



D-Link System, Inc.™
xStack Storage DSN-Series SAN Arrays
Performance Notes & Tips
Rev. B

© 2008 D-Link Systems, Inc. All Rights Reserved

D-Link Systems, 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. D-Link Systems, 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 D-Link Systems, Inc.

The information is provided “as is” without warranty of any kind and is subject to change without notice. The only warranties for D-Link Systems 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. D-Link Systems shall not be liable for technical or editorial errors or omissions contained herein.

Copyright © 2008 D-Link Systems, 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 trademarks of Microsoft Corporation.

Oracle® is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

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.

D-Link Systems, Inc.
17595 Mount Herrmann Street
Fountain Valley, CA 92708
www.DLink.com

Overview

Many customers who purchase one of our D-Link SAN products are curious how higher speeds may be obtained. Attaching a single LAN card to our SAN will often show slow results. The customer will then ask, “How can I get these throughput speeds that D-Link claims are possible?”

The purpose of this document is to provide a simple scenario wherein a baseline speed is achieved with a standard “no-frills” file server, and then simple performance-enhancing tips are tested that show how speed may be increased. Our performance tests make use of the Iometer software package available for free as a download at www.iometer.org.

When configuring your server(s), be sure to use the following tips:

- Increase the number of iSCSI sessions you use. Be sure each session addresses an individual target/volume combination as shown in Testing Configuration 3.
- Use the maximum MTU frame size allowed. In the case of the DSN products, this will be 9,014 bytes.
- Use a network interface card that can provide TCP offload capabilities.
- The RAID level and number of drives used is important. A RAID-0 stripe requires no parity calculations like RAID-5. Also when testing, try and use as many hard drives as possible to spread the data across many locations.
- Using a hard drive with a faster “sustained” throughput rate will help. Some hard drive data sheets include throughput stated as a peak value. (In our examples we used standard Hitachi DeskStar drives with an average sustained throughput of ~65MBps during sequential reads.)
- The server itself can have a great bearing on your testing. Our example uses a low-end Dell PowerEdge SC1430 with a 3GHz CPU. However, if you use a much more powerful server, or a combination of multiple servers working together, then your throughput speed can rise dramatically to near the theoretical limit of the SAN itself.

Please bear in mind that the results shown in this document are simply to provide tips that can be used to increase the speed of the DSN-series SAN array in your own testing. A more powerful server and faster hard drives would have driven our speed much higher.

