



***i*Stor Networks, Inc.™**

GigaStor Spiral Performance

May 10, 2006

© 2006 iStor Networks, Inc.® All Rights Reserved

iStor Networks, Inc. makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. iStor Networks, Inc. shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of iStor Networks, Inc.

The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for iStor products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. iStor shall not be liable for technical or editorial errors or omissions contained herein.

Copyright © 2006 iStor Networks, Inc.®

Trademarks

Adobe® and Acrobat® are trademarks of Adobe Systems, Incorporated.

Java™ is a U.S. trademark of Sun Microsystems, Incorporated.

Microsoft Windows is a U.S. registered trademark of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

All other brand or product names are or may be trademarks or service marks, and are used to identify products or services, of their respective owners.

iStor Networks, Inc.
7585 Irvine Center Drive
Suite 250
Irvine, CA 92618

TABLE OF CONTENTS

TRADEMARKS	2
1. INTRODUCTION	5
1.1. PURPOSE.....	5
1.2. SCOPE.....	5
2. BASELINE TEST.....	6
2.1. CONFIGURATION.....	6
2.2. RESULTS.....	6
3. RAID 5 SCALING.....	8
3.1. CONFIGURATION.....	8
3.2. RESULTS.....	9
4. RAID 0 SCALING.....	11
4.1. CONFIGURATION # 1	11
4.2. RESULTS.....	12
4.3. CONFIGURATION # 2	14
4.4. RESULTS.....	15
5. MULTIPLE VOLUMES PER DRIVE.....	17
5.1. CONFIGURATION.....	17
5.2. RESULTS.....	19
6. WIDE CONNECTIONS	21
6.1. WIDE VOLUME – 4 CONNECTIONS/SESSION	22
6.1.1. Configuration of a Wide RAID 0	22
6.1.2. Results.....	23
6.1.3. Configuration of Added RAID 0 Volumes	24
6.1.4. Results.....	25
6.1.5. Configuration of RAID 5 with RAID 0	26
6.1.6. Results.....	27
7. APPENDIX A.....	28
7.1. TEST EQUIPMENT DESCRIPTION.....	28
7.2. TEST EQUIPMENT DESCRIPTION FOR WIDE CONNECTIONS	30
8. APPENDIX B.....	32
8.1. MICROSOFT FILE SYSTEM WRITES	32

LIST OF FIGURES

Figure 1 - JBOD Baseline Configuration	6
Figure 2 - JBOD Spiral Write	7
Figure 3 - JBOD Spiral Read	7
Figure 4 - RAID 5 Configuration (4 volumes of 4 drives).....	8
Figure 5 - RAID 5 Spiral Write.....	9
Figure 6 - RAID 5 Spiral Read.....	10
Figure 7 – RAID 0 Configuration (4 volumes of 4 drives)	11
Figure 8 – RAID 0 Spiral Write (4 Hosts).....	12
Figure 9 – RAID 0 Spiral Read (4 Hosts).....	13
Figure 10 – RAID 0 Configuration (8 volumes of 2 drives)	14
Figure 11 - RAID 0 Spiral Write (8 Hosts).....	15
Figure 12 - RAID 0 Spiral Read (8 Hosts).....	16
Figure 13 – RAID 5 Configuration (multiple volumes per drive)	18
Figure 14 - RAID 5 Spiral Write (multiple volumes per drive).....	19
Figure 15 - RAID 5 Spiral Read (multiple volumes per drive).....	20
Figure 16 – Configuration of a RAID 0 with wide connections.....	22
Figure 17 – RAID 0 Spiral Write/Read (wide connections).....	23
Figure 18 - Configuration Spiral RAID 0 (wide and narrow connections)	24
Figure 19 – Spiral Read RAID 0 (wide & narrow connections)	25
Figure 20 - Configuration Spiral RAID 5 Wide & Spiral RAID 0 Narrow	26
Figure 21 - Spiral Read RAID 5 & 0 (wide & narrow connections)	27
Figure 22 – Equipment List	28
Figure 23 – iSCSI Parameters	29
Figure 24 - Equipment List - Wide Connections	30
Figure 25 - iSCSI Parameters - Wide Connections	31
Figure 26 - Disk Manager POP-UP.....	32

1.

Introduction

The structure of the paper is to illustrate the configurations and results of a number of differing system configurations and how the GigaStor responded during spiral operations. Spiral operations are sometimes called sequential IO; where the IO patterns move from a starting point to an end point in succession. The following performance is not a result of cache hit numbers in as much as the data is written to/read from the actual disk drives, however the performance is derived from the GigaStor's caching algorithms fulfilling the IOs from within its cache because of the read ahead and write coalescing strategies.

Hardware identifications for the configurations can be found in Appendix A.

1.1. Purpose

The purpose of this document is to illustrate GigaStor's spiral performance.

1.2. Scope

The scope is to show the GigaStor's spiral performance in a variety of configurations. It is not the intent of the document to analyze GigaStor's performance, but rather provide to the reader an idea of the type of performance that the GigaStor should deliver. GigaStor's IOPS performance will be documented in a subsequent paper.

Although this paper deals only with spiral type operations it is noted here that this is not a comprehensive view of the types of configurations that could be configured; rather most of the configurations are base-level configurations that are attempting to illustrate the spiral bandwidth capabilities of the GigaStor. Other configurations may have differing results. The information in this document was designed and tested to maximize the GigaStor's performance and identify maximum product capability.

All tests were performed with GigaStor version 1.1.0.12. The results are the same and apply for firmware releases 1.1.0.15, 1.1.1.1, and 1.2.0.3.

2. Baseline Test

The baseline for all of the RAID level performance is based upon the performance of a single drive. This implies that all drives within the array are of the same type. Once the performance of a single drive is known, all other RAID level's theoretical performance is able to be calculated.

2.1. Configuration

The configuration for the baseline JBOD test is simple as illustrated in Figure 1 below. There is a single host with a single NIC using Microsoft Initiator connected to a single port on the GigaStor through an Ethernet switch. The GigaStor is configured with single JBOD volume sitting behind a single iSCSI Target node. The JBOD is a volume that is wholly contained within a single disk drive. The results of the JBOD testing are illustrated in Figure 2 and Figure 3 for writes and reads respectively.

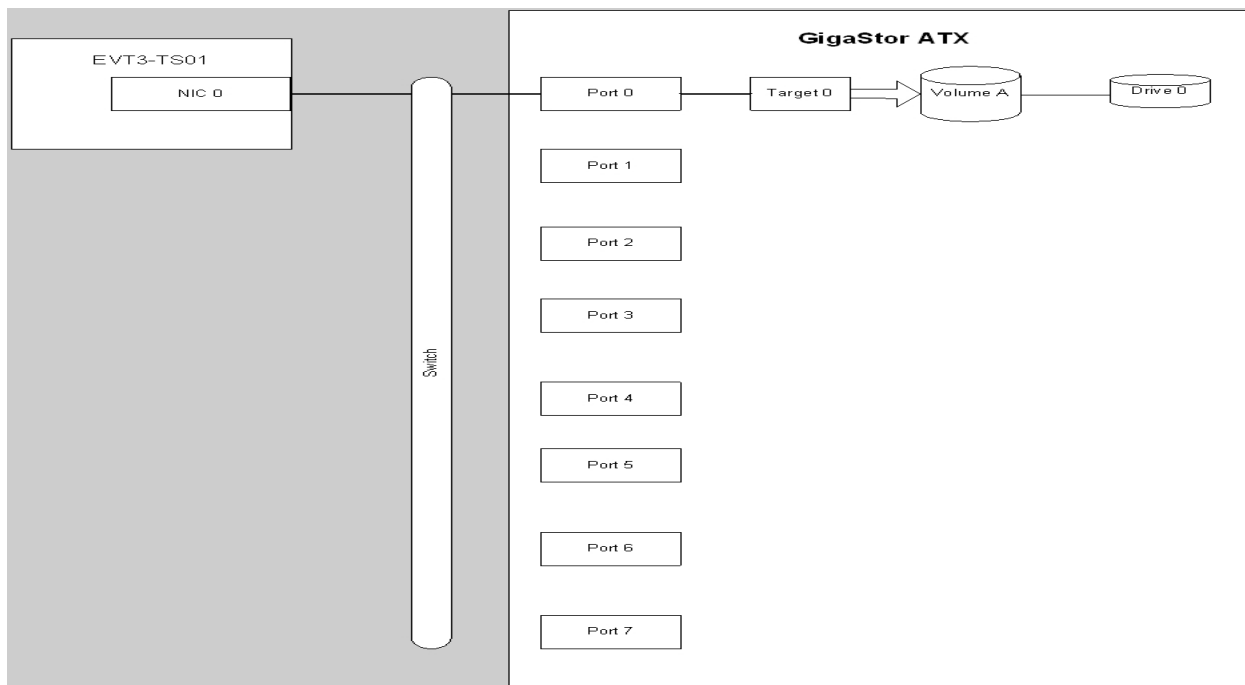


Figure 1 - JBOD Baseline Configuration

2.2. Results

The results illustrate that the drive can sustain 40 MB/sec for writes and 60 MB/sec for reads. As with all the benchmarks, these figures do not represent the data rate of the bits coming off the platter or the transferring of data across the SATA interface, but indicate the drive's sustained bandwidth which includes all drive latencies which are encountered during long spiral operations.

The quick rise of the transfer rates of the disk drives can be attributed to the internal algorithms of the GigaStor. The GigaStor is reading and writing much larger values from the disk drive than what the host is requesting. During the read operations, the GigaStor is fulfilling the request from its cache as the data has been pre-fetched; for writes, the GigaStor terminates the operation with the host as soon as the data arrives and at some other point in time, when it can maximize the utilization of the disk drive, the GigaStor will actually write the data to the drive.

These 2 values will be used during explanations of more complex configurations.

