage of the baseline value to see how the performance was impacted over the course of the testing as a result of the snapshot process.

Figure 9 below shows the results of this testing. The chart compares the relative performance generated during the test when performing the snapshot process to the performance generated when not performing the snapshot process. A value of 100 indicates the no difference in performance between the two cases. In our test configurations, we found that conducting the 30 snapshot copies on the FAS3020 over the course of the test had no sustained impact on the overall performance compared to the test case where no snapshot process was performed. When using the FAS3020, we observed brief periods where performance dropped to approximately 80 percent of that generated without the snapshot process. These periods coincided with the taking of the specific snapshot but the overall performance always recovered back to the levels generated using the baseline configuration.

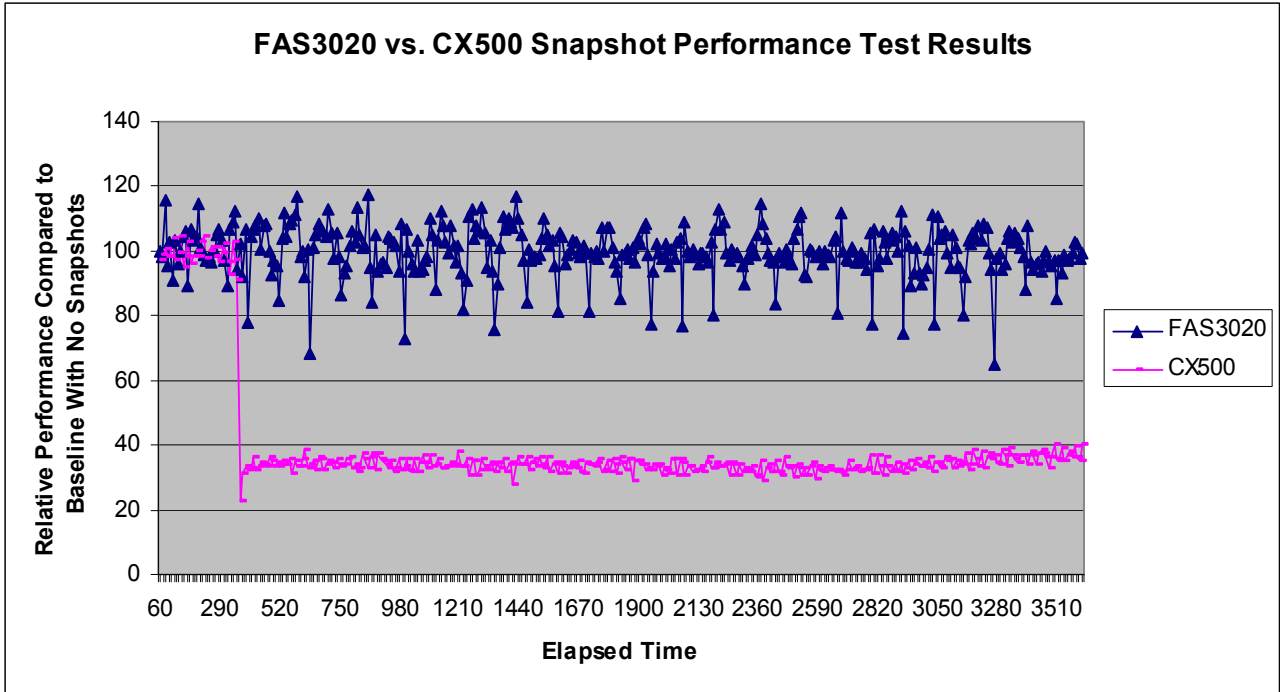


Figure 9. Snapshot Performance Results

For the CX500, we observed a drop in performance of approximately 60 percent after taking the first of the eight snapshot copies during the testing. This is indicated by the drop in figure 9 at approximately 300 seconds into the test. At this point, the CX500 was generating approximately 40 percent of the pre-snapshot performance. During the remainder of the testing, we conducted an additional seven snapshot copies at 5 minute intervals. During these remaining snapshot copies, we observed no additional significant drops in the performance of the CX500. Unlike the FAS3020, the overall performance recorded on the CX500 between the snapshot copies did not recover to pre-snapshot levels.

After conducting the last of the eight snapshot copies, we allowed the test to continue against the CX500 for an additional 20 minutes. During this time, the overall performance of the CX500 begins to recover but does not come back to the level of performance observed before we began the snapshot process.

Performance Test Cases #3-5: Measure Performance of the FAS3020 and CX500 Using 112 Drives

This section provides the results for the performance testing on both the FAS3020 and CX500 using 112 drives with an OLTP workload containing a mixture of 60% random reads and 40% random writes using an 8KB request size size. For the FAS3020 we tested only RAID-DP configurations using both FCP and iSCSI host connections. For the CX500 we tested only FCP using both RAID 5 and RAID 1/0. In February of 2005, EMC announced iSCSI support for the CX500. When we conducted our testing in February of 2005, we attempted to acquire iSCSI support for the CX500 but found out it was not yet available. As a result, we did not conduct tests using iSCSI on the CX500. Please refer to the Testing Methodology section of this report for complete details on how we conducted these tests.

Figure 10 below shows the performance test results generated when both the FAS3020 and the CX500 were configured with 112 disk drives. In this configuration, we found that the FAS3020 configured with RAID-DP generated virtually identical performance measured in IOPS (15,904 vs. 15,722) compared to the CX500 configured with RAID 1/0. Additionally, we found that the FAS3020 configured with RAID-DP generated 16 percent better performance measured in IOPS (15,904 vs. 13,708) compared to the CX500 configured with RAID 5. Finally, we found that the FAS3020 configured with RAID-DP and using iSCSI host connections generated 22 percent lower performance measured in IOPS (12,377 vs. 15,904) compared to the FAS3020 configured with RAID-DP using the Fibre Channel protocol.

In addition to performance in IOPS, figure 10 also shows the average latency in milliseconds reported by IOMeter for each of the test configurations for both the FAS3020 and the CX500. In this test configuration, we found that the FAS3020 using the Fibre Channel protocol reported virtually identical average latency (64 vs. 65 milliseconds) compared to the CX500 when using RAID 1/0. Additionally, we found that the FAS3020 configured with RAID-DP generated 15 percent better average latency (64 vs. 75 milliseconds) compared to the CX500 configured with RAID 5.

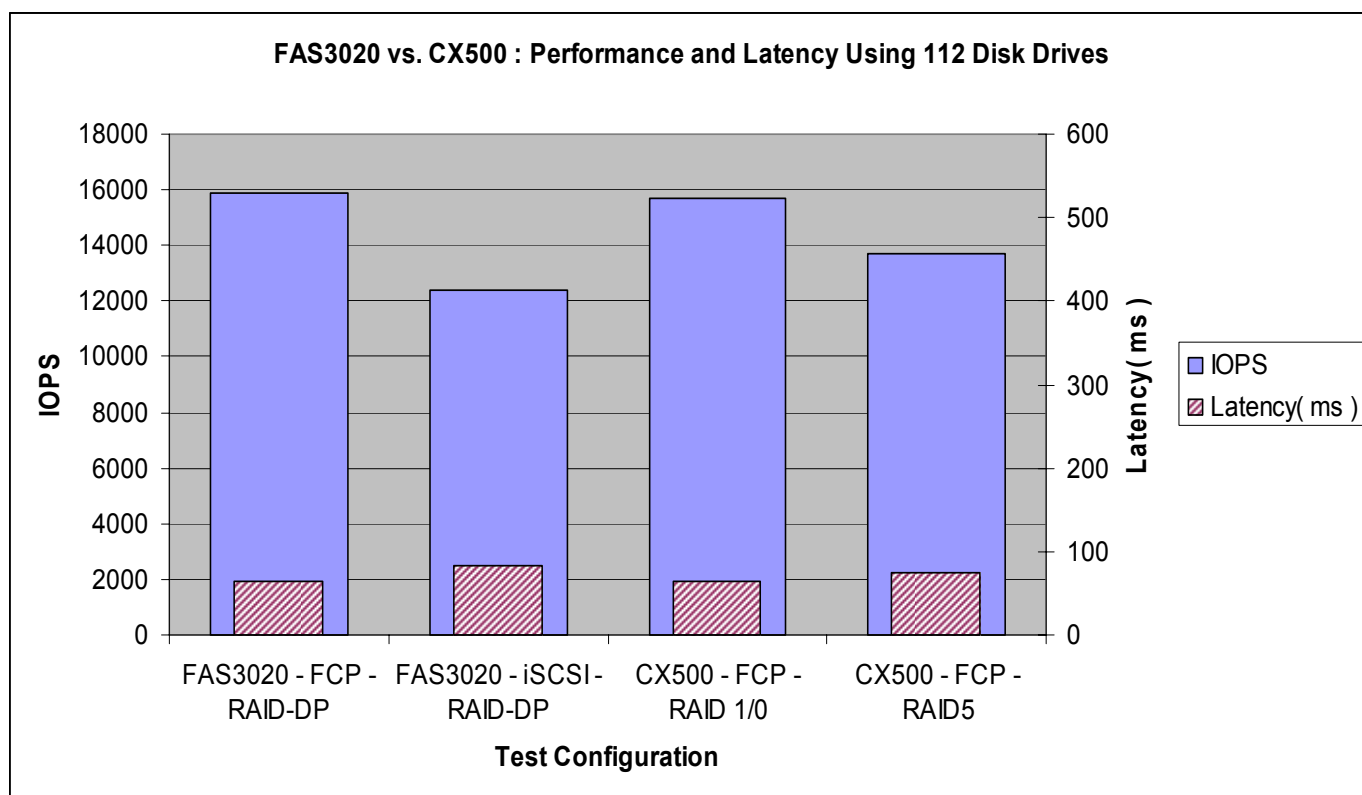


Figure 10. Average IOPS/Latency Results for Performance Testing Using 112 Disk Drives