Testing Environment

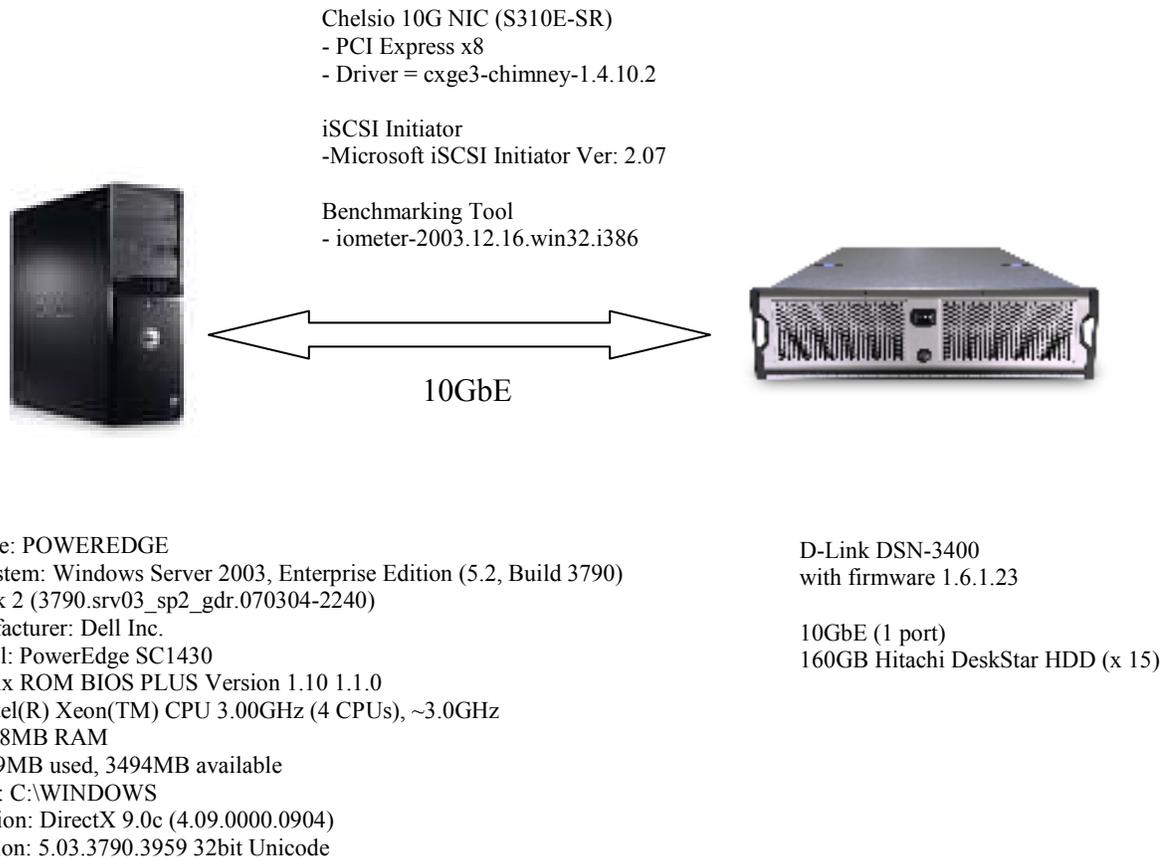


Figure 1

Setting Iometer

- In the Topology pane (see figure 2), select Worker 1. You may have multiple workers in this pane. You will need to repeat these steps for each worker.
- Select the Disk Targets tab. You will see a list of drives. Select the drive letter that matches your xStack storage volume.
- Set the # of Outstanding I/Os to at least 8 (this can be set up to 64). This sets the number of I/Os in the queue (queue depth).