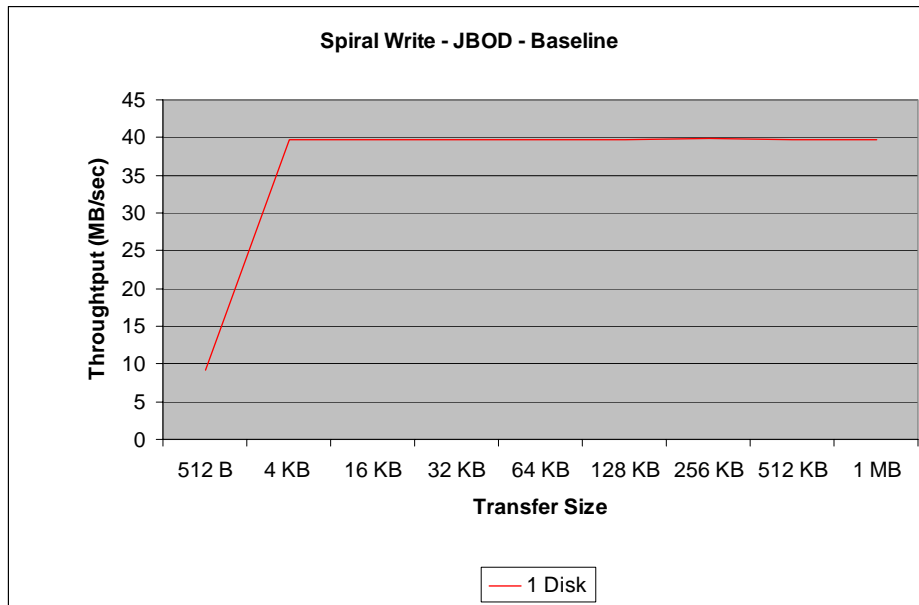


Figure 2 - JBOD Spiral Write

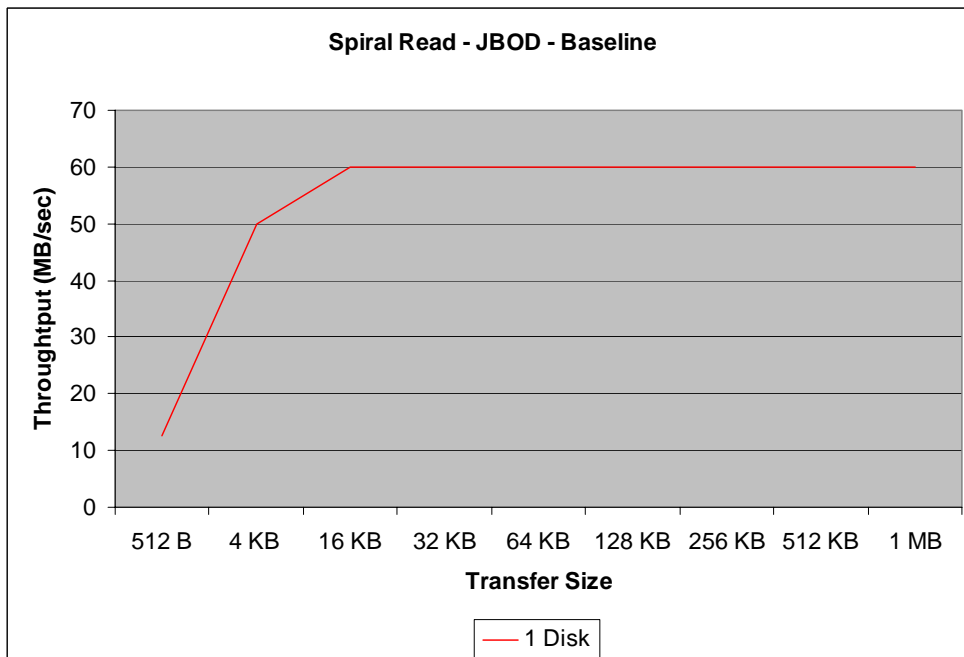


Figure 3 - JBOD Spiral Read

3. RAID 5 Scaling

The current RAID 5 algorithms are such that a RAID 5 volume should be able to produce a drive's bandwidth times N-1, where N is the number of drives participating in the volume during reads and writes.

Note: As performance continues to be enhanced that in future revision of the GigaStor firmware, the bandwidth of read operations should be between N-1 and N depending upon the size of the read-ahead and the current geometry of the drive.

3.1. Configuration

Starting at the drives the GigaStor should be able to consume 120 MB/sec (3 drives * 40 MB/sec) during writes and produce 180 MB/sec (3 drives * 60 MB/sec) during reads. Each volume is mapped to its own set of drives so as to ensure that the operation on each drive is sequential; meaning the only drive seeks required are a part of the action of moving sequentially through the drive. There are 4 such volumes and when running concurrently the GigaStor should be able to consume 480 MB/sec (120MB/sec * 4 volumes) during writes and produce 720 MB/sec (180 MB/sec * 4 volumes) during reads.

Each volume is mapped to its own Target which is connected to 2 Gigabit Ethernet ports that are configured for 2 connections per session. Two Gigabit ports, each capable of 112 MB/sec, should be able to handle the bandwidth of each volume (180 MB/sec reads; 120 MB/sec writes). A representation of the configuration is illustrated in Figure 4. The hosts' have the MC/S parameters set for 'least queue depth', although with these types of IO patterns the 'Round Robin' would have worked just as well.

The intent of the configuration is to ensure that the external elements; drives, volumes, Ethernet links, and hosts are not a limiting factor during the benchmark such that the results indicate the GigaStor's bandwidth capacity.

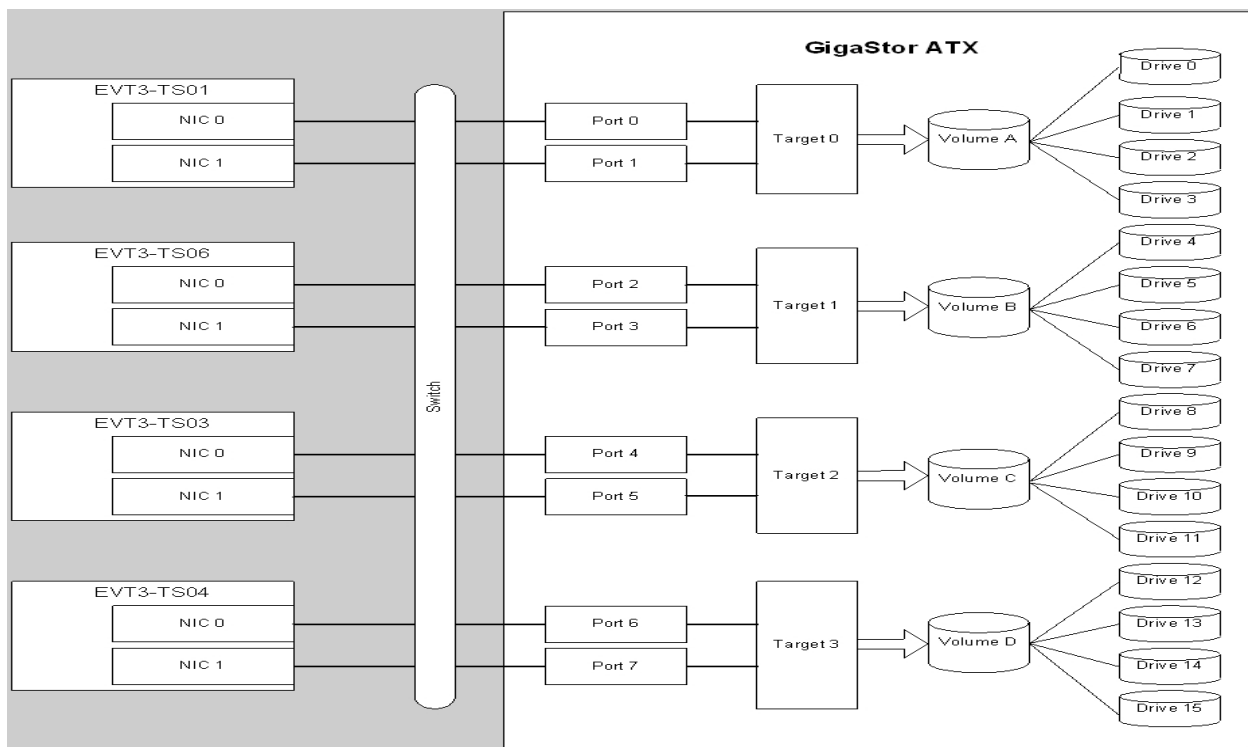


Figure 4 - RAID 5 Configuration (4 volumes of 4 drives)

3.2. Results

The spiral writes have a theoretical maximum transfer rate of 120 MB/sec per volume. Figure 5 illustrates that a single volume reached its peak of 114 MB/sec at the 16 KB transfer size; a 95% scaling factor. Two volumes reached their peak of 229 MB/sec at the 32 KB transfer size; a 95% scaling factor. Three volumes reached their peak of 340 MB/sec at the 64 KB transfer size; a 94% scaling factor. Four volumes reached their peak of 460MB/sec at the 64 KB transfer size; a 96% scaling factor.

The ability of the GigaStor to scale these operations while also dealing with the bandwidth consumption of producing the RAID 5 XOR activities was just about linear.

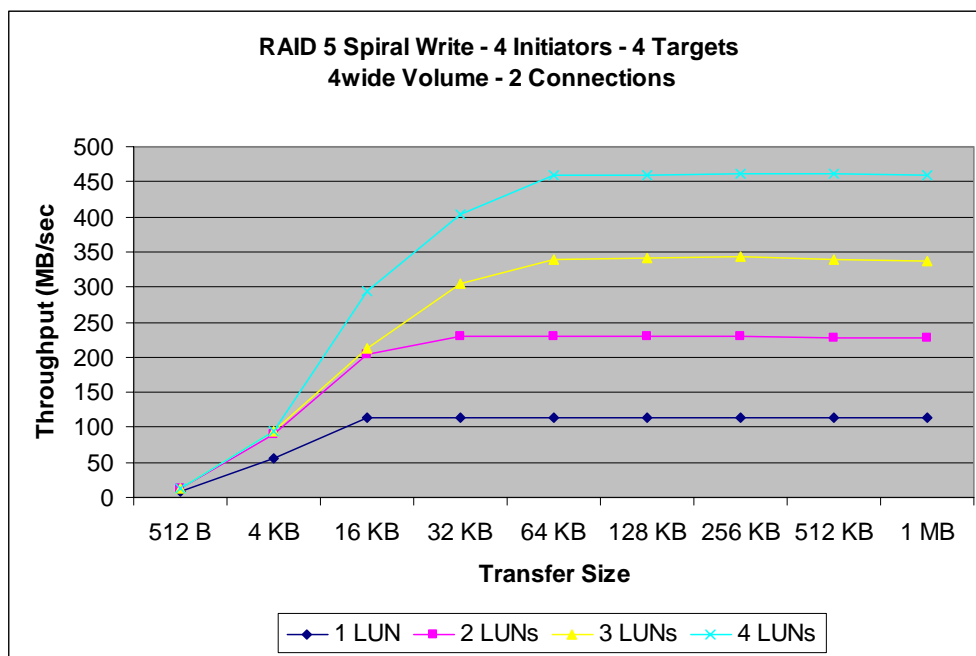


Figure 5 - RAID 5 Spiral Write

GigaStor Spiral Performance



The spiral reads have a theoretical maximum transfer rate of 180 MB/sec per volume. Figure 6 illustrates that a single volume reached its peak of 170 MB/sec at the 64 KB transfer size; a 94% scaling factor. Two volumes reached their peak of 343 MB/sec at the 64 KB transfer size; a 95% scaling factor. Three volumes reached their peak of 514 MB/sec at the 64 KB transfer size; a 95% scaling factor. Four volumes reached their peak of 672MB/sec at the 64 KB transfer size; a 93% scaling factor.

The read benchmark does illustrate anomalies with transfer sizes between 64 KB and 512 KB. The GigaStor scaled the volumes and the ports linearly.