For purposes of this testing, the simulated company is ACME, Inc. ACME has 3000 employees running Windows PCs, one large Exchange database serving its email needs, two internal departmental SQL Server databases for Payroll and HR, and one large Oracle OLTP database. ACME needs two additional copies of each of its SQL Server and Oracle Databases, one for the Quality Assurance (QA) teams and one for the development teams. ACME initially will access all of its storage over a Fibre Channel Storage Network.

ACME's initial storage needs are specified in figure 11 below.

Database Function	Number, Type and Size of Database and Log Files
Microsoft Exchange	One 100 GB database LUN and one 25 GB log file LUN
SQL Server Payroll	Three 50 GB database LUNs – 1 for production, 1 for development, 1 for quality assurance and three 10 GB LUNs for the database log files
SQL Server HR	Three 50 GB database LUNs – 1 for production, 1 for development, and 1 for quality assurance and three 10 GB LUNs for the database log files
Oracle OLTP	Three 100GB database LUNs – 1 for production, 1 for development, and 1 for quality assurance and three 25 GB LUNs for the database logs

Figure 11. Acme Database and Log File LUN Descriptions

Output from this test consists of a table for each storage solution under test specifying the following items relating to how the database and log volumes described above are provisioned:

- Total Number of RAID groups including the RAID type for each(i.e. RAID 4, RAID 1/0)
- Total Number of LUNs
- Total Number of Physical Disk Drives required
- Approximate amount of raw disk space required for deployment
- Actual usable disk space
- The storage processor/controller on which the specific database or log file was deployed

Provisioning Test Case #3: Measure the Time Required to Create RAID Groups, Volumes and LUNs

For this test we measured and recorded the total amount of time required to create the RAID Groups, Volumes and LUNs necessary to deploy the database configuration for ACME as documented in storage provisioning test case #2 above for both the FAS3020 and the CX500.

Provisioning Test Case #4: Measure Steps and Elapsed Time Required to Delete Database Files

For this test, we measured the time and recorded the steps required to free up space on the both the FAS3020 and CX500 by deleting the data and log LUNs for the production, QA, and development instances of the SQL Server HR databases described in provisioning test case #2 above. Any host, storage controller or storage processor based tools provided by either NetApp or EMC were available for use in this testing. This includes, but is not limited to, the NetApp SnapDrive utility and the Web-based Navisphere application from EMC.

Provisioning Test Case #5: Extending the Size of the Oracle OLTP Database

In this test, we measured the amount of time and recorded the number of steps required to add 100 GB of additional space to the Oracle OLTP database LUN and another 50 GB to the Oracle OLTP database logs LUN.

The CX500 provides two different methods for expanding the size of an existing LUN, namely stripe expansion and concatenate expansion. The stripe expansion method actually re-stripes the existing data on the LUN across all of the drives now participating in the expanded LUN. Concatenate expansion adds the additional space to the end of the existing LUN and does not take the time to re-stripe the existing data over the new disk drives. Because these are two vastly different approaches, we investigated both for this test case using the CX500.

Provisioning Test Case #6: Access SQL Server Development Database via iSCSI

For this test, we measured the amount of time and recorded the number of steps required for ACME to move the Payroll SQL Server development environment to a new Windows Server 2003 host system and connect the new host system to both the FAS3020 and the CX500 using iSCSI and a built-in Ethernet adapter instead of Fibre Channel.

In February 2005, EMC announced support for iSCSI using the CX500. However, at the time of this testing, we were unable to obtain the product enhancements required to enable iSCSI support. As a result, we could not perform this test using the CX500.

Provisioning Test Case #7: Measure Time Required to Restore Snapshot Copies

For this test, we measured the amount of time and recorded the number of steps required to create and restore a total of 30 different snapshot copies of the LUN representing the Oracle OLTP database described in Provisioning Test Case #2 above on both the FAS3020 and CX500 configurations. A snapshot is a point in time copy of a LUN that does not change over time even as the LUN from which it was created changes.

On the FAS3020 we were able to perform the 30 snapshot copies. On the CX500 there is an eight snapshot limitation. For the FAS3020, we used SnapDrive under Windows Server 2003. For the CX500, we used EMC's SnapView utility under Windows Server 2003.

Provisioning Test Case #8: Measure Time Required to Clone LUNs

For this test, we measured the amount of time and recorded the number of steps required to clone the LUN representing the Oracle OLTP database described in Provisioning Test Case #2 above on both the FAS3020 and the CX500. For the FAS3020 we used the LUN clone commands accessed through the FAS3020 command line interface. For the CX500 we used the SnapView Clone command.

Performance Testing

During the performance tests, we used the following configurations:

- Against both the FAS3020 and CX500 using the single 100GB OLTP database LUN defined in Provisioning Test Case #2 and created during the provisioning tests
- The FAS3020 configuration populated with 112 physical disk drives each of which was 36GB in size and 15,000 RPM. We tested the FAS3020 using RAID-DP with both FCP and iSCSI host connections. RAID-DP™ (Double Parity) uses two parity disks per RAID group to decrease the likelihood that a double disk failure will cause data loss.
- The CX500 configuration populated with 112 physical disk drives each of which was 36GB in size and 15,000 RPM. We tested the CX500 using only FCP with RAID 5 and RAID 1/0

To generate the load for the performance testing, we used the industry-standard, open-source load generator IOMeter, available from Source Forge at <http://sourceforge.net/projects/iometer/>.

Performance Test Case #1: Performance Using Oracle OLTP Production LUN

We designed this test case to measure the performance of both the FAS3020 and the CX500 by testing each configuration using the single 100GB LUN representing the Oracle OLTP Production LUN created as part of the provisioning tests described above. Please refer to Figures 5 and 6 in this report for specifics of how the Oracle OLTP production database LUN is provisioned on the FAS3020 and the CX500, respectively.

For this test, we employed an IOMeter test script that generated a workload considered comparable to a database application running OLTP (Online Transaction Processing) workloads. This load consisted of a mixture of 60% random reads and 40% random writes using an 8KB request size. We configured the IOMeter test script to use a ramp up of 60 seconds and a run time of 120 seconds. We ran this test script twice for both the FAS3020 and the CX500 and averaged the results of the two tests to generate the results presented in this report. During this testing, we used the following configurations:

- NetApp FAS3020 using Fibre Channel host attach with Fibre Channel disk drives
- NetApp FAS3020 using the iSCSI protocol with Fibre Channel disk drives
- NetApp FAS3020 using Fibre Channel host attach with SATA disk drives
- EMC CX500 using Fibre Channel host attach and RAID-1/0 Fibre Channel drives

The specific IOMeter test parameters used for this test are shown in figure 12 below:

Test Type	# of IOMeter Workers	# of Outstanding Ios Per Worker	Maximum File Size	Ramp Up Time(seconds)	Run Time(seconds)
60% random reads and 40% random writes using an 8KB request size and 8KB IO Alignment	1	256	100GB	60	120

Figure 12. IOMeter Test Parameters For Provisioning Performance Tests for FAS3020 and CX500

For these tests using the Fibre Channel protocol, we configured a Storage Area Network (SAN) consisting of one Dell PowerEdge 2650 host system configured with two 2.8GHz Intel Xeon processors, 4GB of RAM and running Windows Server 2003 Enterprise Edition. The Dell host system contained two QLogic QLA2340 Fibre Channel HBAs. We connected the Dell host system to either the FAS3020 or the CX500 using a Brocade SilkWorm 3800 Fibre Channel switch. We configured each of the host HBAs used during the testing to use a queue depth of 256.

For the FAS3020 only, we used the identical host configuration described above and replaced the Fibre Channel drives on the FAS3020 with a set of 56 SATA drives each 144GB in size with a speed of 7200 RPM. We then re-created the ACME Company database setup on the SATA drives in order to complete the testing using the 100 GB Oracle OLTP production Database LUN.