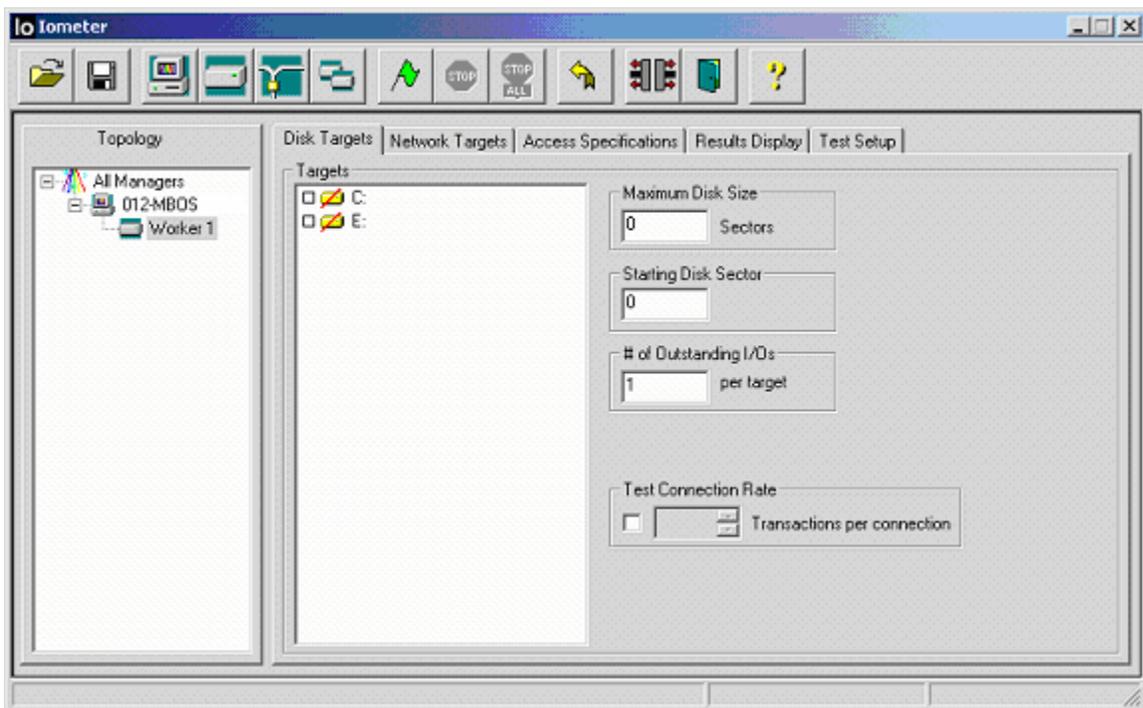


Figure 2

- Select the Access Specifications tab (see figure 3). A list of preset test jobs will appear in the Global Access Specifications pane.
- Select the job labeled “Default.”
- Then select the “Edit” button to the right. The Edit Access Specification window will appear (see figure 4).

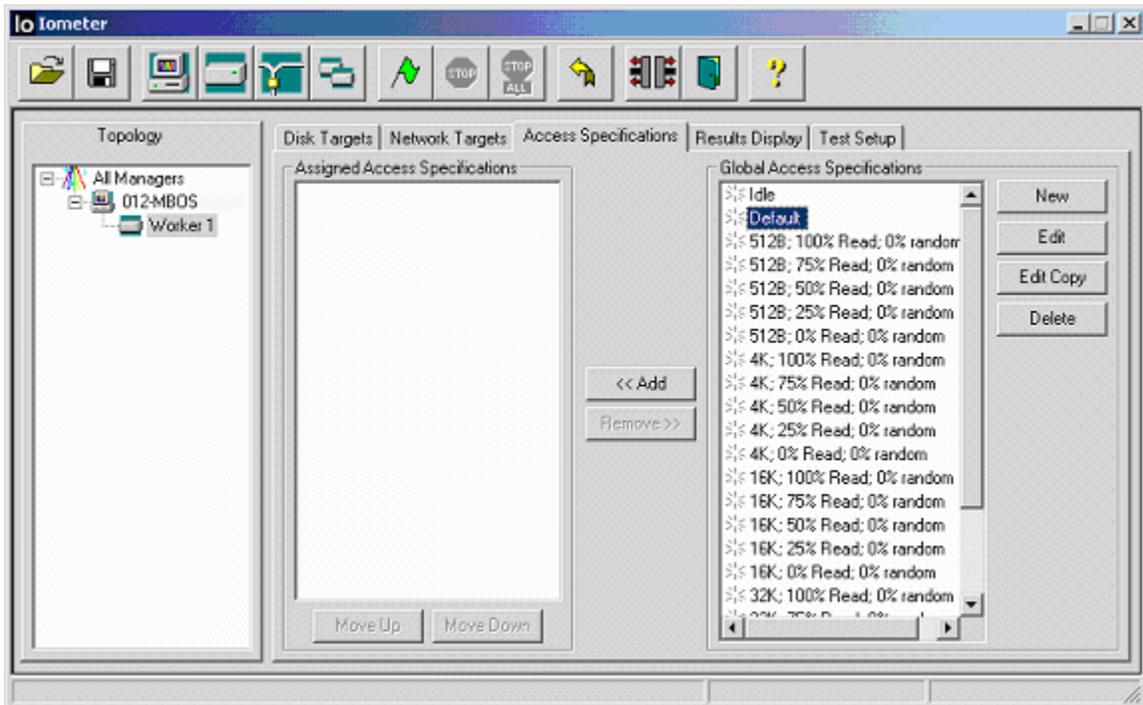


Figure 3

- Set the Transfer Request Size to at least 64K or 256K.
- Set the Percent Random Sequential Distribution slide-bar to 100% Sequential.
- Set the Percent of Access Specification slide-bar to 100 Percent.
- Set the Percent Read/Write Distribution slide-bar to 100% Read.
- Leave everything else at default.
- Click the Okay button at the bottom. You will now go back to the Access Specifications tab in figure 3.