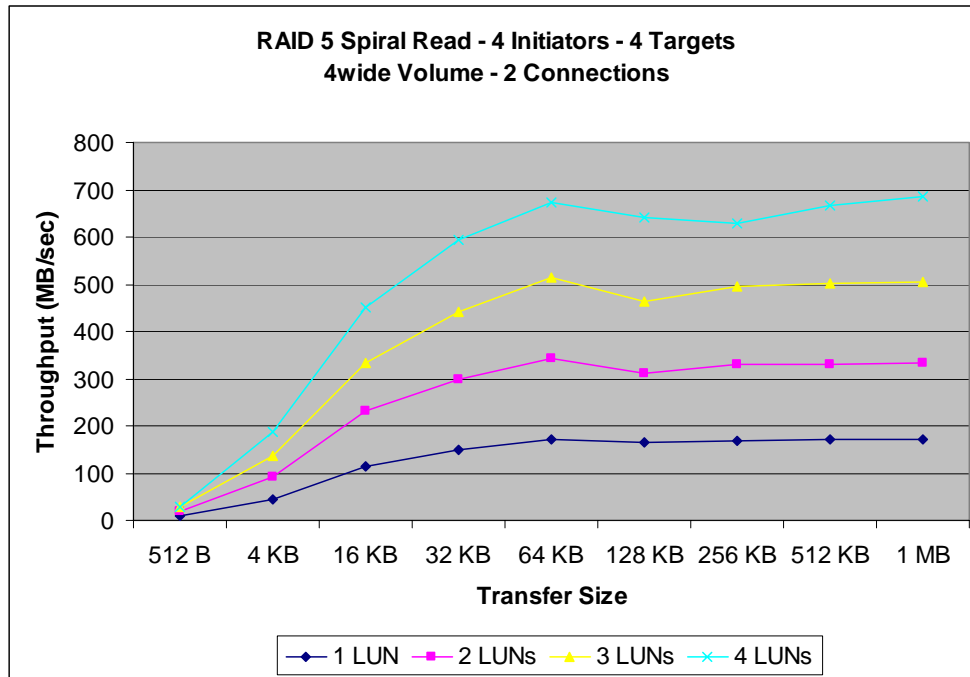


Figure 6 - RAID 5 Spiral Read

4. RAID 0 Scaling

4.1. Configuration # 1

This configuration is identical to the configuration defined in paragraph 3.1, RAID 5 scaling and illustrated in Figure 4, as the intent is to understand the relationship between RAID 0 performance and RAID 5 performance. The major difference is that the RAID 0 volumes have no parity data, so all drives are utilized during all operations. Because of this, the theoretical maximums that should be achieved are different.

Starting at the drives the GigaStor should be able to consume 160 MB/sec (4 drives * 40 MB/sec) during writes and produce 240 MB/sec (4 drives * 60 MB/sec) during reads per volume. Each volume is mapped to its own set of drives so as to ensure that the operation is spiral; meaning the only drive seeks required are a part of the action of moving sequentially through the drive. There are 4 such volumes and when running concurrently the GigaStor should be able to consume 640 MB/sec (160MB/sec * 4 volumes) during writes and produce 960 MB/sec (240 MB/sec * 4 volumes) during reads.

Each volume is mapped to its own Target which is connected to 2 Gigabit Ethernet ports that are configured for 2 connections per session. The two Gigabit ports, each capable of 112 MB/sec, are expected to provide more bandwidth capacity than the bandwidth of each volume during writes (160 MB/sec writes), however the reads are expected to be limited by the Ethernet ports as each volume produces 240 MB/sec versus the 224 MB/sec possible from 2 Ethernet ports. As such the reads should saturate the Ethernet ports. A representation of the configuration is illustrated in

Figure 7. The hosts' have the MC/S parameters set for 'least queue depth', although with these kind of IO patterns the simple 'Round Robin' would have worked just as well.

The intent of the configuration is to ensure that the external elements; drives, volumes, Ethernet links, and hosts are not a limiting factor during the write benchmark and to saturate the Ethernet ports during the reads.

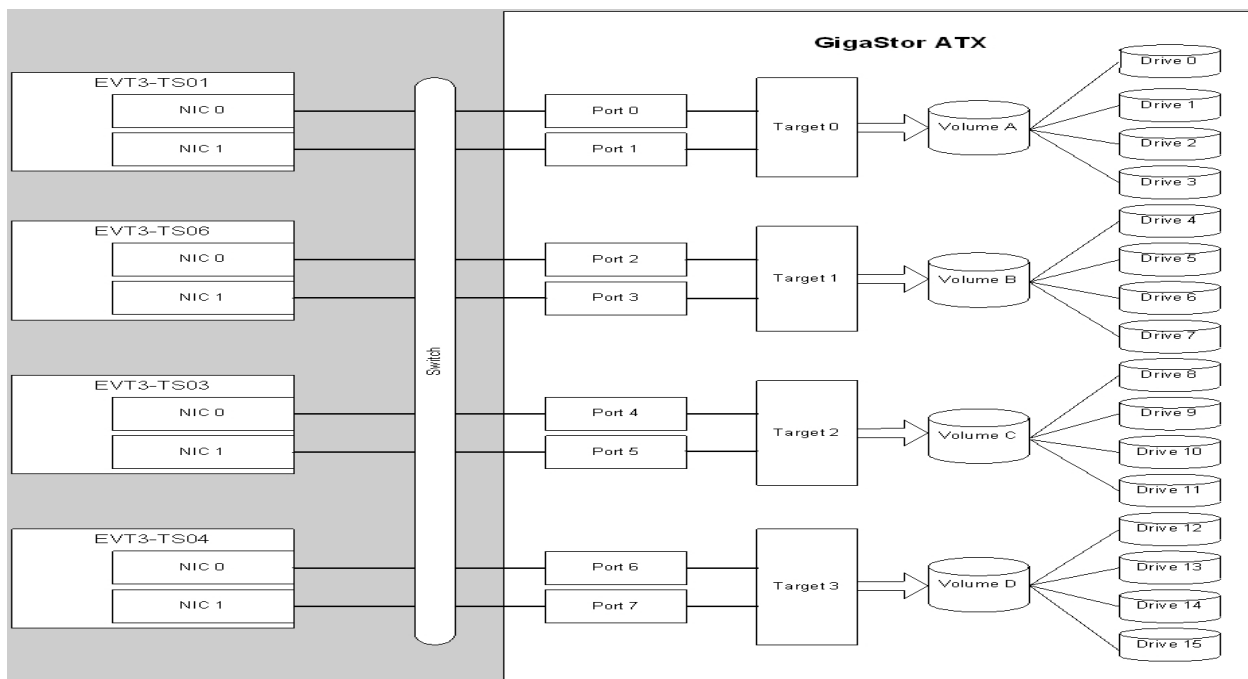


Figure 7 – RAID 0 Configuration (4 volumes of 4 drives)

4.2. Results

The spiral writes have a theoretical maximum transfer rate of 160 MB/sec per volume. Figure 8 illustrates that a single volume reached its peak of 158 MB/sec at the 32 KB transfer size; a 99% scaling factor. Two volumes reached their peak of 304 MB/sec at the 64 KB transfer size; a 95% scaling factor. Three volumes reached their peak of 456 MB/sec at the 256 KB transfer size; a 95% scaling factor. Four volumes reached their peak of 611 MB/sec at the 256 KB transfer size; a 95% scaling factor.

The ability of the GigaStor to scale these operations was just about linear.

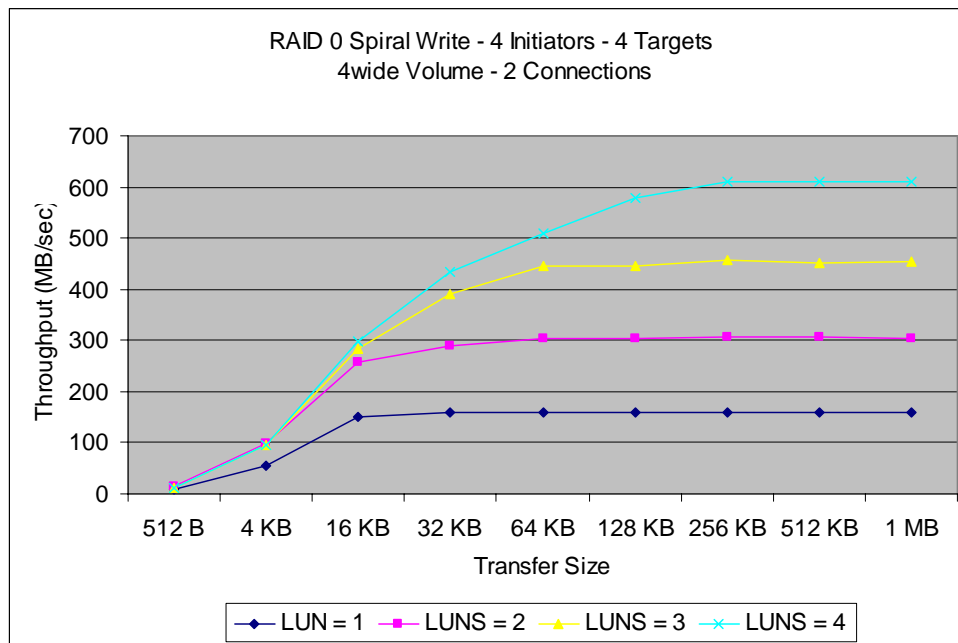


Figure 8 – RAID 0 Spiral Write (4 Hosts)

GigaStor Spiral Performance



The spiral reads have a theoretical maximum transfer rate of 224 MB/sec per volume as the limiting factor is the theoretical performance of the 2 Ethernet ports rather than the bandwidth capacity of the volumes themselves. Figure 9 illustrates that a single volume reached its peak of 183 MB/sec at the 256 KB transfer size; an 82% scaling factor. Two volumes reached their peak of 366 MB/sec at the 256 KB transfer size; an 82% scaling factor. Three volumes reached their peak of 553 MB/sec at the 256 KB transfer size; an 82% scaling factor. Four volumes reached their peak of 738 MB/sec at the 256 KB transfer size; an 82% scaling factor.

Though the actual performance numbers demonstrated high bandwidth from the GigaStor, it also illustrated a problem with the lack of obtaining saturation of each Ethernet port. A closer look at the details of the benchmark results revealed two factors.

1. IOMeter's data indicated that the hosts' dual processors averaged a 60% CPU overhead, and one of the processors was essentially saturated at a 99.8% of CPU overhead.
2. An analyzer view of the Ethernet ports illustrated that the host was having a very difficult time keeping commands queued inside the GigaStor.

Note: This report is based upon GigaStor's ability to scale performance based on the individual disk drives. However, in this case the components within the Ethernet path were limiting the performance of the GigaStor. If one views the benchmark from the standpoint of the performance being limited by the Ethernet's path, the GigaStor scaled quite well as each added path provided 100% more performance (183 MB/sec, 366 MB/sec, 553 MB/sec and 738 MB/sec).

This particular benchmark did not give much insight as to how the GigaStor would perform during RAID 0 spiral operations as the GigaStor was operating in a degraded performance mode because of the lack of queued IOs.

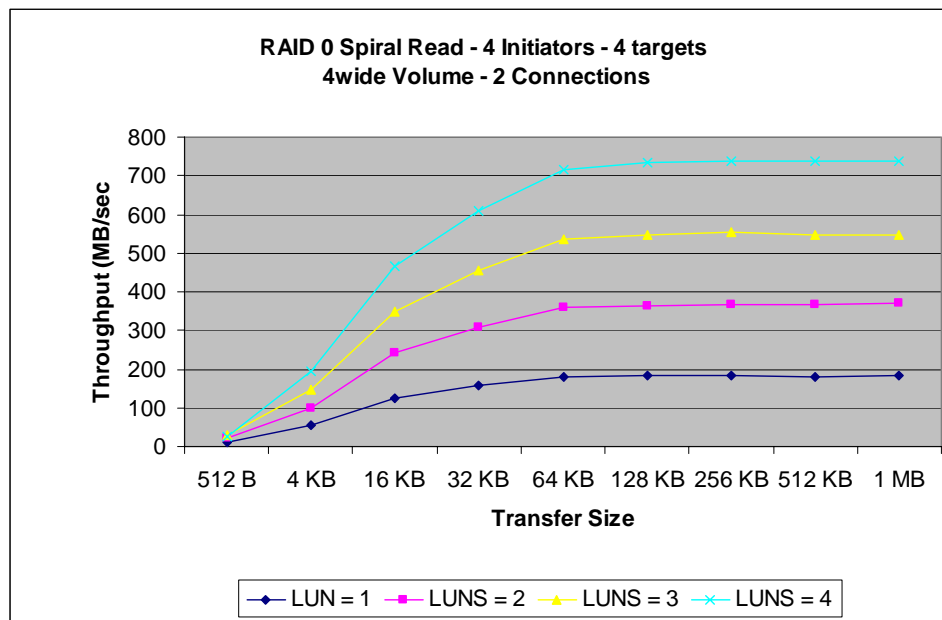


Figure 9 – RAID 0 Spiral Read (4 Hosts)

4.3. Configuration # 2

In light of the results obtained from the configuration defined in paragraph 4.1, and that the limiting factor was the CPUs and/or the NICs that were being used, the configuration was changed to have a single NIC card in 8 host systems instead of four host systems as illustrated in

Figure 10. Also, since the new host configuration was to have 8 initiators, the backend drive configurations modified the RAID 0 widths from four to two so that the bandwidth from the volumes would not saturate a single NIC.