In addition to testing using the Fibre Channel protocol, we conducted the same series of tests against the FAS3020 using iSCSI as the host connect. For these tests using iSCSI, we configured a network consisting of one Dell PowerEdge 2650 host system configured with two 2.8GHz Intel Xeon processors, 4GB of RAM and running Windows Server 2003 Enterprise Edition. The Dell host system contained two Intel Pro1000XT Gigabit Ethernet adapters. We connected the Dell host system to the FAS3020 using a 3COM 3C17400 24-port Gigabit Ethernet switch and the Microsoft iSCSI initiator version 1.06.

Figure 13 below provides a graphic depiction of the network used in the tests with the Fibre Channel protocol for both the FAS3020 and the CX500. Figure 14 below provides a graphic depiction of the network used in the tests with the iSCSI protocol for the FAS3020.

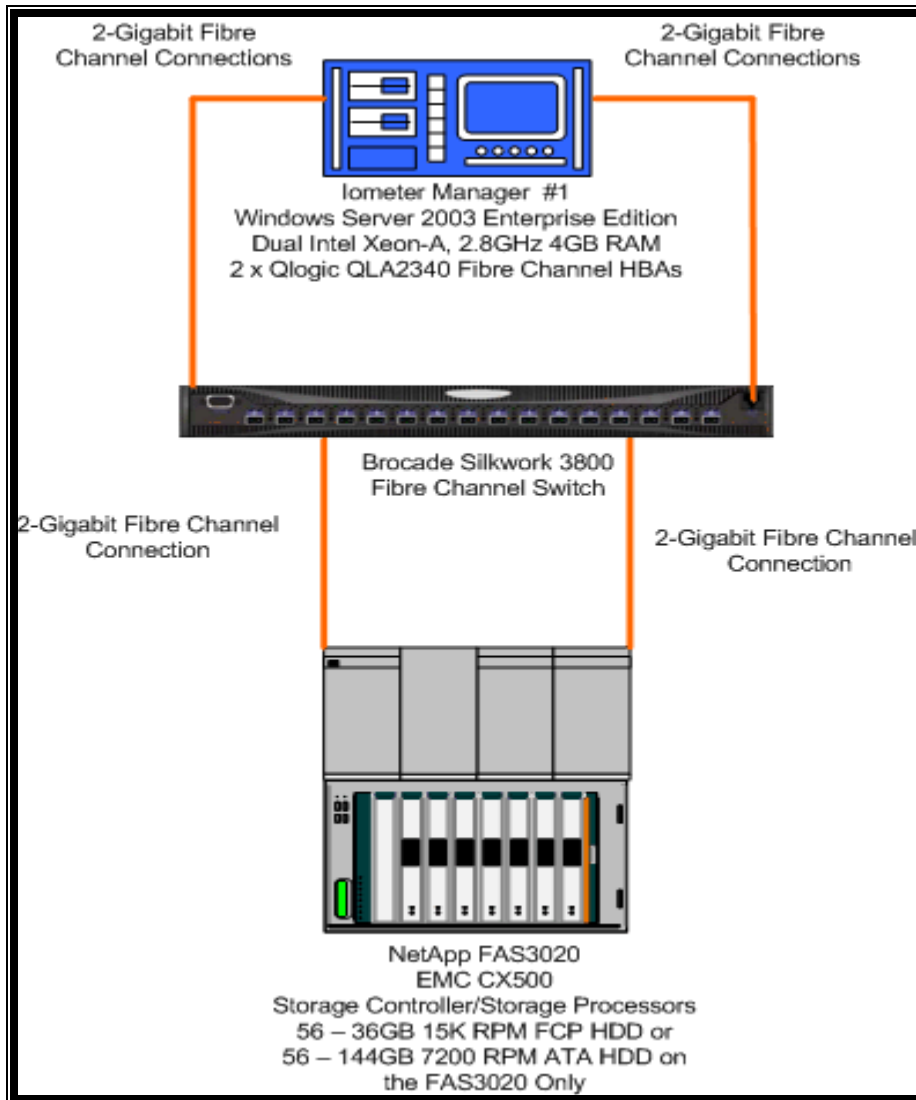


Figure 13: NetApp FAS3020 and EMC CX500- Fibre channel connection diagram

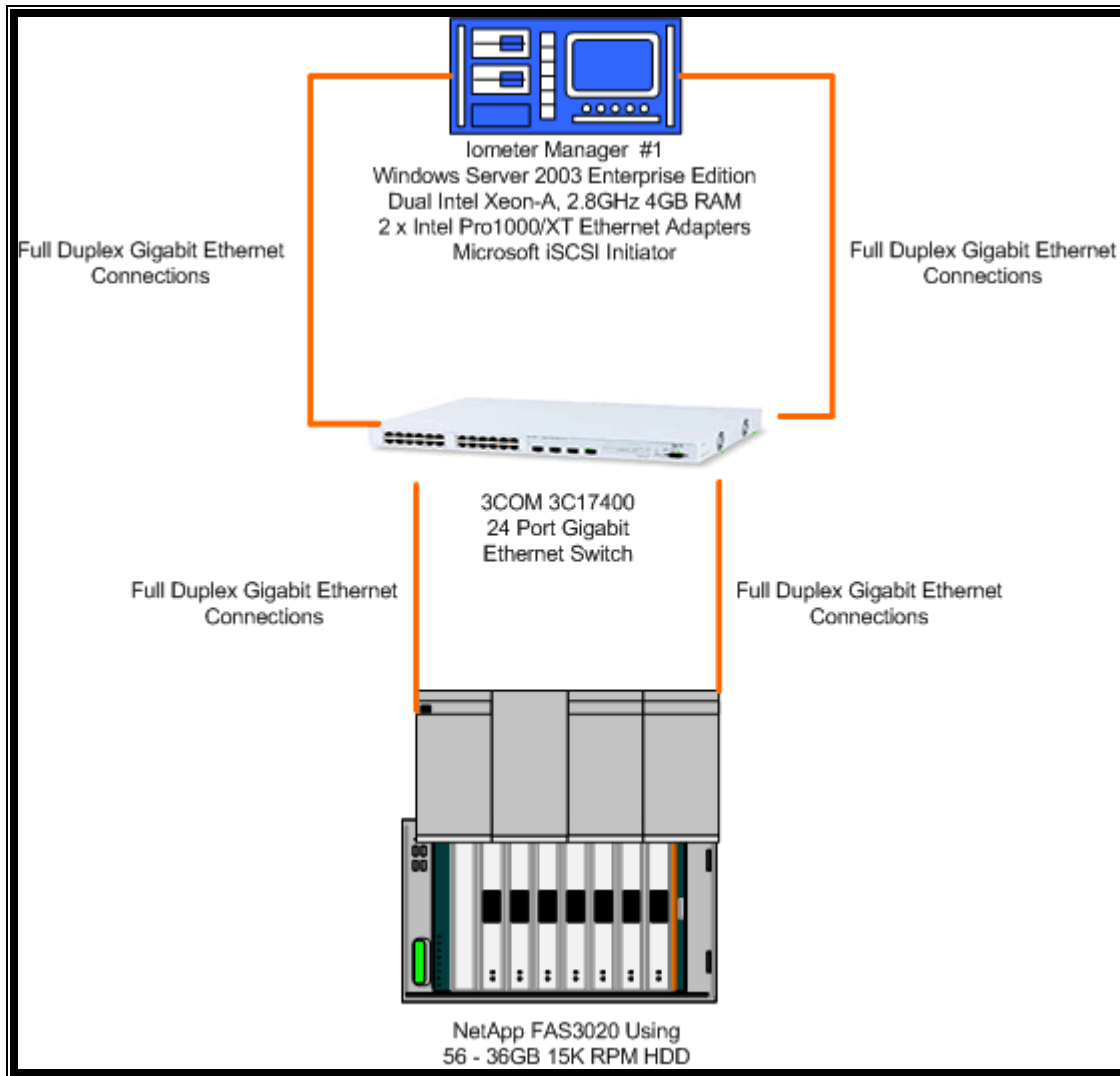


Figure 14: NetApp FAS3020 iSCSI connection diagram

Performance Test Case #2: Measure Performance Impact When Taking Successive Snapshot Copies of the Oracle OLTP Production LUN

For this test, we measured the performance impact on both the FAS3020 and CX500 of creating a series of 30 snapshot copies of the 100GB LUN representing the Oracle OLTP production database described above while the LUN was under a constant load of read and write traffic. To generate the load for the test, we used an IOMeter test script that generated a mixture of 60% random reads and 40% random writes using an 8KB request size. We set the run time in the IOMeter test script to 1 hour and let the test run continuously while taking the series of snapshot copies over the course of the hour.

