STORAGETEK®

APPLICATION NOTE

Configuring Storage Routers for Maximum Efficiency

DECEMBER 2004

Executive summary

Many IT installation customers have trouble determining the best configuration for connecting SCSI tape drives to SANs via Fibre Channel routers. Their needs can vary, from getting the best capacity from their infrastructure, to getting the best performance. Some customers throw hardware at a performance problem, when instead, by understanding how a device works and making adjustments accordingly, they can improve performance and save money.

There are different advantages with the various possible configurations for connecting SCSI tape drives to SANs. Understanding these differences, along with the capabilities and benefits of the hardware involved, can help you determine what your best configuration will be. By using a configuration designed specifically for your needs, you can save considerable time and money.

One good methodology for configuring an effective routed storage solution entails identifying the bottlenecks in your system architecture and optimizing the other components to deal with this limitation. This document describes the use of Linear Tape Open (LTO) Ultrium tape drives configured with StorageTek's StorageNet® 3300 storage routers and StorageNet 3400 multi-protocol routers to a SAN to illustrate this methodology. However, this same methodology can also be applied to any device.

1 The challenge of SAN configuration

A good understanding of the application environment is important to configuring a SAN. Poorly understanding the environment can lead to either over-configuring the SAN, resulting in too much cost, or under-configuring it, resulting in an inadequate storage configuration. Developing an appreciation of your application environment provides the foundation for creating a SAN appropriate to that environment for an appropriate cost.

When connecting tape drives to the SAN, a significant consideration is what the tape drive's function will be. The majority of time it is for backup solutions, but equally important is the host capability and data volume being backed up. If the hosts getting backed up cannot deliver the data as fast as the router can handle it, you do not have to worry about the aggregate bandwidth of the router. The router will have excess capacity relative to the load being delivered by the application server. But if these hosts and data are of a high-performance nature, evaluate where any possible bottlenecks can be and how to avoid them.

The key to evaluating performance bottlenecks is to model and understand where the potential bottlenecks are, and what impact they will have in your specific deployment. The bottlenecks are driven by the application mix that is being served. Understanding that mix in the context of the storage technology characteristics sets the foundation for more accurately planning your storage configuration and growth needs more accurately.

2 Applying the solution

2.1 Tape drive evaluation using the LTO Gen 2

First, let's look at the IBM LTO Ultrium Generation 2 tape drives. Their performance characteristics are specified as:

- Native sustained data transfer rate of 35 megabytes per second
- Compressed data transfer rate of up to 70 megabytes per second

The "up to" variable is important to understand, even though every tape drive manufacturer uses this potentially confusing terminology. This variable is dependent upon many circumstances, one of which is the compressibility of the data. This can also skew all of the calculations of bandwidth utilization, but the fact is that very few applications are going to push the native rate to the maximum, much less the compressed rate to the maximum. Seventy megabytes per second is the maximum compressed data rate that the drive can sustain regardless of the compressibility of the data. The gating factor becomes the speed of the compression chip. Although the example calculation of compressed data rates is based upon 2:1 compression, in fact, it could be greater than that.

It is possible that an application will burst to these speeds, but it is unlikely two drives will peak at the same time.

Deciding on what specification—native or compressed—you should use depends on the goal of your configuration. If it's to gain maximum speed, use the compressed rate when planning what router to use and the native rate for the tape drive. This will give you maximum bandwidth on the router, and no shortage of tape drives. If cost is a factor and you want to be conservative with equipment expenses, you may want to use native speeds when planning the router and compressed performance on the tape drive. This will result in a smaller router and fewer tape drives. These are the extreme boundaries that delineate a range of options between performance and cost.

A good rule-of-thumb is to base all tape drive speeds on maximum native performance specs. This should leave some room for bursting.

To fine tune the rule-of-thumb, it's always a good idea to ask the backup administrators what transfer rate they are currently experiencing, and then adjust your rates accordingly.

2.2 Router evaluation using the StorageNet 3300 storage routers

Just as with the tape drives, it is important to dissect the SN3300 router specifications:

- One Fibre Channel port (2 Gb at 80% efficiency transfers 160 MB/sec)
- Two SCSI buses, one port per bus; each bus is an Ultra 160 type (160 MB/sec at 90% efficiency transfers 144 MB/sec)

The FC port is rated at 160 MB/sec. This means:

- 160 MB/sec divided by 35 MB/sec (native rate) = 4.57 drives
- 160 MB/sec divided by 70 MB/sec (2:1 compression) = 2.2 drives

One SCSI bus is rated at 144 MB/sec. This means:

- 144 MB/sec divided by 35 MB/sec (native rate) = 4.1 drives per bus or 2 buses times 4.1 drives per bus = 8.2 drives
- 144 MB/sec divided by 70 MB/sec (2:1 compression) = 2.05 drives bus or 2 buses times 2.05 drives per bus = 4.05 drives

Looking at the two components—SCSI bus and FC port—we can determine the number of drives needed using either specification: native speed versus compressed speed.