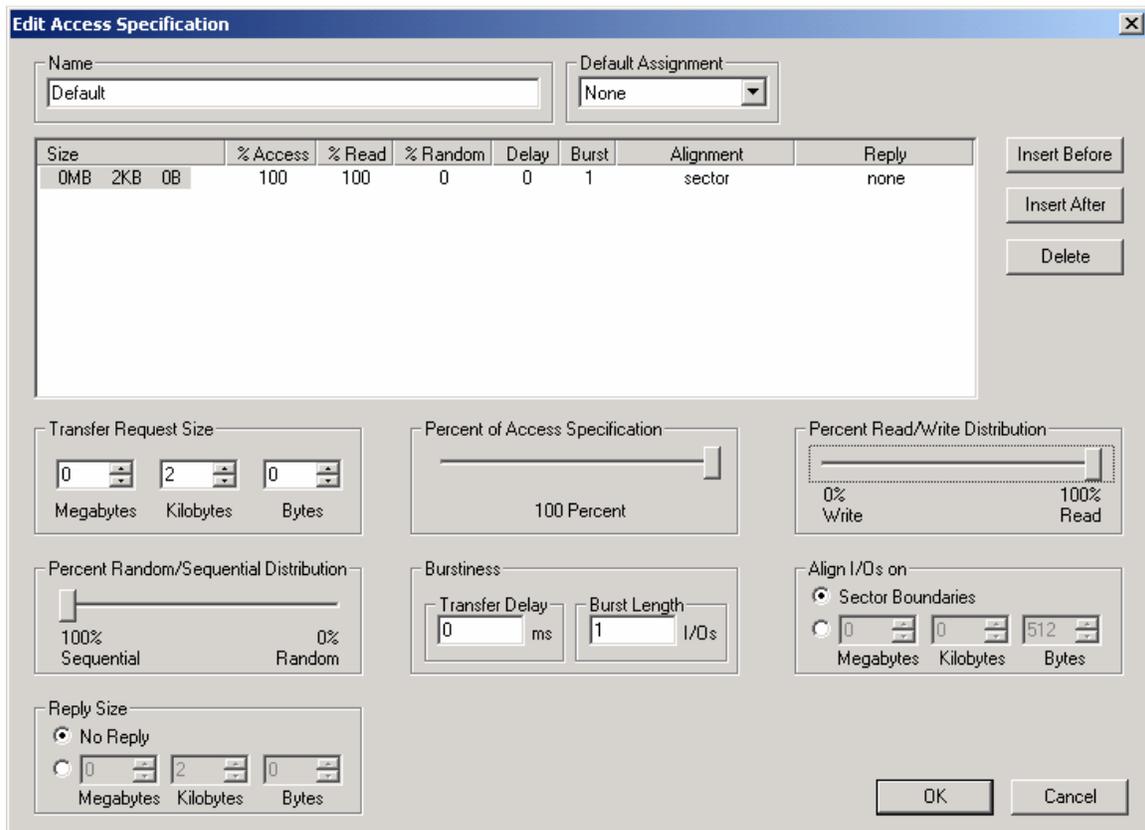


Figure 4

- With the “default” job highlighted in the Global Access Specifications pane (see figure 5), click the “Add” button in the middle of the screen. This will move the “Default” job to the Assigned Access Specifications pane. This will be the job we run.