We configured each of the host HBAs used during the testing to use a queue depth of 256. The specific IOMeter test parameters used for this test are shown in figure 15 below:

Test Type	# of IOMeter Workers	# of Outstanding Ios Per Worker	Maximum File Size	Ramp Up Time(seconds)	Run Time(seconds)
60% random reads and 40% random writes using an 8KB request size and 8KB IO Alignment	1	256	100GB	60	3600

Figure 15. IOMeter Test Parameters For Provisioning Performance Tests for FAS3020 and CX500

For the FAS3020 we performed a series of 30 snapshot copies at 2 minute intervals during the course of the 60 minute test run. Because the CX500 has a limit of a total of 8 snapshot copies for any given LUN, we created a series of 8 snapshot copies at 5 minute intervals and then allowed the test to run for the last 20 minutes with no additional snapshot copies created.

In addition to executing the test described above, we also ran the same hour long IOMeter test script but did not perform the snapshot process during the test period. We used these baseline performance results to help assess the performance impact of the snapshot process on each system.

During the 60 minute duration of the test, we ran the Performance Monitor application from Microsoft on the host system running Windows Server 2003 to measure the read and write activity on the logical volumes being accessed on both the FAS3020 and the CX500 during the testing. We configured Performance Monitor to capture information related to the following logical disk counters at 10 second intervals during the testing:

- Number of disk reads per second
- Number of disk writes per second
- Average read latency in seconds per operation
- Average write latency in seconds per operation

By monitoring these counters while taking the snapshot copies, we were able to determine the impact on the overall performance as a result of the snapshot process. This was not possible using IOMeter alone as it reports only a single average IOPS metric calculated over the entire test run time.

To compute the results presented for this test, we recorded the total IOPS values generated at each of the 10 second intervals using Performance Monitor over the course of the test when not performing the snapshot process and used these values as our baseline.

We then recorded the total IOPS values generated at each of the 10 second intervals using Performance Monitor over the course of the test when performing the snapshot process. This data represented the overall performance of the test configuration when the snapshot process was performed. At each data point, we computed the difference in the number of IOPS between the baseline configuration when no snapshot process was performed and the configuration where we conducted the snapshot process. We calculated the difference as a percentage of the baseline value to see how the performance was impacted over the course of the testing as a result of the snapshot process.

For example, if a specific data point had value of 5000 IOPS when not using a snapshot and 4950 IOPS at the same point in time when performing the snapshot process, we calculated that the performance when conducting the snapshot process at that specific point in time was 99 percent of the baseline performance as follows:

- $100 - ((5000 - 4950) / 5000) * 100$

For this test, a value of 100 indicates that there was no difference between the performance recorded during the testing when taking a snapshot and when not conducting the snapshot process. Data points less than 100 indicate that there was performance degradation between the baseline configuration when no snapshot process was performed and the configuration where we conducted the snapshot process over the course of the test. Lower values indicate a greater performance difference between the baseline configuration and the configuration where the snapshot process was performed.

Performance and Latency Tests Using 112 Drives

We designed these test cases to measure the performance of both the FAS3020 and the CX500 by testing each configuration with a total of 112 physical drives installed. This allowed us to configure the CX500 with the maximum possible number of disk drives.

To generate the load for these tests, we employed an IOMeter test script that generated a workload considered comparable to a database application running OLTP (Online Transaction Processing) workloads. This load consisted of a mixture of 60% random reads and 40% random writes using an 8KB request size size.

During the performance testing, we used the following configurations. Because we did not have 112 SATA drives at the time of the testing, we did not use SATA drives in these tests.

- NetApp FAS3020 using the Fibre Channel protocol
- NetApp FAS3020 using the iSCSI protocol
- EMC CX500 using Fibre Channel and RAID 1/0
- EMC CX500 using Fibre Channel and RAID 5

For these tests using the Fibre Channel protocol, we configured a Storage Area Network (SAN) consisting of two identical Dell PowerEdge 2650 host systems each configured with two 2.8GHz Intel Xeon processors, 4GB of RAM and running Windows Server 2003 Enterprise Edition. Each of the Dell host systems contained two Qlogic QLA2340 Fibre Channel HBAs. We connected both of the Dell host systems to either the FAS3020 or the CX500 using a Brocade SilkWorm 3800 Fibre Channel switch. Each storage system had two Fibre Channel connections to the Brocade switch. All Fibre Channel connections on the host and targets were 2Gbps.

In addition to testing using the Fibre Channel protocol, we conducted the same series of tests against the FAS3020 using iSCSI as a host attach. For these tests using iSCSI, we configured a network consisting of two identical Dell PowerEdge 2650 hosts systems each configured with two 2.8GHz Intel Xeon processors, 4GB of RAM and running Windows Server 2003 Enterprise Edition. Each of the Dell host systems contained two Intel Pro1000XT Gigabit Ethernet adapters. We connected both of the Dell host systems to the FAS3020 using a 3COM 3C17400 24-port Gigabit Ethernet switch and the Microsoft iSCSI initiator version 1.06. The FAS3020 had two Gigabit Ethernet connections to the 3Com switch.

For complete details of the systems used in these tests, including driver related information, please refer to Appendix B of this report.

Figure 16 below provides a graphic depiction of the network used in the tests with the Fibre Channel protocol for both the FAS3020 and the CX500. Figure 17 below provides a graphic depiction of the network used in the tests with the iSCSI protocol for the FAS3020.

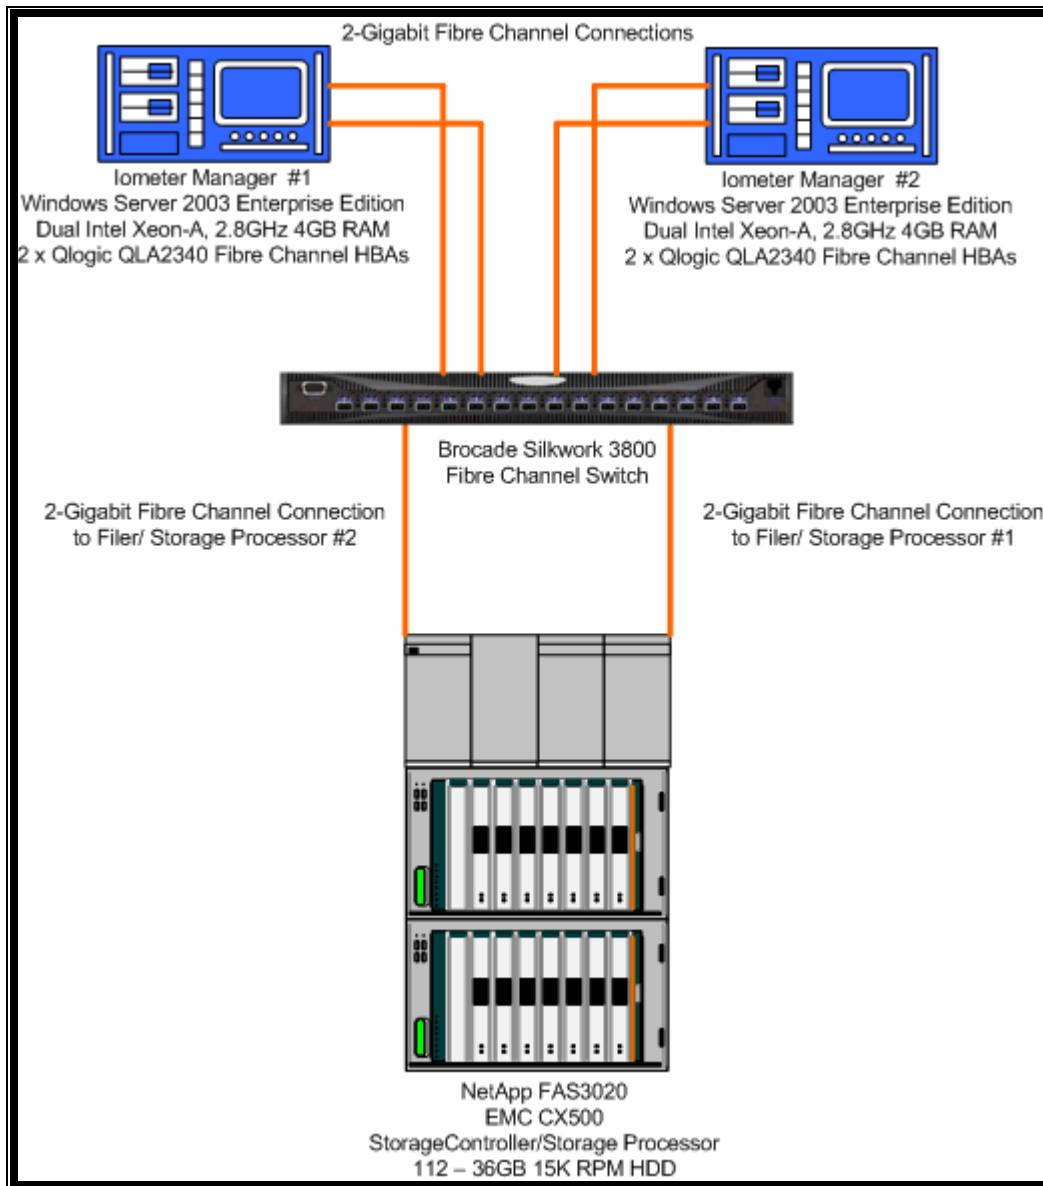


Figure 16: NetApp FAS3020 and EMC CX500- Fibre channel connection diagram




Figure 17: NetApp FAS3020 iSCSI connection diagram

Performance Test Case #3: Measure Performance of CX500 Using 112 Drives and RAID 1/0

The goal of this test is to measure the performance of the CX500 using an OLTP-like workload when populated with the maximum number of physical disk drives in a RAID 1/0 configuration. This load consisted of a mixture of 60% random reads and 40% random writes using an 8KB request size. We configured the IOMeter test script to use a ramp up of 60 seconds and a run time of 120 seconds. We ran this test script twice on the CX500 and averaged the results of the two tests to generate the results presented in this report.