Starting at the drives the GigaStor should be able to consume 80 MB/sec (2 drives * 40 MB/sec) during writes and produce 112 MB/sec during reads per volume. The volume can produce 120 MB/sec during reads (2 drives * 60 MB/sec), however each NIC has a theoretical bandwidth of 112 MB/sec. Each volume is mapped to its own set of drives so as to ensure that the operation is spiral; meaning the only drive seeks required are a part of the action of moving sequentially through the drive. There are 8 volumes and when running concurrently the GigaStor should be able to consume 640 MB/sec (80MB/sec * 8 volumes) during writes and produce 896 MB/sec (112 MB/sec * 8 Ethernet ports) during reads.

Each volume is mapped to its own Target which is connected to a single Gigabit Ethernet port that are configured for 1 connection per session. One Gigabit port capable of 112 MB/sec, should be able to handle the bandwidth of each volume during writes (80 MB/sec writes), however the reads should be limited by the Ethernet ports as each volume produces 120 MB/sec versus the 112 MB/sec possible from the single Ethernet port. As such the reads should saturate the Ethernet ports.

The intent of this test is to illustrate that the previous benchmark's values were limited by the Host and/or the NIC cards.

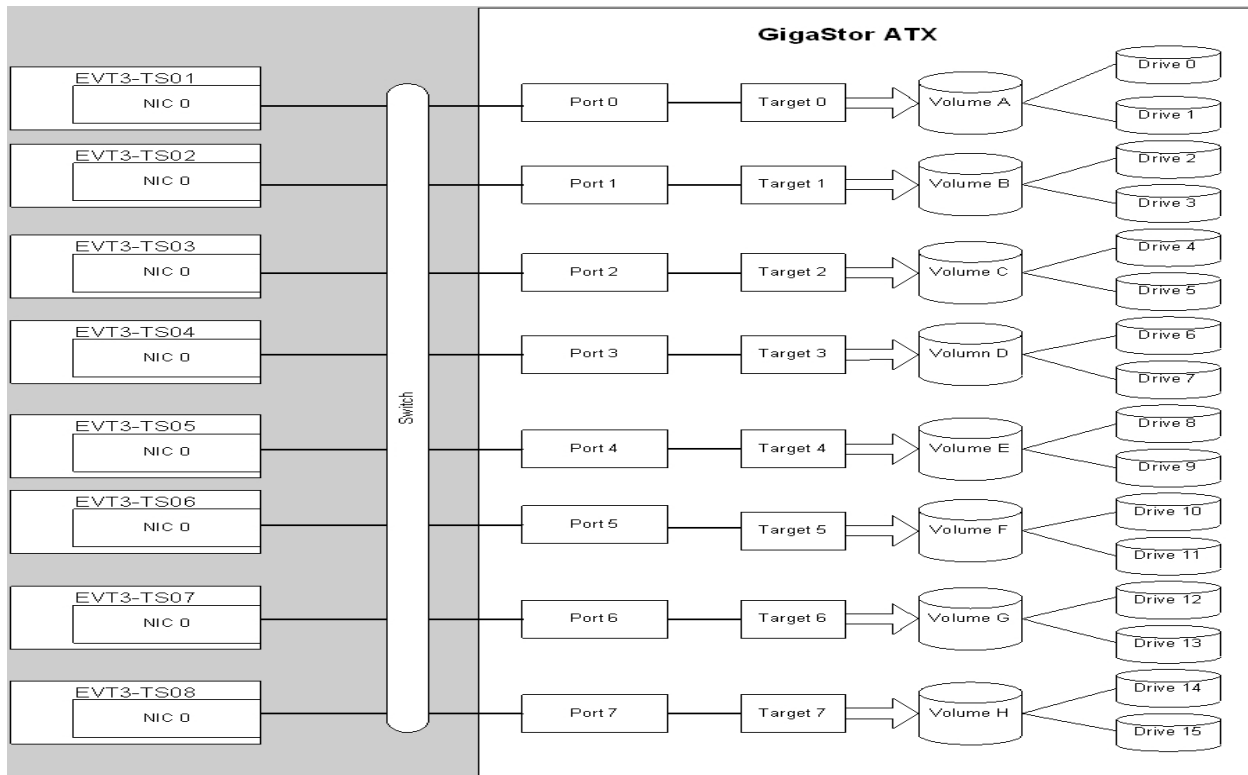


Figure 10 – RAID 0 Configuration (8 volumes of 2 drives)

4.4. Results

The spiral writes have a theoretical maximum transfer rate of 80 MB/sec per volume. Figure 11 illustrates that a single volume reach its peak of 80 MB/sec at the 32 KB transfer size; a 100% scaling factor. Two volumes reached their peak of 163 MB/sec at the 32 KB transfer size; a 100% scaling factor. Four volumes reached their peak of 326 MB/sec at the 32 KB transfer size; a 100% scaling factor. Eight volumes reached their peak of 620 MB/sec at the 1 MB transfer size; a 97% scaling factor.

Although this configuration was created to determine if the Host and/or NICs were impeding the read performance, there is a slight improvement for the writes.

The ability of the GigaStor to scale these operations was just about linear.

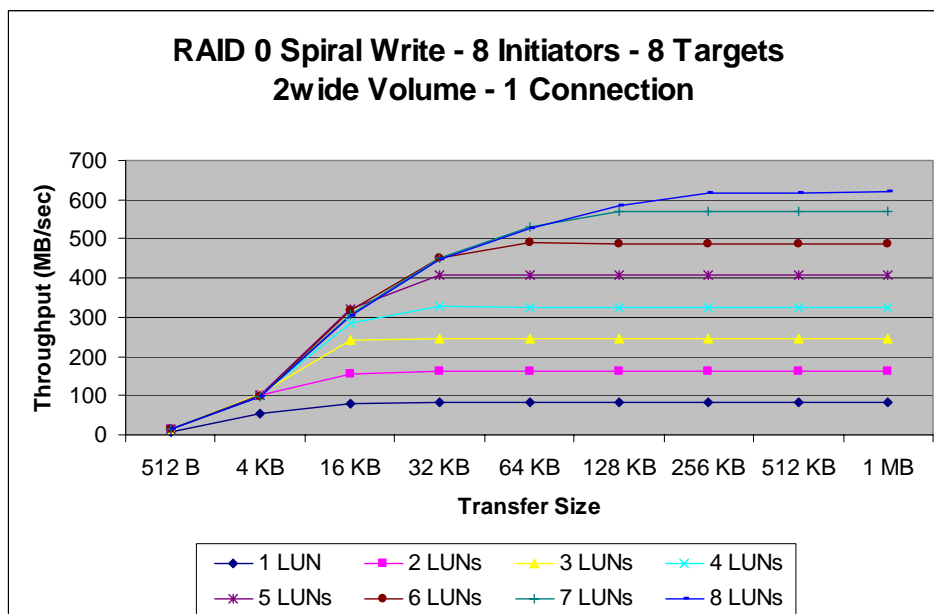


Figure 11 - RAID 0 Spiral Write (8 Hosts)

The spiral reads have a theoretical maximum transfer rate of 112 MB/sec per volume as the limiting factor is the theoretical performance of the Ethernet port. Figure 12 illustrates that a single volume reached its peak of 112 MB/sec at the 32 KB transfer size; a 100% scaling factor. Two volumes reached its peak of 224 MB/sec at the 32 KB transfer size; a 100% scaling factor. Four volumes reached their peak of 450 MB/sec at the 32 KB transfer size; a 100% scaling factor. Eight volumes reached their peak of 875 at the 32 KB transfer size; a 98% scaling factor.

In this configuration the GigaStor was able to scale the performance linearly. The width of the volumes may have been halved, but the number of volumes involved in the test was doubled such that the number of drives involved was a constant. The only variable in the test was the number of hosts and that each host was unburdened from managing two NICs. In conclusion, the configuration defined in paragraph 4.1 indicates that the hosts were limiting the performance to 738 MB/sec, whereas with this configuration the performance reached 875 MB/sec; a 16 % improvement. Not only was the scaling better along with the absolute performance numbers, the GigaStor was able to reach these performance numbers at a much smaller transfer size.

A closer look at the details of the benchmark results again revealed two factors.

1. IOMeter's data illustrated that the hosts' dual processors averaged a 42% CPU overhead, and neither processor was saturated.
2. An analyzer view of the Ethernet ports illustrated that each host was able to keep at least 2 commands queued inside the GigaStor.

The conclusion is that the Host and/or the NICs were the limiting factor in the benchmark configuration defined in paragraph 4.1. The GigaStor illustrated its ability to provide scalable performance across its volumes and its ports.

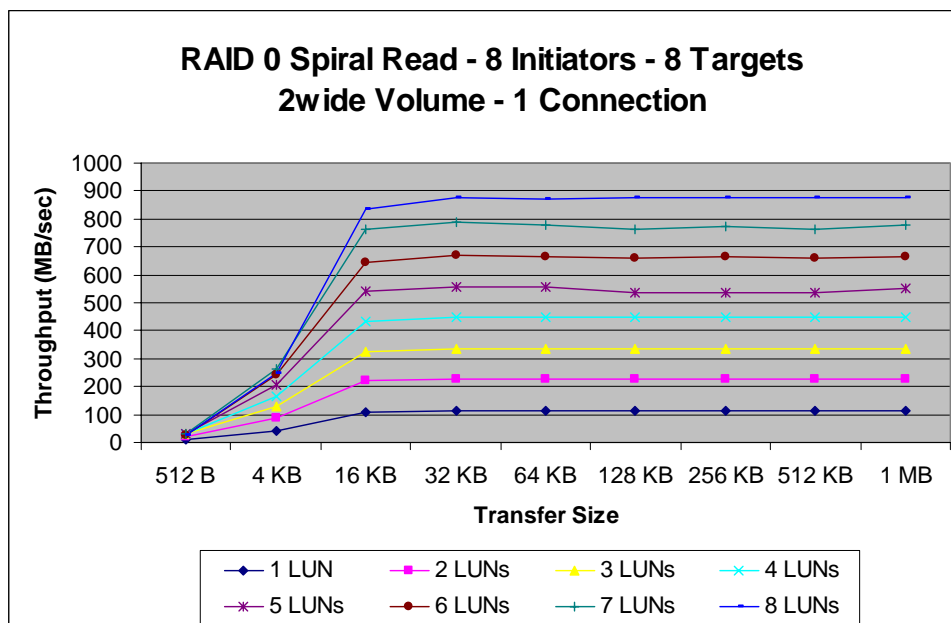


Figure 12 - RAID 0 Spiral Read (8 Hosts)

5. Multiple Volumes per Drive

All the previous benchmarks had the Initiator to volume mappings such that all transactions to the disk drives were contiguous. There was a single volume extent on the drive and the benchmark marched through the extent starting at the beginning and moving towards the end of the volume. The drives were not instructed to seek to other positions on the disk drive as part of the benchmark.