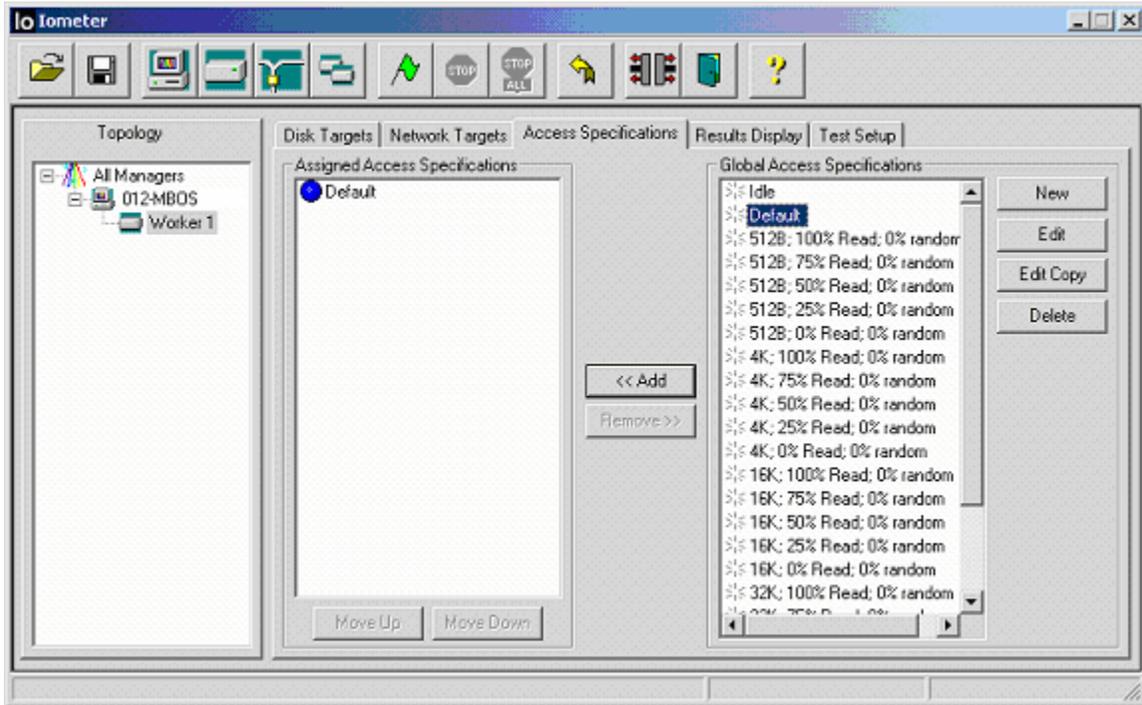


Figure 5

- Reminder: Be sure to repeat the previous steps for each worker appearing in the Topology pane. Then move to the next step.
- Select the Results Display tab as shown in figure 6.
- Move the Update Frequency (seconds) bar to the number “1.”
- Click the green start flag at the top to start the testing. You will see values immediately begin to appear.
- Click the Stop sign to halt the test.