We configured the CX500 from scratch and provisioned the system such that a series of 32 LUNs each with a size of 50GB were created and spread evenly across a total of 112 disk drives using RAID 1/0 with the following layout. Each of the 112 disk drives had a capacity of 36GB and a speed of 15,000 RPM.

- On Storage Processor A, we created a total of 4 RAID 1/0 groups each using 14 physical drives
- On Storage Processor B, we created a total of 4 RAID 1/0 groups each using 14 physical drives

- On Storage Processor A, we created and bound a total of 4 LUNs to each of the 4 RAID groups encompassing the entire usable space in each RAID group. We assigned each of the 16 LUNS a size of 50GB.
- On Storage Processor B, we created and bound a total of 4 LUNs to each of the 4 RAID groups encompassing the entire usable space in each RAID group. We assigned each of the 16 LUNS a size of 50GB.
- On Storage Processor A, we created a storage group and added each of the 16 LUNs to the storage group.
- On Storage Processor B, we created a storage group and added each of the 16 LUNs to the storage group

We then configured the two Dell host systems used in the test so that each host system had access to a total of 16 of the LUNs created on the CX500. For the IOMeter test script, we configured a total of 16 workers for each IOMeter manager system and assigned one of the 16 LUNs mapped on the CX500 to each of the workers as a disk target. The result was a 1 to 1 mapping between the 16 IOMeter workers and the 16 mapped LUNs on each of the Dell host systems for a total of 32 IOMeter workers targeting 32 individual LUNs spread evenly across the two CX500 storage processors utilizing a total of 112 physical disk drives.

For the test, we configured IOMeter to populate each of the 32 LUNs with a 15GB test file. This ensured that enough data was being accessed so that it could not be cached by the CX500. Figure 18 below shows the main IOMeter test parameters used for this test. Because each of the 32 IOMeter workers used a total of 32 outstanding IOs, we configured each of the four host HBAs used during the testing to use a queue depth of 256. This allowed the queue depths available using all of the HBAs to match the total number of possible outstanding IOs using the IOMeter workers.

Test Type	# of IOMeter Workers	# of Outstanding Ios Per Worker	Maximum File Size	Ramp Up Time(seconds)	Run Time(seconds)
60% random reads and 40% random writes using an 8KB request size and 8 KB I/O alignment	32	32	15GB	60	120

Figure 18. CX500 IOMeter Test Parameters using 112 Drives and RAID 1/0

Between runs of the test script we restored the data files used by IOMeter on all 32 LUNs used during the test to their initial state so that the tests used the identical IOMeter data file for both test iterations. To save time during testing, we had planned to use the snapshot capabilities of both the FAS3020 and CX500 to restore the target LUNs used by IOMeter. However, during the testing it became evident that maintaining an active snapshot on the CX500 during the testing caused a significant decrease in the performance of the CX500. As a result, we only used the snapshot technology on the FAS3020 to restore the IOMeter test LUN to its initial state. On the CX500, we deleted the IOMeter work file between runs and forced IOMeter to re-create the file each time between tests that employed write operations.

Performance Test Case #4: Measure Performance of CX500 Using 112 Drives and RAID 5

The goal of this test is to measure the performance of the CX500 using an OLTP-like workload when populated with the maximum number of physical disk drives in a RAID 5 configuration. This load consisted of a mixture of 60% random reads and 40% random writes using an 8KB request size size. We configured the IOMeter test script to use a ramp up of 60 seconds and a run time of 120 seconds. We ran this test script twice on the CX500 and averaged the results of the two tests to generate the results presented in this report.

For this test, we configured the CX500 from scratch and provisioned the system such that a series of 32 LUNs were created and spread evenly across a total of 112 disk drives using RAID 5 with the following layout. Each of the 112 disk drives had a capacity of 36GB and a speed of 15,000 RPM.

- On Storage Processor A, we created a total of 8 RAID 5 groups each using 7 physical drives
- On Storage Processor B, we created a total of 8 RAID 5 groups each using 7 physical drives
- On Storage Processor A, we created and bound two LUNs to each of the 8 RAID groups. We assigned each of the 16 LUNs a size of 50GB.
- On Storage Processor B, we created and bound two LUNs to each of the 8 RAID groups. We assigned each of the 16 LUNs a size of 50GB.
- On Storage Processor A, we created a storage group and added each of the 16 LUNs to the storage group.
- On Storage Processor B, we created a storage group and added each of the 16 LUNs to the storage group

We then configured the two host systems used in the test so that each host system had access to a total of 16 of the LUNs created on the CX500. After configuring the CX500 as described, we executed the workload described above against the new configuration using the testbed described in Figure 16.

For each workload test script on each host, we configured a total of 16 workers for each IOMeter manager system and assigned one of the 16 LUNs mapped on the CX500 to each of the workers as a disk target. The result was a 1 to 1 mapping between the 16 IOMeter workers and the 16 mapped LUNs on each of the Dell host systems for a total of 32 IOMeter workers targeting 32 individual LUNs spread evenly across the two CX500 storage processors utilizing a total of 112 physical disk drives.

For the test, we configured IOMeter to populate each of the 32 LUNs with a 15GB test file. This ensured that enough data was being accessed so that it could not be cached by the CX500. Figure 19 below shows the main IOMeter test parameters used for this test. Because each of the 32 IOMeter workers used a total of 32 outstanding IOs, we configured each of the four host HBAs used during the testing to use a queue depth of 256. This allowed the queue depths available using all of the HBAs to match the total number of possible outstanding IOs using the IOMeter workers.

Test Type	# of IOMeter Workers	# of Outstanding Ios Per Worker	Maximum File Size	Ramp Up Time(seconds)	Run Time(seconds)
60% random reads and 40% random writes using an 8KB request size and 8 KB I/O alignment	32	32	15GB	60	120

Figure 19. CX500 IOMeter Test Parameters using 112 Drives and RAID 5

Between runs of the test script we restored the data files used by IOMeter on all 32 LUNs during the test to their initial state so that the tests used the identical IOMeter data file for both test iterations. To save time during testing, we had planned to use the snapshot capabilities of both the FAS3020 and CX500 to restore the target LUNs used by IOMeter. However, during the testing it became evident that maintaining an active snapshot on the CX500 during the testing caused a significant decrease in the performance of the CX500. As a result, we only used the snapshot technology on the FAS3020 to restore the IOMeter test LUN to its initial state. On the CX500, we deleted the IOMeter work file between runs and forced IOMeter to re-create the file each time between tests that employed write operations.

Performance Test Case #5: Measure Performance of FAS3020 Using 112 Drives and RAID-DP With FCP and iSCSI

The goal of this test is to measure the performance of the FAS3020 using an OLTP-like workload when populated with 112 physical disk drives in a RAID-DP configuration using both Fibre Channel and iSCSI host connections. This load consisted of a mixture of 60% random reads and 40% random writes using an 8KB request size size. We configured the IOMeter test script to use a ramp up of 60 seconds and a run time of 120 seconds. We ran this test script twice on the FAS3020 and averaged the results of the two tests to generate the results presented in this report.

For this test, we configured the FAS3020 from scratch and provisioned the system such that a series of 32 LUNs were created and spread evenly across a total of 112 disk drives using RAID-DP with the following layout. Each of the 112 disk drives had a capacity of 36GB and a speed of 15,000 RPM.

- On storage controller A, we created a single aggregate and volume containing 56 physical drives.
- On storage controller B, we created a single aggregate and volume containing 56 physical drives.
- On storage controller A, we created a total of 4 RAID-DP groups each using 14 physical drives
- On storage controller B, we created a total of 4 RAID-DP groups each using 14 physical drives
- On storage controller A, we created a total of 16 distinct LUNs each with a size of 50GB.
- On storage controller B, we created a total of 16 distinct LUNs each with a size of 50GB.