The following configuration produces multiple volume extents on each drive. The benchmark will then start marching through each extent with no knowledge of the other extents. This causes the commands that are received by the disk drives to move the heads from extent to extent as each volume is being processed; thus creating an IO stream that appears to be random in nature. The GigaStor's read-ahead and write coalescing algorithms should recognize the multiple sequential streams from the random IOs that are received and minimize the amount of seeking by the individual disk drives.

5.1. Configuration

The configuration defined in Figure 13 creates 4 volumes of a RAID 5 that are 16 wide. Although the GigaStor allows for flexibility in the creation of volumes, this particular configuration creates volumes in a very fixed manner; much like RAID controllers in the past.

Each volume consists of the same physical block numbers on each of the 16 drives. This means that volume 'A' is the first volume created on the drives. The space it consumes on each drive is identical. Following volume 'A' in ascending logical blocks is volume 'B', followed by volume 'C' and finally volume 'D'.

Each volume is mapped to its own Target which is connected to 2 Gigabit Ethernet ports that are configured for 2 connections per session. The hosts' have the MC/S parameters set for 'least queue depth'.

The intent of the configuration is to illustrate how the GigaStor deals with recognizing that the random IOs are essentially a number of spiral IO streams to different volumes that reside on the same disk drives; just how fair is the GigaStor when it comes to scheduling commands for different volumes that have extents residing on the same drive? This configuration; not the width of the volume, but rather the fact that a number of volumes are participating on a single drive is much more of a real world configuration. The previous configurations were illustrating the GigaStor's ability to move bandwidth, while this configuration illustrates GigaStor's ability to deal with multiple spiral streams, such a one might utilize to move multiple streams of rich data content.

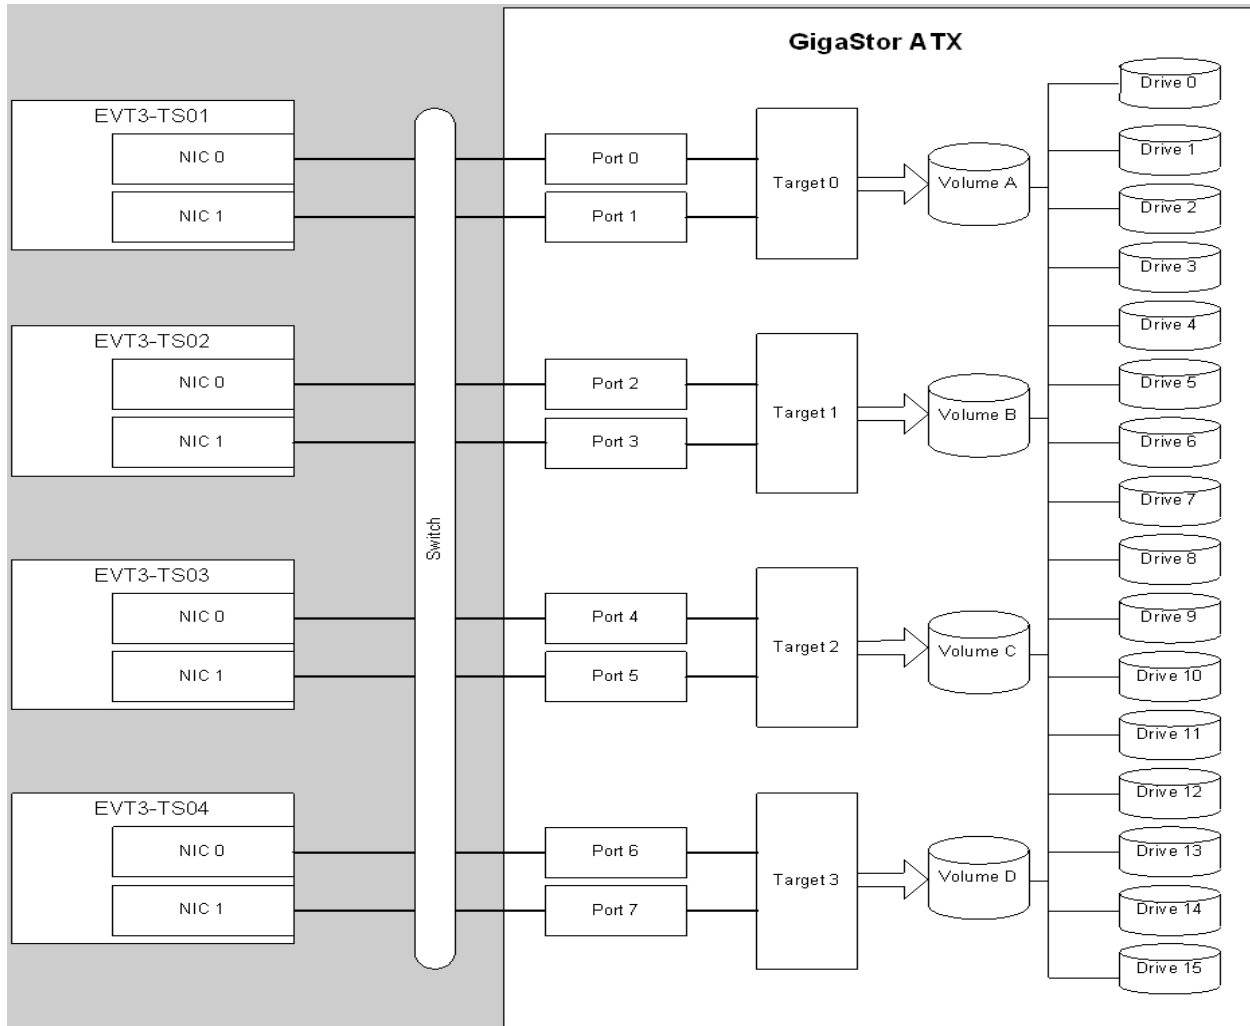


Figure 13 – RAID 5 Configuration (multiple volumes per drive)

5.2. Results

The scaling factors in this paragraph deal with the Ethernet ports and not the drives. The bandwidth capacity of the volume (15 drives * 40 MB/sec for writes or *60 MB/sec for reads) far exceeds the Ethernet port configuration.

Figure 14 illustrates that a single volume reached its peak of 185 MB/sec at the 1 MB transfer size; an 83% scaling factor. Two volumes reached their peak of 347 MB/sec at the 1 MB transfer size; a 78% scaling factor. Three volumes reached their peak of 478 MB/sec at the 512 KB transfer size; a 71% scaling factor. Four volumes reached their peak of 468 MB/sec at the 256 KB transfer size; a 52% scaling factor.

The scaling factors do not attempt to understand the time required seeking from volume to volume; they are based upon the Ethernet ports' capacity to move data. Considering the seek times the first three volumes performed very well. Obviously it can be seen that the 4th volume did not scale as well as the first three volumes.

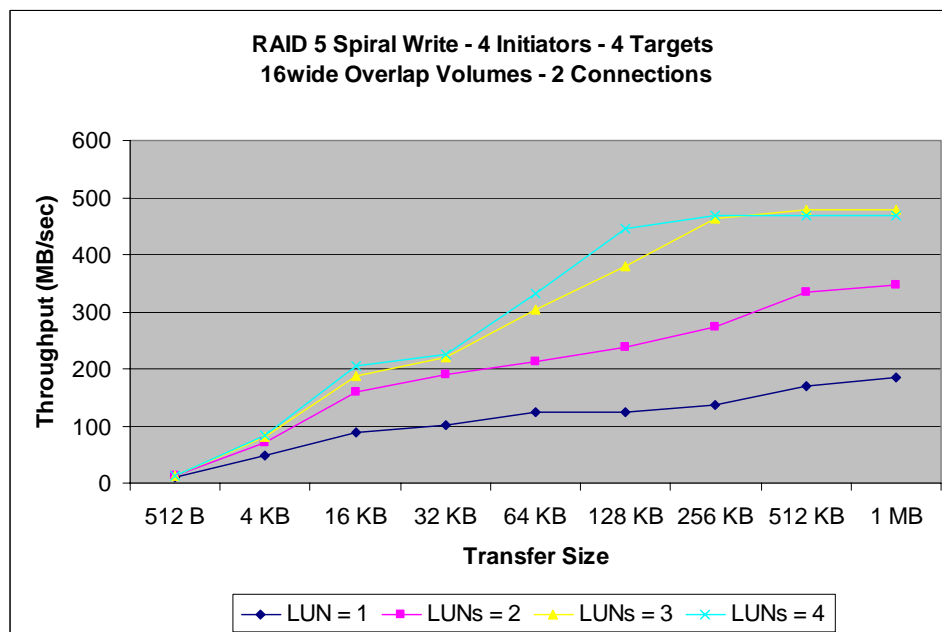


Figure 14 - RAID 5 Spiral Write (multiple volumes per drive)

Figure 15 illustrates that a single volume reached its peak of 182 MB/sec at the 128 KB transfer size; an 81% scaling factor. Two volumes reached their peak of 364 MB/sec at the 256 KB transfer size; an 81% scaling factor. Three volumes reached their peak of 547 MB/sec at the 256 KB transfer size; an 81% scaling factor. Four volumes reached their peak of 663 MB/sec at the 512 KB transfer size; a 74% scaling factor.

The above scaling factors do not attempt to understand the time required seeking from volume to volume; they are based upon the Ethernet ports' capacity to move data. Though the 4th volume caused a dip in the scaling factor, the four RAID 5 volumes performed consistently well delivering more than 600 MB/sec while seeking from volume to volume.

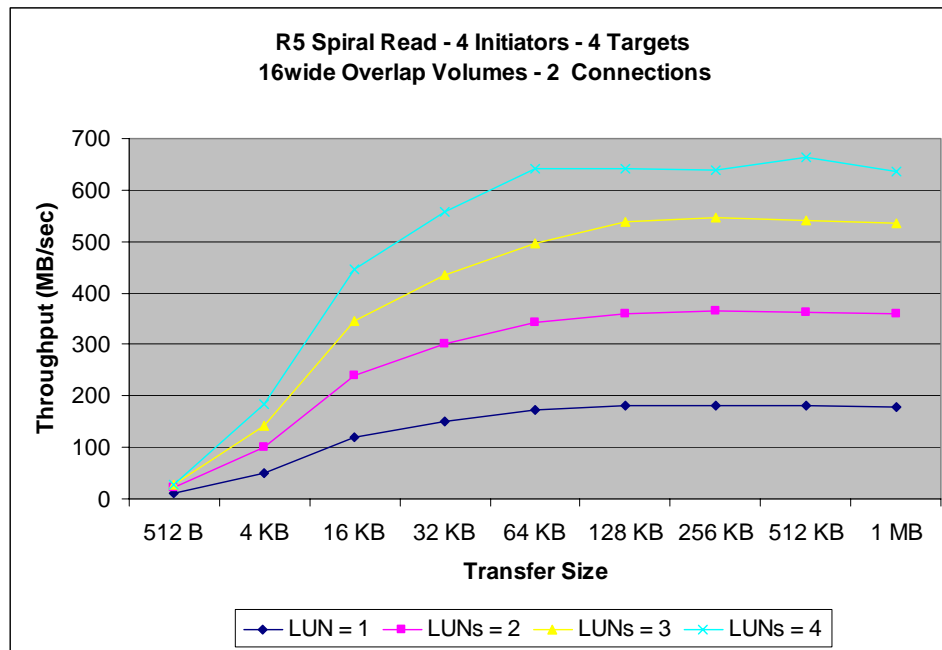


Figure 15 - RAID 5 Spiral Read (multiple volumes per drive)

6. Wide Connections

The previous benchmark sections illustrate the ability of the GigaStor to move large amounts of data, what is called bandwidth. It centered on the GigaStor's internal bandwidth, how much data the GigaStor could deliver to a variety of clients.