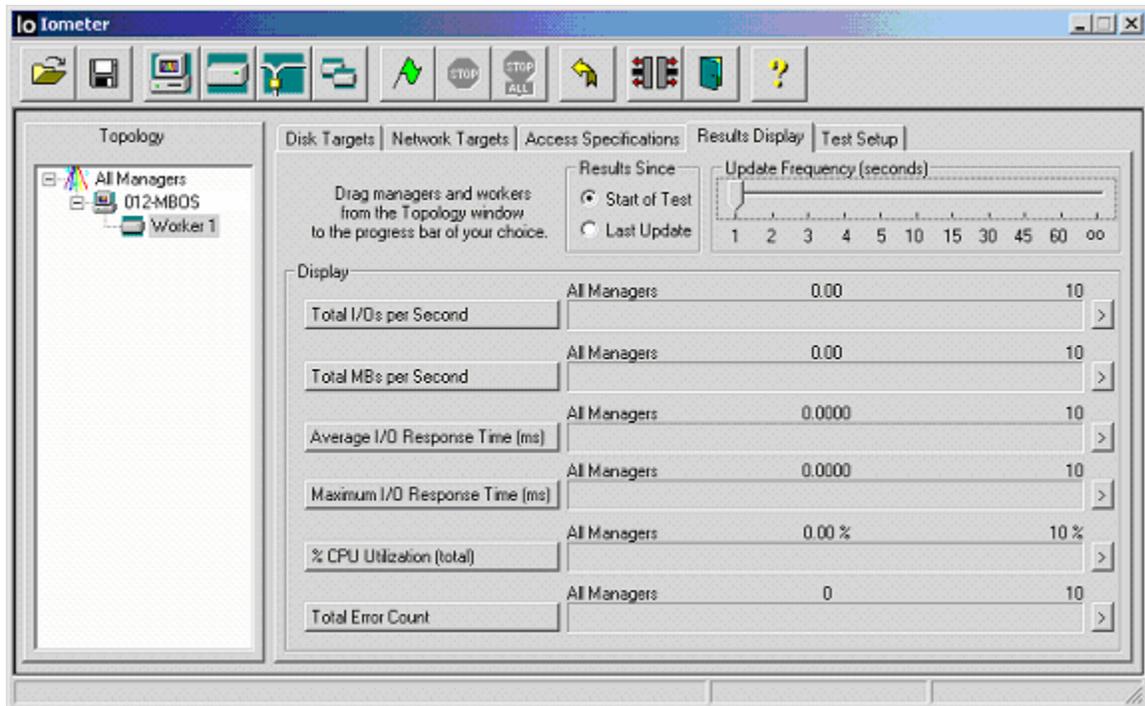


Figure 6

Volume Creation Note

When you create a volume in the xStack Storage GUI and are asked for a chunk size, always select 1MB.

MTU Size Note

If you decide to use an MTU size of 9014 (jumbo frames), then be sure the server/workstation, xStack Storage and the switch itself (if you are using one) are all set to use them.

Volume Initialization Note

Be aware that the xStack Storage unit requires an initialization period for preparing volumes. If you have just created a RAID volume and immediately begin testing, then your throughput will likely suffer. Wait until the volume has finished initializing before running Iometer.