We then configured the two host systems used in the test so that each host system had access to a total of 16 of the LUNs created on each of the FAS3020 storage controllers. After configuring the FAS3020 as described, we executed the workload described above against the new configuration using the testbed described in Figure 16 when using a FCP host connect and Figure 17 when using an iSCSI host connect.

For each workload test script on each host, we configured a total of 16 workers for each IOMeter manager system and assigned one of the 16 LUNs mapped on each of the FAS3020 storage controllers to each of the 16 IOMeter workers as a disk target. The result was a 1 to 1 mapping between the 16 IOMeter workers and the 16 mapped LUNs on each of the Dell host systems for a total of 32 IOMeter workers targeting 32 individual LUNs spread evenly across the two FAS3020 storage controllers utilizing a total of 112 drives.

For the test, we configured IOMeter to populate each of the 32 LUNs with a 15GB test file. This ensured that enough data was being accessed so that it could not be cached by the FAS3020. Figure 20 below shows the main IOMeter test parameters used for this test. Because each of the 32 IOMeter workers used a total of 32 outstanding IOs, we configured each of the four host HBAs used during the testing to use a queue depth of 256. This allowed the queue depths available using all of the HBAs to match the total number of possible outstanding IOs using the IOMeter workers.

Test Type	# of IOMeter Workers	# of Outstanding ios Per Worker	Maximum File Size	Ramp Up Time(seconds)	Run Time(seconds)
60% random reads and 40% random writes using an 8KB request size and 8 KB I/O alignment	32	32	15GB	60	120

Figure 20. FAS3020 IOMeter Test Parameters using 112 Drives and RAID-DP

Between runs of the test script we restored the data files used by IOMeter during the test to their initial states using the snapshot technology on the FAS3020. This ensured that all test cases started at the same initial state with regard to the IOMeter data file ensuring consistent results between test iterations.

Appendix A. Usability and Provisioning Test Result Details

This section provides the details of the provisioning test cases including the specific steps required to complete each of the test cases along with tester comments and feedback logged during the provisioning test process. Please refer to the Testing Methodology section of this report for complete details on how we conducted each of these tests.

Provisioning Test Case #1: Software Installation and FCP Protocol Configuration

For this test we configured the Fibre Channel protocol on both the Network Appliance FAS3020 and EMC CX500 storage processors and connected each to a Windows Server 2003 host system in preparation for normal storage provisioning. Success criteria for this test consisted of accessing a LUN on the specific storage solution under test using the Fibre Channel Protocol from the Windows Server 2003 host system. During this process we recorded the individual steps and total elapsed time required to complete the process on both the Network Appliance FAS3020 and EMC CX500 configurations

Execution Steps and Elapsed Time for Network Appliance FAS3020

Configuring the FCP protocol on the FAS3020 consumed a total of **60 minutes** of time and required that we complete the following steps:

1. Check Network Appliance documentation on Fibre Channel (F/C) Host Bus Adapter (HBA) compatibility, firmware requirements and driver requirements as well as possible operating system (O/S) patches/hotfixes.
2. Physically install F/C HBA.
 - Check firmware revision of F/C HBA and update if necessary
3. Install F/C drivers for O/S
 - Download latest driver from manufacturer's web site.
4. Physically connect the F/C HBA in the test host to a switch
5. Physically connect the FAS3020 to the switch
6. Load SnapDrive on test host
7. Install any O/S patches or hotfixes

We found that the documentation around configuring the Fibre Channel Protocol was very good for the FAS3020. We also really liked the fact that the documentation was located directly on the FAS3020 allowing us to save time during the FCP protocol configuration.

Currently, there does not appear to be a link between the www.NetApp.com to the <http://NOW.NetApp.com> site as searches done on the NOW site do not reference information on the public site. In addition, supplying customers with the necessary Microsoft Hotfixes will save them time in the setup process. Finally, since the functionality provided by the aggregate storage pool versus a traditional setup is a new concept, supplying a white paper on the advantages and disadvantages of these approaches would be very valuable.

Execution Steps and Elapsed Time for EMC CX500

Configuring the FCP protocol on the CX500 consumed a total of **80 minutes** of time and required that we complete the following steps:

1. Check the EMC documentation on Fibre Channel HBA compatibility, firmware requirements and driver requirements as well as possible operating system patches/hotfixes.
2. Physically install the HBA
 - Check firmware revision of F/C HBA and update if necessary.
3. Install F/C drivers for O/S.
 - Download latest driver from manufacturer's web site.

LUN Description	Storage controller #	RAID Level	LUN Size
Exchange Logs	1	DP	25
Exchange DB	2	DP	100
HR QA DB	2	DP	50
HR QA Logs	1	DP	10
Oracle Prod DB	2	DP	100
Oracle Prod Logs	1	DP	25
Payroll Development DB	1	DP	50
Payroll Development Log	2	DP	10
Oracle Development DB	2	DP	100
Oracle Development Logs	1	DP	25
Oracle QA DB	1	DP	100
Oracle QA Logs	2	DP	25
Payroll Prod DB	2	DP	50
Payroll Prod Log	1	DP	10
HR Development DB	1	DP	50
HR Development Logs	2	DP	10
HR Prod DB	1	DP	50
HR Prod Logs	2	DP	10
Payroll QA DB	1	DP	50
Payroll QA Log	2	DP	10
Totals Size of DB and Log LUNs			860GB

Figure 21. FAS3020 Database and Log File LUN Layout

When conducting this exercise on the CX500, we utilized available EMC documentation and best practices guides listed above to ensure an optimal layout. These documents contained detailed worksheets to help with planning, provisioning and sizing a deployment similar to the one we completed for this testing. Figure 22 below shows how we configured the CX500 to handle the ACME database layout. In addition to determining the space for the database and log volumes, we also factored in 20 percent of additional space over and above the minimum space required for the database and log volumes to use as an area to store snapshot copies. This 20 percent is the same as the default snapshot reserve space used during the layout of the FAS3020.

As figure 22 shows, we found it required a total of 82 physical disk drives to deploy the ACME database environment compared to the 56 physical disk drives required for the same deployment performed on the FAS3020.

Applications	SP	LUN Size in GB	Total Used Space (GB)	RAID Group	RAID Level	BUS	# of Physical Disks	Approx Raw Space (GB)
Exchange Logs	A	25	31	RAID Group 0	5	0	5	165
HR QA Logs	A	10	12.5					
Oracle Development Logs	A	25	31					
Oracle QA Logs	A	25	31					
Payroll QA Log	A	10	12.5					
HR Development Logs	B	10	12.5	RAID Group 1	5	1	5	165
HR Prod Logs	B	10	12.5					
Oracle Prod Logs	B	25	31					
Payroll Development Log	B	10	12.5					
Payroll Prod Log	B	10	12.5					
Exchange DB	B	100	200	RAID Group 2	1/0	1	8	264
Oracle Prod DB	A	100	200	RAID Group 3	1/0	0	8	264
Oracle QA DB	A	100	200	RAID Group 4	1/0	0	8	264
Oracle Development DB	B	100	200	RAID Group 5	1/0	1	8	264
HR QA DB	A	50	100	RAID Group 6	1/0	0	10	330
Payroll Development DB	A	50	100					
Payroll Prod DB	A	50	100					
HR Development DB	B	50	100	RAID Group 7	1/0	1	10	330
HR Prod DB	B	50	100					
Payroll QA DB	B	50	100					
Snapshot Reserve	A	N/A	N/A	RAID Group 8	1/0	0	10	330
Snapshot Reserve	B	N/A	N/A	RAID Group 9	1/0	1	10	330
Totals		860	1599				82	2706

Figure 22: CX500 Database Provisioning

Provisioning Test Case #3: Measure the Time Required to Create RAID Groups, Volumes and LUNs

For this test we measured and recorded the total amount of time required to create the RAID Groups, Volumes and LUNs necessary to deploy the database configuration for ACME as documented in storage provisioning test case #2 above for both the FAS3020 and the CX500.

Execution Steps and Elapsed Time for Network Appliance FAS3020

Creating all of the necessary aggregates, volumes and LUNs for deploying the ACME database scenario consumed **35 minutes** of time and required that we complete the following steps:

1. Review the NetApp documentation to determine the appropriate LUN setup
2. Create a volume aggregate on each of the FAS3020 storage controllers containing a total of 28 physical disk drives
3. Create a flexible volume under each of the two volume aggregates contain the LUNs for database and log files.
4. Create the individual LUNs for each of the database and log files

For this test case, we utilized the aggregate feature of ONTAP 7.0 to create an area to house multiple

volumes and their associated LUNs. After creating the containing aggregate, we found it required only 20 seconds to create, format and mount a 100GB LUN. Additionally, using a flexible volume improves the performance over traditional RAID Groups by allowing any LUN created in the flexible volume to utilize all the physical disk spindles contained in the aggregate. We found the SnapDrive plug in for Windows to be extremely useful when interacting with the FAS3020 by providing a GUI based interface for creating and managing LUNs in our Windows environment. The interface aids in the management and adds a nice ease-of-use touch to the device.