However, the flip side of this issue is to understand how much bandwidth can the GigaStor deliver to a single user? How can the GigaStor take a wide volume (the wider the volume the more bandwidth capacity that exists), and deliver all of its bandwidth to a single user? For example, a volume that is 8 drives wide could deliver 480 MB/sec to a single user. iSCSI provides a couple of ways of doing this; One is to configure a host with a number of NICs or iSCSI adapters and have the Operating System aggregate the data across those ports using a number of iSCSI sessions. In the Microsoft world, that would take the form of utilizing the MPIO driver. Another mechanism would be to again install a number of NICs or iSCSI adapters and configure them to run multiple connections per session.

This section will illustrate how the GigaStor performs when configured for many connections per session. This was chosen because while the MPIO driver appeared to work rather well with 2 sessions, anything greater than two caused performance to degrade. A look at the Ethernet ports showed that when configured for more than 2 sessions with MPIO that the performance on the Ethernet wire was slow and erratic. There were times when many milliseconds would elapse before the GigaStor would get its next command. It didn't seem to matter what type of NIC or iSCSI adapter that was used, in all cases the phenomenon was the same. iStor is not making any claims about the MPIO driver, just that the bandwidth from the volumes could not be exposed to a single host using MPIO.

The performance tests and configurations that are run in this section are not available in the currently shipping product; as the current product ships with a maximum number of connections set at two. For experimental basis, we raised that value to 4 connections per session to perform the following tests. A future enhancement that is planned for release of the GigaStor firmware will set that value to eight.

As mentioned earlier in this paper, the CPU utilization was very high when we had 2 NICs in the system. The hosts and NICs were exchanged for a higher end host with iSCSI accelerators. The iSCSI accelerators have illustrated extremely low host CPU utilization.

6.1. Wide Volume – 4 Connections/session

6.1.1. Configuration of a Wide RAID 0

Starting at the drives the GigaStor should be able to consume 280 MB/sec (7 drives * 40 MB/sec) during writes and produce 420 MB/sec (7 drives * 60 MB/sec) during reads. Figure 16 illustrates a single volume mapped to its own Target which is connected to 4 Gigabit Ethernet ports that are configured for 4 connections per session. Four Gigabit ports, each capable of 112 MB/sec, should be able to handle the bandwidth of the volume. The single host has the MC/S parameters set for 'least queue depth'.

The intent of the configuration is to ensure that the external elements; drives, volumes, Ethernet links, and hosts are not a limiting factor while delivering high bandwidth to a single user.

Unlike the previous benchmark runs, this run only tests the GigaStor at the 64 KB transfer size as this test was not intended to test the GigaStor's ability to scale, but rather its ability to move the bandwidth from wide volumes to a single user.

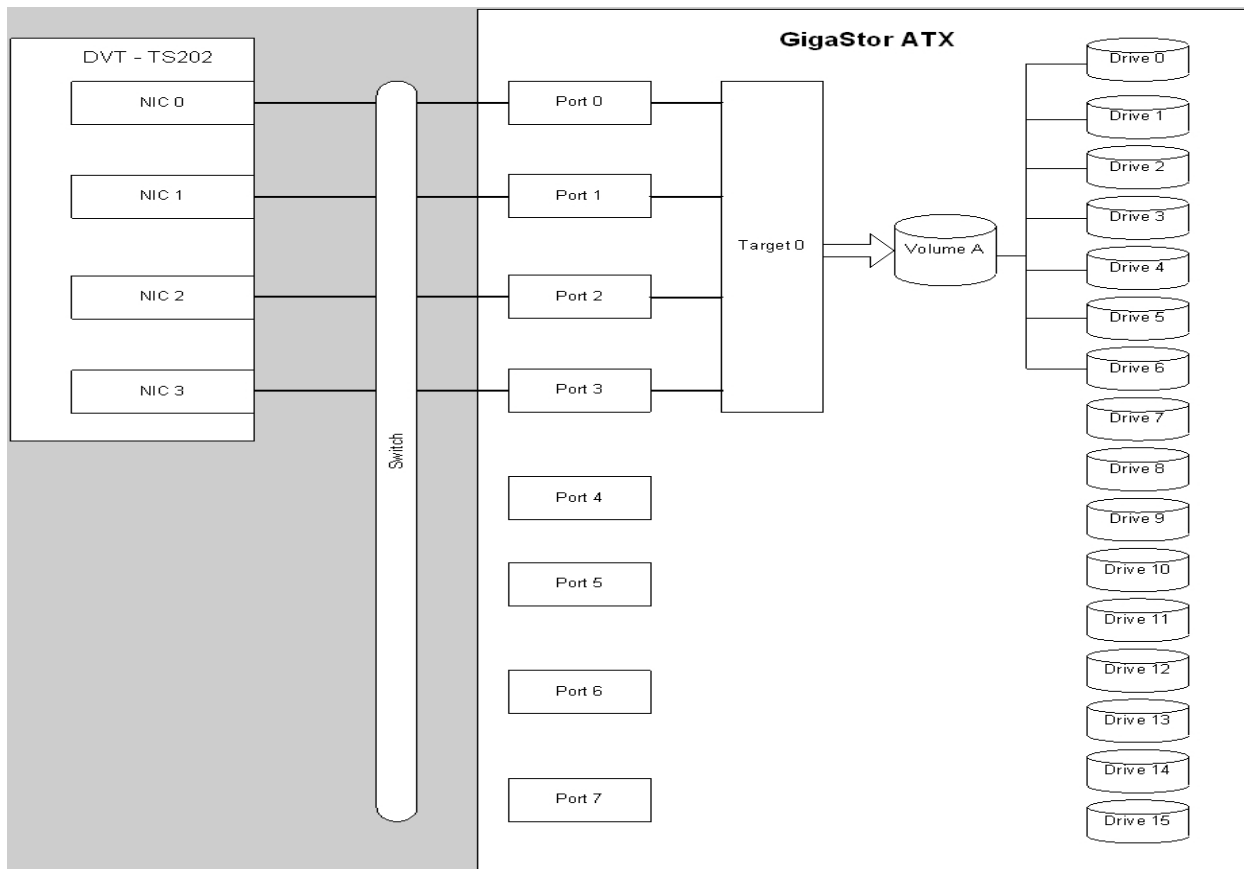


Figure 16 – Configuration of a RAID 0 with wide connections

6.1.2. Results

Figure 17 illustrates that the spiral write resulted in bandwidth of 97 MB/sec during a single connection, 176 MB/sec during two connections, 241 MB/sec during three connections and 271 MB/sec during four connections. The volume is capable of consuming 280 MB/sec (7 drives * 40 MB/sec); this corresponds to a 97% scaling factor.

The spiral read resulted in bandwidth of 112 MB/sec during a single connection, 225 MB/sec during two connections, 334 MB/sec during three connections and 401 MB/sec during four connections. The first three connections scaled at 100%, while the 4th connection scaled at 95%. Note that the first 3 connection's calculations were based upon the performance of each connection. The 4th calculation used the volume's bandwidth capacity of 420 MB/sec (7 drives * 60 MB/sec) rather than the capacity of 4 connections because the volume's configuration could not produce enough bandwidth to saturate the 4 connections.

Theoretically the GigaStor should have delivered approximately 420 MB/sec during the connections. The limiting factor found was the amount of bandwidth available from the 7 wide RAID 0 volume. To understand if the GigaStor's bandwidth capacity was causing this problem, a number of other hosts were added to the configuration as illustrated in paragraph 6.1.3.

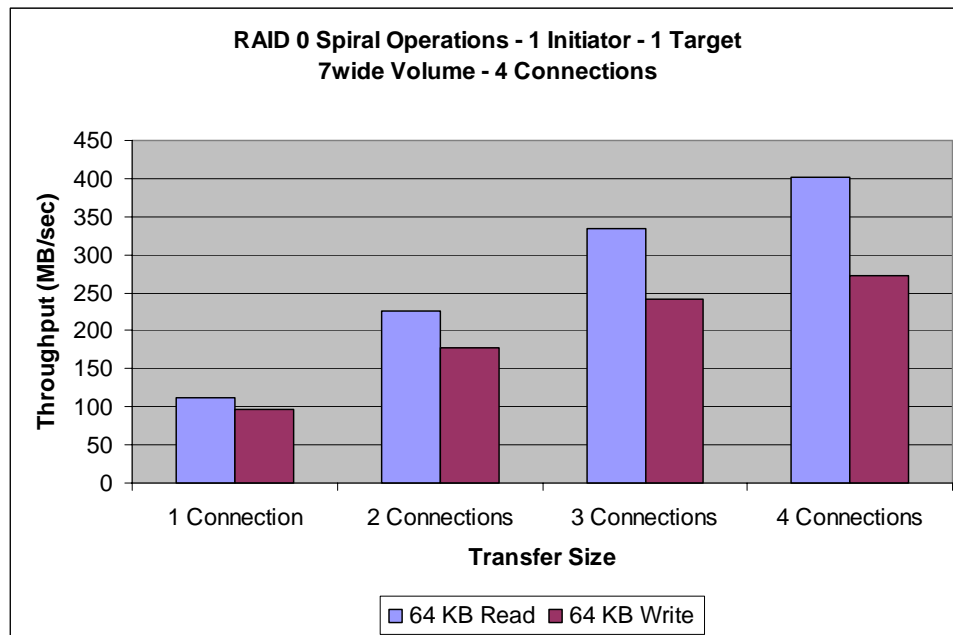


Figure 17 – RAID 0 Spiral Write/Read (wide connections)

6.1.3. Configuration of Added RAID 0 Volumes

Whereas the configuration in paragraph 6.1.1 did not scale adequately in the manner which one would have expected, the configuration depicted in

Figure 18 has added four 2-wide RAID 0 volumes along with 4 more host, each configured with a single NIC to the already configured 7 wide RAID 0. The intent is to see if the internals of the GigaStor were impeding the ability to reach the theoretical values.