Testing Configuration1

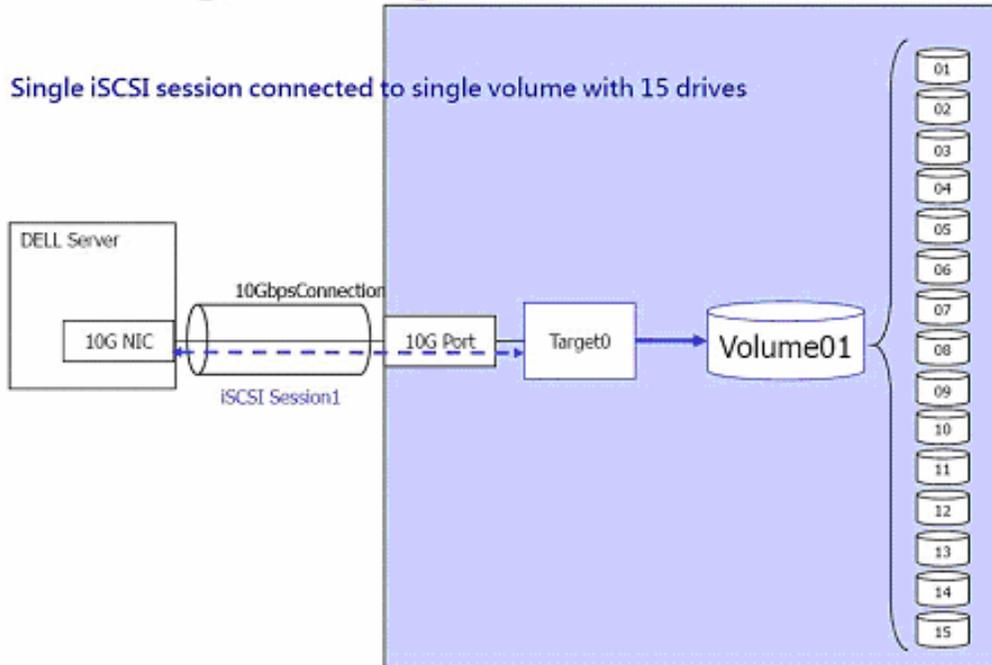


Figure 7

Testing Configuration2

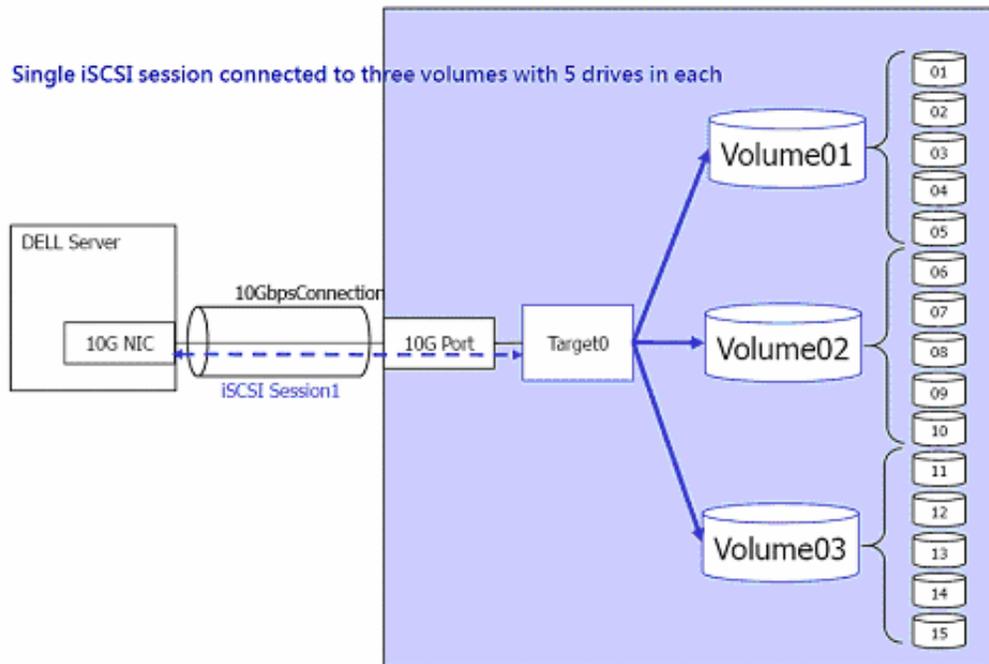


Figure 8

Testing Configuration3

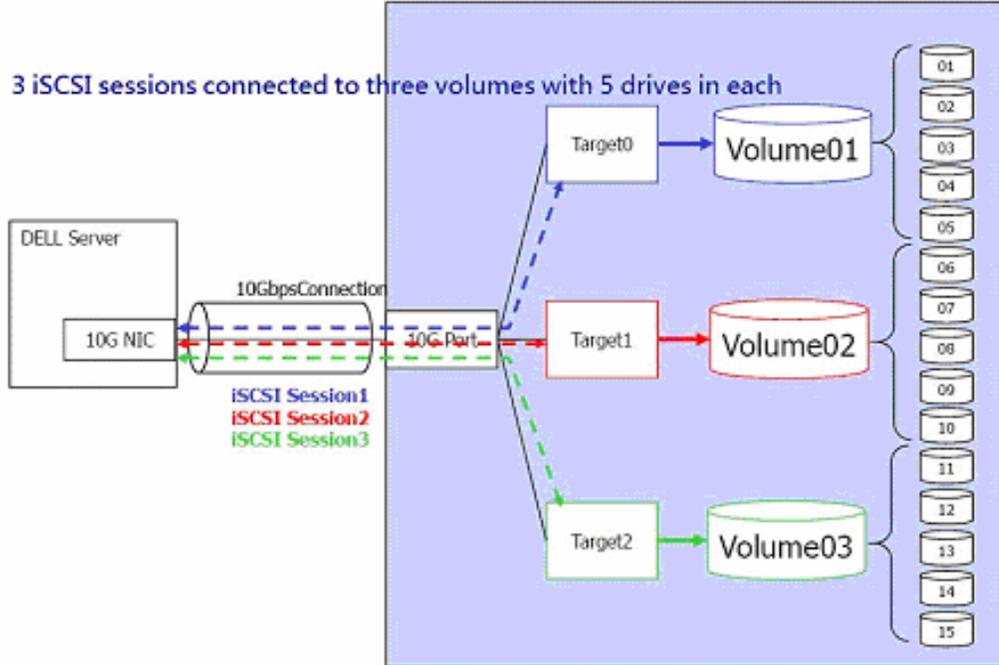


Figure 9

Comparison under different configurations

	Sequential Read (MB/s)	Sequential Write (MB/s)
Testing Configuration 1	224.17	345.91
Testing Configuration 2	N/A	N/A
Testing Configuration 3	710.12	754.80

Increasing the number of iSCSI sessions as shown in Testing Configuration 3 (iSCSI session to a target to a volume) will significantly increase performance.

■Conditions

- Testing Configuration 3
- DELL PowerEdge1340
- 100% Sequential Read, Write
- iometer Transfer size=1MB
- MTU Size=9014 bytes
- Chelsio Driver = cxge3-chimney-1.4.10.2

Comparison under different MTU configurations

Test Scenario 3	Sequential Read (MB/s)	Sequential Write (MB/s)
MTU 1500	656.21	524.15
MTU 9014	710.12	754.80

Setting the MTU size to highest allowed by the DSN SAN array, the higher performance will be. Therefore, we recommend 9014 bytes MTU size be used.

■Conditions

- Testing Configuration 3
- DELL PowerEdge1340
- 100% Sequential Read, Write
- iometer Transfer size=1MB