The FAS3020 required the creation of a CIFS share on the storage controller in order to use the SnapDrive interface. Because this is different from what most customers would be accustomed to with Fibre Channel disk arrays, it could be considered confusing to the user. Out of band management on a traditional Fibre Channel disk array generally takes place transparently via TCP ports from a client to the device.

In addition, we initially found the time required to create LUNs via SnapDrive was unusually long and even created time out situations. After talking to Network Appliance support, we were provided with the following set of registry entries to set on the Windows host system to remedy this issue.

[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\SWSvc\Parameters\PreferredFilerIPAddresses] "netapppri"="10.160.37.251"

[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\SWSvc\Parameters\PreferredFilerIPAddresses] "netappsec"="10.160.37.241"

Changing the registry entries resolved the issue. We felt that documenting this upfront would be very valuable to a customer.

Execution Steps and Elapsed Time for EMC CX500

Creating all of the necessary RAID Groups and LUNs for deploying the ACME database scenario on the CX500 consumed **1 hour and 25 minutes** of time and required that we complete the following steps:

1. Review the EMC provided CX500 Configuration Planning Guide to determine the appropriate LUN setup.
2. Complete the associated load worksheets and determine appropriate configuration
3. Create RAID Groups to contain the database and log file LUNs as detailed in the planning guide
4. Create the individual LUNs to hold the database and log files
5. Create a set of Storage Groups and add the individual LUNs to each as appropriate

For this exercise, we used the EMC Navisphere Java Applet to administer and monitor the physical layout of the array while implementing the ACME database layout. We found the Navisphere interface to be easy to learn and utilize.

The EMC Configuration Planning Guide provides detailed guidance on how to configure the CX500. When creating the LUNs for this test, we found it can take a considerable amount of time to complete the LUN binding process on the CX500. For example, it took nearly 20 minutes to bind a 100GB LUN. However, Navisphere provides the capability to bind multiple LUNs in the same RAID Group and of the same size simultaneously. This saved some time compared to executing the binding process on each individual LUN. Additionally, we did not find a simple way to capture the current layout of the storage array for printing or otherwise. A feature that allowed the administrator to print the current layout would be very useful.

Provisioning Test Case #4: Measure Steps and Elapsed Time Required to Delete Database Files

For this test, we measured the time and recorded the steps required to free up space on the both the FAS3020 and CX500 by deleting the data and log LUNs for the production, QA, and development instances of the SQL Server HR databases. Any host, storage controller or storage processor based tools provided by

either NetApp or EMC were available for use in this testing. This includes, but is not limited to, the NetApp SnapDrive utility and the Web-based Navisphere application from EMC.

Execution Steps and Elapsed Time for Network Appliance FAS3020

Deleting the specified database files on the FAS3020 consumed **3 minutes** of time and required us to complete the following steps:

1. Highlight the virtual drive entries under the SnapDrive heading in the “Computer Management” panel.
2. Right click and select “Delete”.
3. Answer “OK” to question on whether or not to delete the LUN
4. Repeat for all requested LUNs

Overall, we found the process of deleting LUNs on the FAS3020 to be quick and easy to accomplish using the SnapDrive utility.

Execution Steps and Elapsed Time for EMC CX500

Deleting the specified database files on the CX500 consumed **5 minutes** of time and required us to complete the following steps:

1. Click on “Storage Group” icon in Navisphere.
2. Right click on the storage group and click on “Select LUNs”
3. Highlight on the LUNs to be removed, click on the left pointing arrow.
4. Click OK.
5. Once back on the main menu, click on the LUNs to be removed (ctrl click to select multiple LUNs).
6. Right click over one of the LUNs and choose “Unbind LUN”
7. Click ‘Yes’ followed by OK and the LUN is gone.

Overall, we found the process of deleting LUNs on the CX500 to be quick and easy to accomplish using the Navisphere application.

Provisioning Test Case #5: Extending the Size of the Oracle OLTP Database

In this test, we measured the amount of time and recorded the number of steps required to add 100 GB of additional space to the Oracle OLTP database LUN and another 50 GB to the Oracle OLTP database logs LUN.

The CX500 provides two different methods for expanding the size of an existing LUN, namely stripe expansion and concatenate expansion. The stripe expansion method actually re-stripes the existing data on the LUN across all of the drives now participating in the expanded LUN. Concatenate expansion adds the additional space to the end of the existing LUN and does not take the time to re-stripe the existing data over the new disk drives. Because these are two vastly different approaches, we investigated both for this test.

Execution Steps and Elapsed Time for Network Appliance FAS3020

Increasing the size of the Oracle OLTP database and log file required **30 seconds** and required we complete the following steps on the FAS3020:

1. Select Virtual disk in the “Computer Management” interface.

2. Right click on the Virtual disk and select “Expand Disk”
3. Enter size to expand by and click “OK”

We found the process of expanding the size of the LUNs to be very fast and easy to use. It also provide the additional benefit of performing the complete operation including expanding the file system.

We found the language used for this expansion process somewhat confusing. A suggestion would be to offer something such as “Size of New LUN” rather than asking for the “Space to Expand By”. This would eliminate the need for any potential mathematical operations during this process.

Execution Steps and Elapsed Time for EMC CX500

Increasing the size of the Oracle OLTP database and log file required **20 minutes** using the concatenation expansion process and **70 minutes** when using the stripe expansion and required we complete the following steps on the CX500:

1. Create a new LUN of the same RAID type and desired size to add. Right click on the LUN to be expanded and click “Expand”.
2. The expansion wizard appears and gives the user the option of Stripe Expansion or Concatenate Expansion
 - Stripe Expansion will re-stripe the data
 - Concatenate Expansion will add the additional space to the LUN

We found the process of expanding the size of the LUNs to be very fast and easy to use. However, when compared to the FAS3020, the LUN expansion process required significantly more time on the CX500.

Provisioning Test Case #6: Access SQL Server Development Database via iSCSI

For this test, we measured the amount of time and recorded the number of steps required for ACME to move the Payroll SQL Server development environment to a new Windows Server 2003 host system and connect the new host system to both the FAS3020 and the CX500 using iSCSI and a built-in Ethernet adapter instead of Fibre Channel.

At the time of this testing, EMC had not yet released iSCSI support for the CX500. Such support was announced in February, 2005. As a result, we were unable to complete this test for the CX500.

Execution Steps and Elapsed Time for Network Appliance FAS3020

To complete this test on the FAS3020 consumed **40 minutes** of time and required that we complete the following steps:

1. Install and Setup the Microsoft software iSCSI Initiator on new test host.
2. Configure the iSCSI initiator to ‘see’ the FAS3020 interface on new test host.
3. Load SnapDrive on new test host.
4. Install any O/S patches or hotfixes.
5. Disconnect the Virtual disks from the original test host.
 - Open “Computer Management”, open SnapDrive and select Virtual disk to disconnect.
 - Right click on target virtual disk and select “Disconnect Disk”
6. On the new host open the “Computer Management” window.

7. Expand the “SnapDrive” item by clicking on it.
8. Make sure there is an iSCSI session connected to the FAS3020 by expanding the “iSCSI Management” item and confirming the active sessions.
9. Select the “Disks” item under “SnapDrive”.
10. Right click on “Disks” and select “Connect Disk”
11. Answer dialogue boxes to connect to the Virtual Disk.

We found the process of connecting to the Payroll SQL Server Database LUN to be almost as easy as “flipping a switch”. Additionally, we found the documentation to be very clear and easy to follow. This makes the versatility of the Network Appliance Storage Server of great interest to organizations evaluating the use of multiple protocols. With the rapid emergence of iSCSI, we believe this will be a significant selling point advantage for Network Appliance.

Provisioning Test Case #7: Measure Time Required to Restore Snapshot Copies

For this test, we measured the amount of time and recorded the number of steps required to create and restore a total of 30 different snapshot copies of the LUN representing the Oracle OLTP database on both the FAS3020 and CX500 devices. A snapshot is a point in time copies of a LUN that does not change over time even as the LUN from which it was created changes. For the FAS3020, we used SnapDrive™ under Windows Server 2003. For the CX500, we used the EMC SnapView utility under Windows Server 2003.

During the course of the testing, we discovered that the CX500 has a hard limitation of eight (8) snapshot copies per LUN. As a result, we only created and restored a total of eight snapshot copies on the CX500.

Execution Steps and Elapsed Time for Network Appliance FAS3020

To complete this test on the FAS3020 consumed **5 minutes** of time to create the 30 snapshot copies, **2 minutes** of time to restore each of the 30 snapshot copies and required that we complete the following steps:

1. Open “Computer Management” window.
2. Expand “SnapDrive” item.
3. Expand “Disks” item.
4. Expand target Virtual Disk.
5. Select “Snapshots”.
6. Right click “Snapshots” and select “Create Snapshot”.
7. Answer dialogue box with name of new snapshot.
8. Repeat the above 29 more times.
9. To recover a snapshot, select a snapshot created in the above process.
10. Right click on the target snapshot and select “Restore disk from snapshot”.
11. Click “Yes” on dialogue box.
12. Repeat above step for the other 29 snapshot copies.