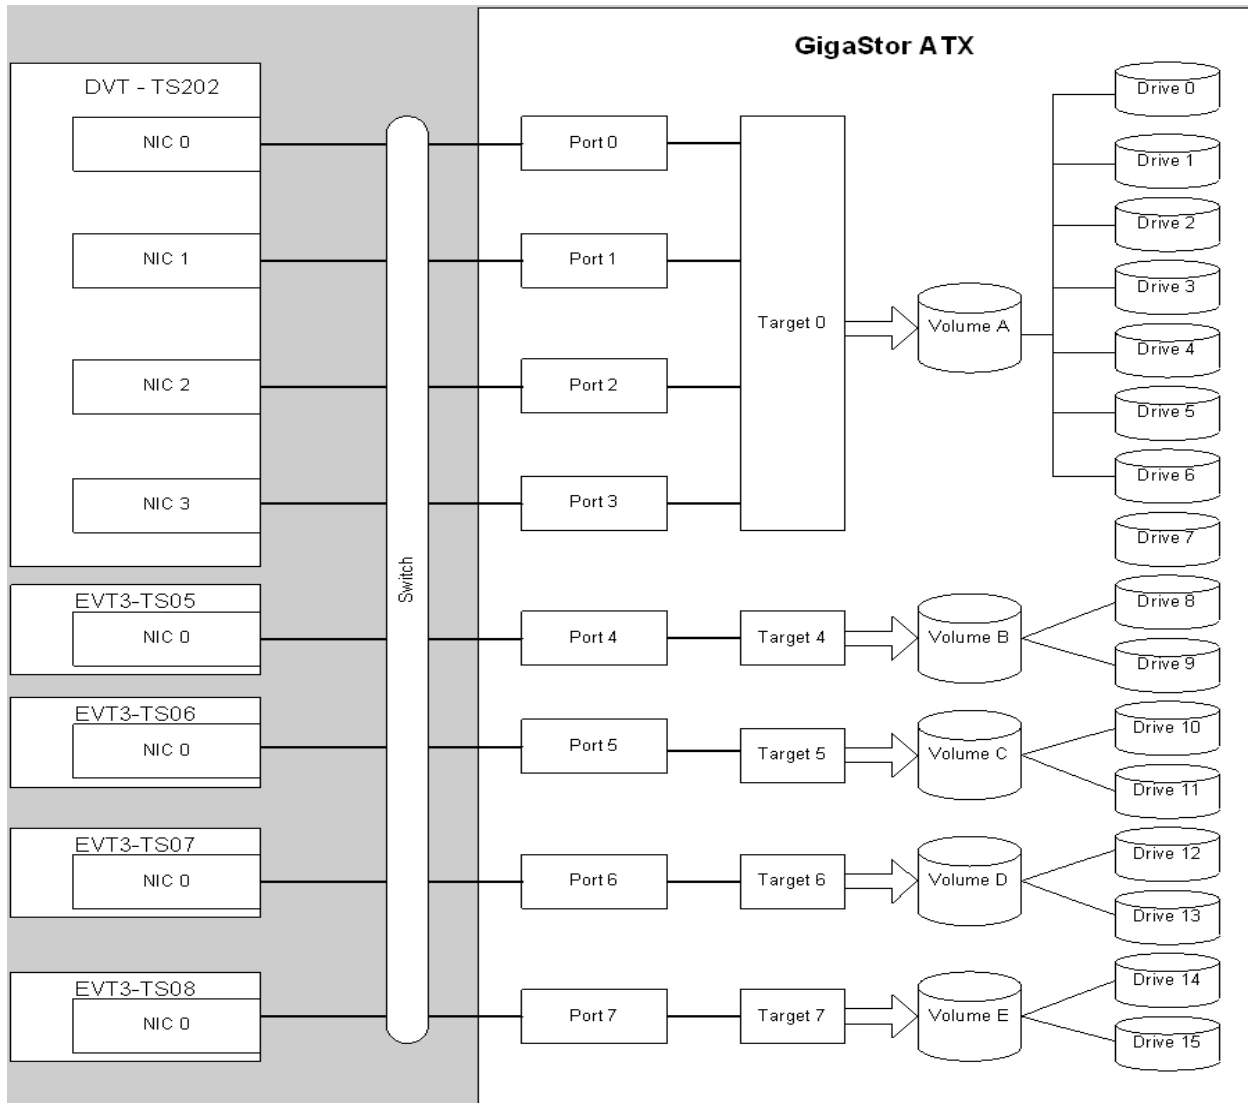


Figure 18 - Configuration Spiral RAID 0 (wide and narrow connections)

6.1.4. Results

The GigaStor was able to continue to provide more bandwidth as the additional computers and connections were added. The system with the 4 NICs utilizing 4 connections per session continued to provide approximately the same amount of performance as in the earlier test (385 MB/sec); however the 4 additional hosts were able to receive close to their maximum bandwidths as the GigaStor was able to provide the additional bandwidth of 107 MB/sec, 110 MB/sec, 105 MB/sec and 107 MB/sec for the four host respectively.

This illustrated that the GigaStor was able to provide much more bandwidth than the 420 MB/sec that was seen in the configuration defined in paragraph 6.1.1.

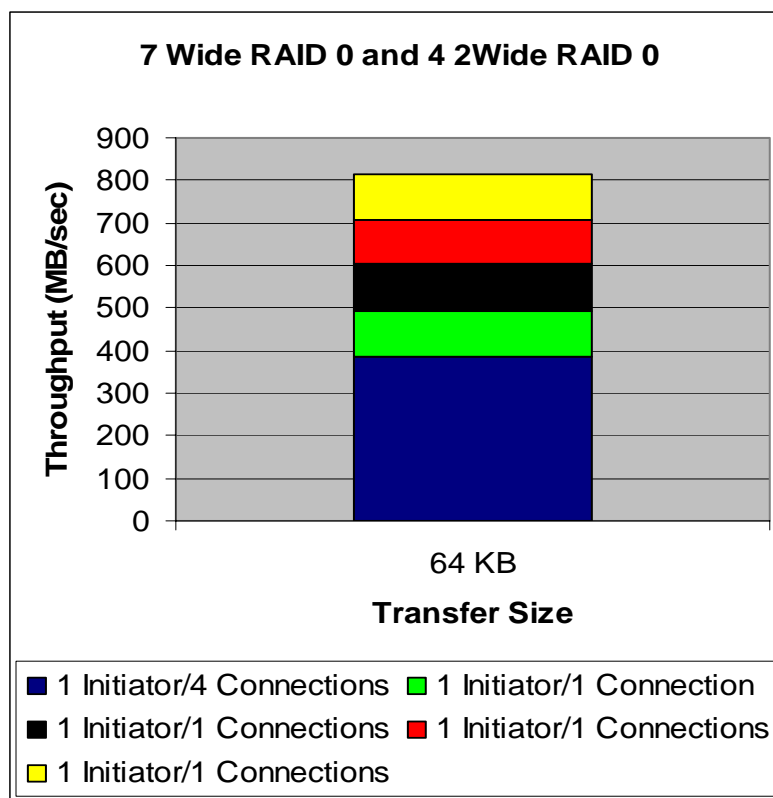


Figure 19 – Spiral Read RAID 0 (wide & narrow connections)

6.1.5. Configuration of RAID 5 with RAID 0

The configuration defined in

Figure 20 was chosen to understand the performance relationships between RAID 0 and RAID 5 when configured with many connections per session. In essence how well would the GigaStor perform when the wide volume was a RAID 5? The only difference between this configuration and the one defined in paragraph 6.1.3 is that eight drives are used with the RAID 5; only 7 were used with the RAID 0.

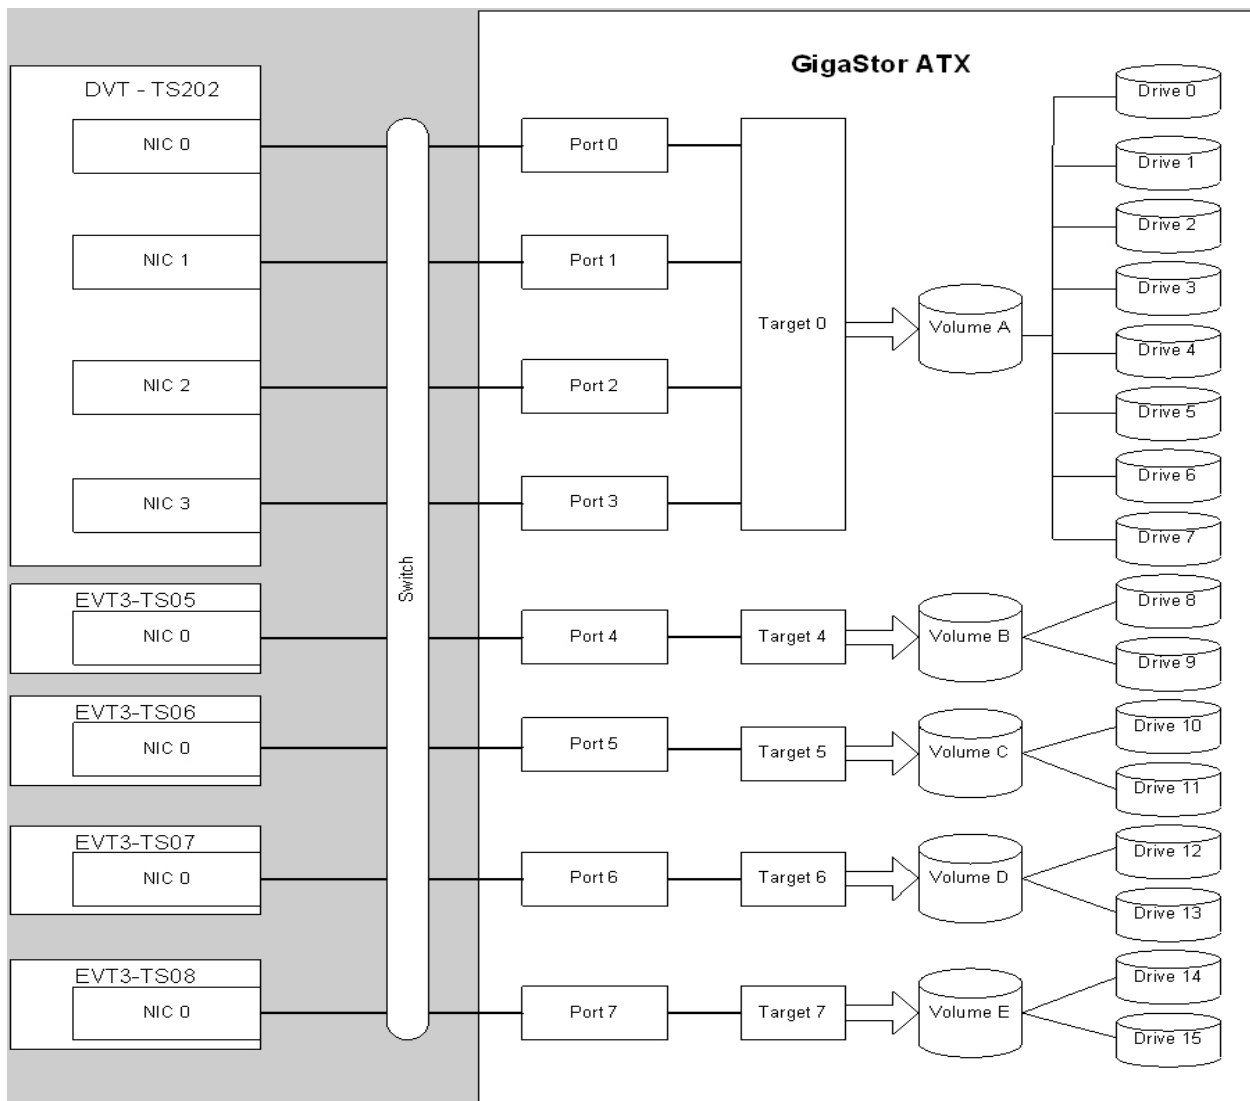


Figure 20 - Configuration Spiral RAID 5 Wide & Spiral RAID 0 Narrow

6.1.6. Results

Figure 21 illustrates that the results were almost identical. The system with the 4 connections was able to sustain 375 MB/sec while the 4 RAID 0 systems had bandwidth of 107 MB/sec, 109 MB/sec, 107 MB/sec and 107 MB/sec respectively during spiral reads.