- MTU Size=9014 bytes
- Chelsio Driver = cxge3-chimney-1.4.10.2

Comparison of TCP/Offload effect

Test Scenario 3	Sequential Read (MB/s)	Sequential Write (MB/s)
TCPo Disabled	478.74	516.35
TCPo Enabled	710.12	754.80

The use of a TCP offload engine improves performance and reduces the host's CPU utilization rate. Therefore, we recommend the use of a TCP offload engine.

■Conditions

- Testing Configuration 3
- DELL PowerEdge1340
- 100% Sequential Read, Write
- iometer Transfer size=1MB
- MTU Size=9014 bytes
- Chelsio Driver = cxge3-chimney-1.4.10.2

Comparison of RAID level

Test Scenario 3	Sequential Read (MB/s)	Sequential Write (MB/s)
RAID 5	574.22	612.74
RAID 0	710.12	754.80

Using a RAID 0 volume will yield better IOPS but at the cost of redundancy. (A failed drive in a RAID 0 “stripe” will cause all of your data to be lost.)

As a note: RAID 10 (1+0) is a “stripe of mirrors” and will yield the highest throughput with the safety of redundancy. However, it requires twice the disk space of a simple RAID 0 volume.

■Conditions

- Testing Configuration 3
- DELL PowerEdge1340
- 100% Sequential Read, Write
- iometer Transfer size=1MB
- MTU Size=9000 bytes
- Chelsio Driver = cxge3-chimney-1.4.10.2