We found the snapshot function was very easy to perform. The visual interface in SnapDrive makes the snapshot process far more understandable and manageable. In addition, we found the documentation on disk space requirements related to the snapshot process was clear and easy to understand.

Execution Steps and Elapsed Time for EMC CX500

To complete this test on the CX500 consumed **5 minutes** of time to create the 8 snapshot copies, **2 minutes** of time to restore each of the 8 snapshot copies and required that we complete the following steps:

1. Create two additional LUNs to be used as “Reserved Private LUNs”

2. Right click on one of the service processors “Reserved LUN Pool” (A or B)
3. Select “Configure”
4. Add the two additional LUNs to the SPs, one to each side.
5. Right click on the desired LUN, select SnapView, select “Persistent”, then “Create Snapshot”, enter a name for the snapshot, and click OK.
6. Repeat the above 8 more times.
7. Right click on the LUN to be snapped, select “Create SnapView Session”, label the session and click OK.
8. To recover a snapshot, right click a snapshot created in the above process and select “Activate Snapshot”.
9. Right click on the snapshot session and select “Start Rollback”
10. Label the rollback, select the priority level, and click OK.
11. Repeat above steps for the other 8 snapshot copies.

We found that the visual interface in Navisphere provides an easy to follow method of creating and restoring snapshot copies. Additionally, we found the documentation is clear on how to perform the operations.

We found there are many steps involved in performing this operation on the CX500. Also, we were somewhat unclear about the connection between the snapshot and the snapshot session. Finally, we could find no information on how much disk space is required to perform the snapshot process.

Provisioning Test Case #8: Measure Time Required to Clone LUNs

For this test, we measured the amount of time and recorded the number of steps required to clone the LUN representing the Oracle OLTP database on both the FAS3020 and the CX500. For the FAS3020 we used the LUN clone commands. For the CX500 we used the SnapView Clone command.

Execution Steps and Elapsed Time for Network Appliance FAS3020

To complete this test on the FAS3020 consumed **30 seconds** of time and required that we complete the following steps:

1. Open command line interface with FAS3020 Storage Servers.
2. Issue “vol clone create” command.
3. Issue “vol status” command to confirm.

For this test, we used the FlexClone feature on the specific OLTP database LUNs. Flexible cloning does not require bringing down the production database to make the copy. This is critical in a production environment. Another benefit of flexible cloning is that the cloned database uses the same blocks on disk as the original database so only the changed blocks require additional space. The commands to clone the LUNs happen almost instantaneously. We liked the fact that the documentation to support this process is located on board the Storage Server.

The documentation states that you must choose one method of managing your FAS3020; command line, ONTAP or SnapDrive. However, there are operations that **MUST** occur either via command line or ONTAP. For example, creation of the aggregate, flexible volume and the management share are all done on the storage controller via command line OR the ONTAP web interface. Cloning is also done only via command line interface. This can create confusion on the management of the device as SnapDrive’s capabilities are also available in the other management modes with easy access menus. It would be valuable to either expand the capabilities of the SnapDrive interface or simply hide the ONTAP menu options. The Command Line Interface may be best left only for those really looking to get their hands dirty.

Execution Steps and Elapsed Time for EMC CX500

To complete this test on the CX500 consumed **65 minutes (40 minutes to create the clone and 25 minutes to sync the clone)** of time and required that we complete the following steps:

1. Right click on the LUN to be cloned and select the SnapView option then "Create Clone Group".
2. When the window comes up, name the clone group, select the LUN to be cloned, and click OK. Command completes almost instantaneously.
3. Create two new LUNs of the same RAID type and size
4. Create an additional LUN of the same RAID type and size
5. Once bound, right click on the CX serial number, select SnapView, then select "Clone Feature Properties".
6. Select the two new LUNs and add them to the Clone Private LUNs.
7. Right click on the Clone Group and select "Add Clone".
8. Select the third newly created LUN, change the sync priority to high, and click OK

We found that the CX500 documentation provides a clear set of steps related to the cloning process. However, compared to the FAS3020, we found the cloning process on the CX500 required significantly more time and steps to complete.

Appendix B: System disclosures

Network Appliance FAS3020	
Storage Processor Unit –1 w/2 SPU	Network Appliance FAS3020 Storage Servers
Cache size	4GB
Disk Arrays – 2	Disk Array Enclosures
Disk Drives	14 per DAE – 36GB 15K RPM Fibre Channel
Base software version	Network Appliance Data ONTAP Release 7.0.1X9

Figure 23: Network Appliance FAS3020 Disclosure Information

EMC CX500 CLARiiON Storage Server	
Storage Processor Unit –1 w/2 SPU	EMC CX500 Storage Processor Enclosure Array
Cache size	4GB
Disk Arrays – 2	EMC DAE2-OS Disk Array Enclosure and Supplement Power Supplies
Disk Drives	14 per DAE – 36GB 15K RPM Fibre Channel
EMC PowerPath	W2k.2.0.0_b034
EMC Navisphere Manager	6.7.3.2.0

Figure 24: EMC CX500 Storage Server Disclosure Information

Dell Computer Corporation PowerEdge 2650 – IOMeter Manager	
Processor / Speed / # of CPUs	Dual Intel 2.8Ghz Xeon, 4800 MHz (2 x 2400)
System RAM / Type / # of Slots	4 GB (PC600 RDRAM)
Network Adapter	2- Intel 82544 PRO/1000 XT Gigabit Ethernet Controller
OS for IOMeter tests	Microsoft Windows 2003 Enterprise Edition
Installed HBAs	2 x Qlogic 2340 version 9.0.1.12 (SCSI miniport) Driver Date: 10/10/2004
iSCSI Initiator Used	Microsoft iSCSI Software Initiator v1.06

Figure 25: Dell PowerEdge 2650 Servers Used for IOMeter Manager Systems

Networking Equipment	
Fibre Channel Switch	Brocade SilkWorm 3800 Firmware v3.1.3 ^a
Network Switch (for iSCSI Tests)	3COM Gig-E Managed Switch

Figure 26: Networking Equipment Used

VeriTest (www.veritest.com), the testing division of Lionbridge Technologies, Inc., provides outsourced testing solutions that maximize revenue and reduce costs for our clients. For companies who use high-tech products as well as those who produce them, smoothly functioning technology is essential to business success. VeriTest helps our clients identify and correct technology problems in their products and in their line of business applications by providing the widest range of testing services available.

VeriTest created the suite of industry-standard benchmark software that includes WebBench, NetBench, Winstone, and WinBench. We've distributed over 20 million copies of these tools, which are in use at every one of the 2001 Fortune 100 companies. Our Internet BenchMark service provides the definitive ratings for Internet Service Providers in the US, Canada, and the UK.

Under our former names of ZD Labs and eTesting Labs, and as part of VeriTest since July of 2002, we have delivered rigorous, objective, independent testing and analysis for over a decade. With the most knowledgeable staff in the business, testing facilities around the world, and almost 1,600 dedicated network PCs, VeriTest offers our clients the expertise and equipment necessary to meet all their testing needs.

For more information email us at info@veritest.com or call us at 919-380-2800.

Disclaimer of Warranties; Limitation of Liability:

VERITEST HAS MADE REASONABLE EFFORTS TO ENSURE THE ACCURACY AND VALIDITY OF ITS TESTING, HOWEVER, VERITEST SPECIFICALLY DISCLAIMS ANY WARRANTY, EXPRESSED OR IMPLIED, RELATING TO THE TEST RESULTS AND ANALYSIS, THEIR ACCURACY, COMPLETENESS OR QUALITY, INCLUDING ANY IMPLIED WARRANTY OF FITNESS FOR ANY PARTICULAR PURPOSE. ALL PERSONS OR ENTITIES RELYING ON THE RESULTS OF ANY TESTING DO SO AT THEIR OWN RISK, AND AGREE THAT VERITEST, ITS EMPLOYEES AND ITS SUBCONTRACTORS SHALL HAVE NO LIABILITY WHATSOEVER FROM ANY CLAIM OF LOSS OR DAMAGE ON ACCOUNT OF ANY ALLEGED ERROR OR DEFECT IN ANY TESTING PROCEDURE OR RESULT.

IN NO EVENT SHALL VERITEST BE LIABLE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH ITS TESTING, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL VERITEST'S LIABILITY, INCLUDING FOR DIRECT DAMAGES, EXCEED THE AMOUNTS PAID IN CONNECTION WITH VERITEST'S TESTING. CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES ARE AS SET FORTH HEREIN.