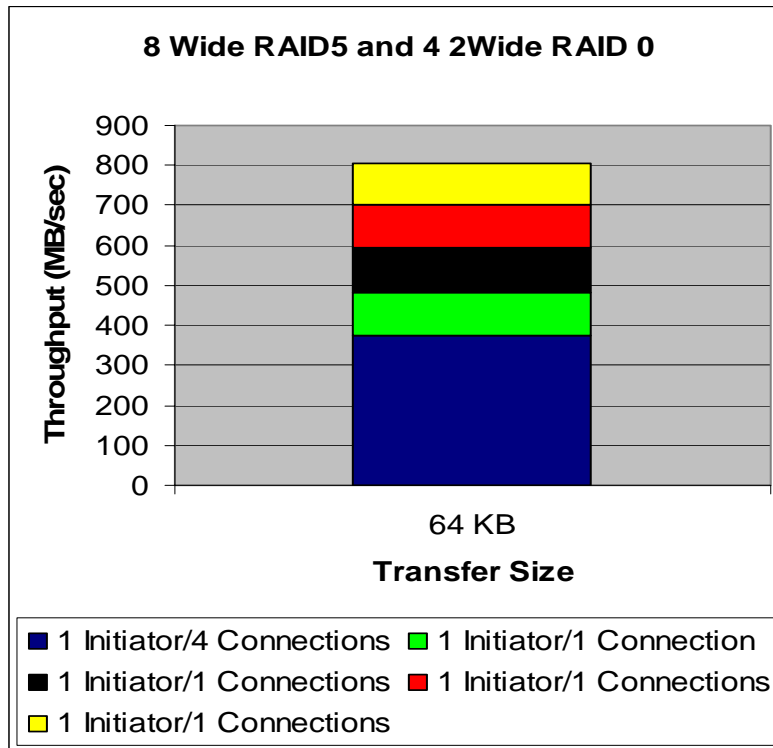


Figure 21 - Spiral Read RAID 5 & 0 (wide & narrow connections)

7. Appendix A

7.1. Test Equipment Description

All the tests performed utilized the below equipment except for those described in paragraph 6, Wide Connections.

Host Computer				
Name	CPUs	Memory	MS Initiator	Operation System
EVT3-TS01	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Enterprise Edition, SP1; DELL 850
EVT3-TS02	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Enterprise Edition, SP1; DELL 850
EVT3-TS03	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Enterprise Edition, SP1; DELL 850
EVT3-TS04	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Enterprise Edition, SP1; DELL 850
EVT3-TS05	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Enterprise Edition, SP1; DELL 850
EVT3-TS06	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Enterprise Edition, SP1; DELL 850
EVT3-TS07	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Enterprise Edition, SP1; DELL 850
EVT3-TS08	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Enterprise Edition, SP1; DELL 850
Network Interface Cards				
Name	Model	Product Description	Version	Description
Broadcom	NetXtreme	Gigabit Ethernet	7.100.0.0	64-bit 133MHz PCI-X, single-port, copper
Broadcom	NetXtreme	Gigabit Ethernet	7.100.0.0	64-bit 133MHz PCI-X, single-port, copper
Broadcom	NetXtreme	Gigabit Ethernet	7.100.0.0	64-bit 133MHz PCI-X, single-port, copper
Broadcom	NetXtreme	Gigabit Ethernet	7.100.0.0	64-bit 133MHz PCI-X, single-port, copper
Broadcom	NetXtreme	Gigabit Ethernet	7.100.0.0	64-bit 133MHz PCI-X, single-port, copper
Broadcom	NetXtreme	Gigabit Ethernet	7.100.0.0	64-bit 133MHz PCI-X, single-port, copper
Broadcom	NetXtreme	Gigabit Ethernet	7.100.0.0	64-bit 133MHz PCI-X, single-port, copper
Storage Controller				
Name	Model	Buffer Memory	Version	Description
hn0000000052300018	GigaStor ATX	512 MB	1.1.0.12	8 - 1 Gb ports, Hardware Rev. 00-01
Disk Drives for 8x1GbE				
Name	Model	Serial Number	Version	Description
Drive 0	7L250S0	L50558TH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 1	7L250S0	L504WALH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 2	7L250S0	L5054CTH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 3	7L250S0	L5053FLH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 4	7L250S0	L504W33H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 5	7L250S0	L5053MLH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 6	7L250S0	L504W30H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 7	7L250S0	L504W1QH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 8	7L250S0	L505J10H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 9	7L250S0	L504W21H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 10	7L250S0	L5063H3H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 11	7L250S0	L5063FVH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 12	7L250S0	L504W4MH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 13	7L250S0	L5054P1H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 14	7L250S0	L5053TDH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 15	7L250S0	L505J6NH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150

Figure 22 – Equipment List

iSCSI Parameters	
Name	Negotiated Value
MaxRecvDataSegmentLength (Target)	8192
MaxRecvDataSegmentLength (Initiator)	65536
HeaderDigest	None
DataDigest	None
OfMarker	No
IfMarker	No
OfMarkInt	Irrelevant
IfMarkInt	Irrelevant
MaxBurstLength	262144
FirstBurstLength	Irrelevant
DefaultTime2Wait	2
DefaultTime2Retain	0
MaxOutstandingR2T	1
MaxConnections	As stated in the configuration of each test
ErrorRecoveryLevel	0
InitialR2T	Yes
ImmediateData	No
DataPDUInOrder	Yes
DataSequenceInOrder	Yes

Figure 23 – iSCSI Parameters

7.2. Test Equipment Description for Wide Connections

Host Computer				
Name	CPUs	Memory	MS Initiator	Operation System
DVT-TS202	4 - 2.7 GHz Xeon	2 GB	2	Windows Server 2003, Standard Enterprise
EVT3-TS05	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Enterprise Edition, SP1; DELL 850
EVT3-TS06	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Enterprise Edition, SP1; DELL 850
EVT3-TS07	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Enterprise Edition, SP1; DELL 850
EVT3-TS08	3GHz PentiumD (dual core)	512 MB	2	Windows Server 2003, Standard Edition, SP1; DELL 850
Network Interface Cards				
Name	Model	Product Description	Version	Description
Alacritech Accelerator	SES1001T	100007 Rev.50	7.1.1.0	64-bit 133MHz PCI-X, single-port, copper
Alacritech Accelerator	SES1001T	100007 Rev.50	7.1.1.0	64-bit 133MHz PCI-X, single-port, copper
Alacritech Accelerator	SES1001T	100007 Rev.50	7.1.1.0	64-bit 133MHz PCI-X, single-port, copper
Alacritech Accelerator	SES1001T	100007 Rev.50	7.1.1.0	64-bit 133MHz PCI-X, single-port, copper
Broadcom	NetXtreme	Gigabit Ethernet	7.100.0.0	64-bit 133MHz PCI-X, single-port, copper
Broadcom	NetXtreme	Gigabit Ethernet	7.100.0.0	64-bit 133MHz PCI-X, single-port, copper
Broadcom	NetXtreme	Gigabit Ethernet	7.100.0.0	64-bit 133MHz PCI-X, single-port, copper
Broadcom	NetXtreme	Gigabit Ethernet	7.100.0.0	64-bit 133MHz PCI-X, single-port, copper
Storage Controller				
Name	Model	Buffer Memory	Version	Description
hn0000000052300018	GigaStor ATX	512 MB	1.1.0.12	8 - 1 Gb ports, Hardware Rev. 00-01
Disk Drives for 8x1GbE				
Name	Model	Serial Number	Version	Description
Drive 0	7L250S0	L50558TH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 1	7L250S0	L504WALH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 2	7L250S0	L5054CTH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 3	7L250S0	L5053FLH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 4	7L250S0	L504W33H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 5	7L250S0	L5053MLH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 6	7L250S0	L504W30H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 7	7L250S0	L504W1QH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 8	7L250S0	L505J10H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 9	7L250S0	L504W21H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 10	7L250S0	L5063H3H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 11	7L250S0	L5063FWH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 12	7L250S0	L504W4MH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 13	7L250S0	L5054P1H	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 14	7L250S0	L5053TDH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150
Drive 15	7L250S0	L505J6NH	BANC1E00	Maxtor SATA drive, 250 GB, 7200 RPM, 16 MB buffer, SATA 150

Figure 24 - Equipment List - Wide Connections

iSCSI Parameters	
Name	Negotiated Value
MaxRecvDataSegmentLength (Target)	8192
MaxRecvDataSegmentLength (Initiator)	65536
HeaderDigest	None
DataDigest	None
OfMarker	No
IfMarker	No
OfMarkInt	Irrelevant
IfMarkInt	Irrelevant
MaxBurstLength	262144
FirstBurstLength	Irrelevant
DefaultTime2Wait	2
DefaultTime2Retain	0
MaxOutstandingR2T	1
MaxConnections	4 for DVT - TS202; 1 for EVT-TS05/EVT-TS06/EVT-TS07/EVT-TS08
ErrorRecoveryLevel	0
InitialR2T	Yes
ImmediateData	No
DataPDUInOrder	Yes
DataSequenceInOrder	Yes

Figure 25 - iSCSI Parameters - Wide Connections

8. [Appendix B](#)

8.1. Microsoft File System Writes

The GigaStor should provide spiral read and write performances that will saturate the network transports. Of course, IO sizes that are very small, less than 8KB, may show some variances as the processing overhead becomes a greater percentage of the entire time to process the command and move the data.

These benchmarks were run with IOMeter and they were running against a RAW disk drive, meaning that the file system was not involved. Other benchmarks have shown that when running through the file system that the write performance will drop off suddenly, approximately 90% at the 512 KB and 1 MB transfer sizes. To avoid this type of system behavior please implement the following directions.

Microsoft Initiator version 2.0 or later is required. It can be downloaded from Microsoft's website at <http://www.microsoft.com/downloads>.

iStor recommends that users only manage disks on any Windows platform using Disk Manager accessed through the "Disk Management" option in the Computer Manager. Disk Manager is aware of the operating system, iSCSI, and MPIO and therefore displays the proper options for managing disks. The options you are looking for appear on the Properties dialog on the pop-up menu when the user right mouse clicks on a disk in Disk Manager as shown in Figure 26.

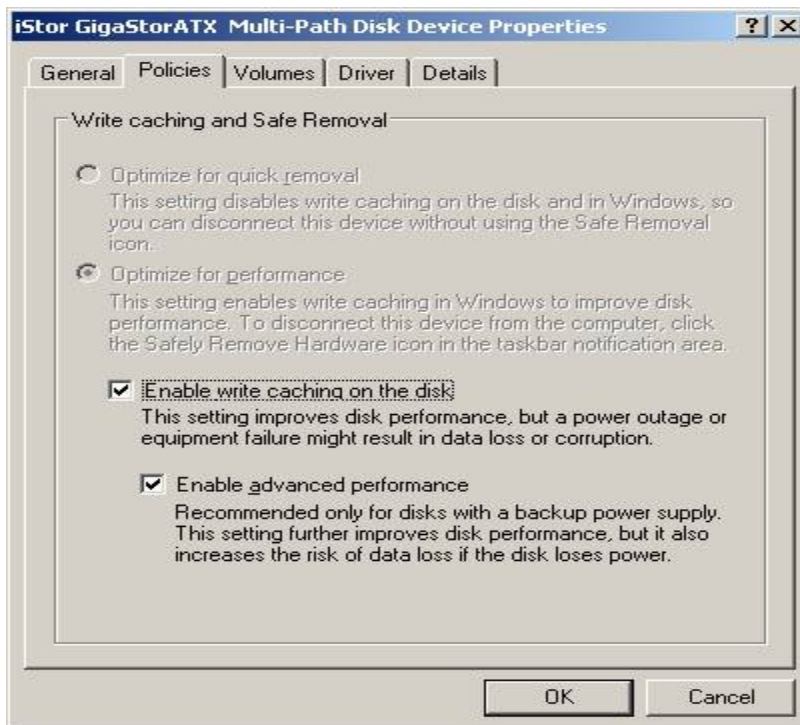


Figure 26 - Disk Manager POP-UP

All GigaStor products include power backup schemes that provide for data protection in the event of a power loss. Users that want full write performance should check both boxes.