The component with the lowest number of supported drives at native speed is the FC port itself at 4.57 drives:

$$\frac{160 \text{ MB/sec}}{35 \text{ MB/sec}} = 4.57 \text{ drives per FC port}$$

The component with the lowest number of supported drives at compressed (2:1) speed is also the FC port at 2.2 drives:

$$\frac{160 \text{ MB/sec}}{70 \text{ MB/sec}} = 2.2 \text{ drives per FC port}$$

Depending on your requirements, a safe compromise could be to use two or three drives maximum. Even if two drives were running maximum **native** speed (35 megabytes per second), the third drive could still run at maximum **compressed** speed (70 megabytes per second) with some comfort.

Knowing the application or data type is key to calculating how you should build your configuration and what type of throughput you could expect. If the data is not compressible, more resources in terms of tape drives and router capacity are required to meet the application's backup requirements.

2.3 Router evaluation using the StorageNet 3400 multi-protocol routers

The SN3400 router is very flexible in its configuration; it can handle up to four I/O modules. The two types of module available are:

Fibre Channel module – Two 2-Gbit FC connectors

SCSI module – Four SCSI Ultra 3 buses/connectors

The default configuration is one of each module.

Let's break down the specifications for the SN3400:

- One Fibre Channel module with 2 ports can handle (at 80% efficiency) 160 MB/sec times 2 = 320 MB/sec
- One SCSI module with 4 SCSI buses (Ultra 3) can handle 160 MB/sec

The internal router's aggregate bandwidth is 300 MB/sec, which means:

- 300 MB/sec divided by 35 MB/sec (native) = 8.5 drives per card
- 300 MB/sec divided by 70 MB/sec (2:1 compression) = 4.2 drives per card

The SCSI bus has 4 SCSI buses per card, which means:

- One SCSI bus rated at 160 MB/sec (at 90% efficiency) = 144 MB/sec per bus
- 144 MB/sec times 4 buses = 576 MB/sec per card
- 576 MB/sec divided by 35 MB/sec (native) = 16.45 drives per card
- 576 MB/sec divided by 70 MB/sec (2:1 compression) = 8.22 drives per card

The dual FC port can handle 320 MB/sec, which means:

- 320 MB/sec divided by 35 MB/sec (native) = 9.14 drives per card
- 320 MB/sec divided by 70 MB/sec (2:1 compression) = 4.57 drives per card

Looking at the three components—SCSI bus, FC port and internal processing—we can determine the number of drives needed for both native speed and compressed speed.

The component with the lowest number of supported drives at native speed is the router itself at 8.5 drives.

The component with the lowest number of supported drives at compressed (2:1) speed is the router at 4.2 drives.

These calculations make it clear how the native and compressed throughput can have a big impact on your configuration. Knowing the application or data type is key to calculating how you should build your configuration and what type of throughput you could expect. If the data is not compressible, more resources in terms of tape drives and router capacity are required to meet the application's backup requirements.

Depending on your requirements, a safe compromise could be to use a maximum of five or six drives. Even if four drives are running maximum **native** speed (35 megabytes per second), drives five and six could still run at maximum **compressed** speed (70 megabytes per second) with some comfort.

3 Benefits of optimizing

You can apply this methodology, which breaks down a device or network to determine bottlenecks, to most devices and networks. Whether it's a host platform (CPU, bus, I/O architecture) or an entire network (LAN type, WAN connection, router type), identifying the slowest component will help you ascertain the cost-benefit ratio of improving that component.

Once you know where your potential bottlenecks are, you can properly assess their potential impact in your environment. This will help you determine the scalability of your configuration and when you will reach your limit and need to reconfigure or add additional resources. By having your configuration designed to your specific needs, you avoid over-purchasing new equipment and optimize the equipment you already have.

A baseline understanding of your application environment provides a foundation for anticipating and managing storage configuration changes. For instance, if new applications are planned, you can model and forecast the storage network changes required to support those applications. The ability to accurately forecast these changes allows your enterprise to accurately budget for them and minimize their cost.

Cheater's Corner

As a starting basis without performing the calculations, here are two shortcut rules-of-thumb you can use. The aggressive configuration is appropriate when maximum performance is of primary importance. Use the conservative configuration when you want to avoid unnecessary equipment expenses.

	Aggressive configuration (performance bias)	Conservative configuration (cost bias)
SN3300 router	Four LTO Gen 2 drives	Two LTO Gen 2 drives
SN3400 router	Eight LTO Gen 2 drives	Four LTO Gen 2 drives

STORAGETEK®

ABOUT STORAGETEK®

Storage Technology Corporation (NYSE: STK) is a $2 billion global company that enables businesses, through its information lifecycle management strategy, to align the cost of storage with the value of information. The company's innovative storage solutions manage the complexity and growth of information, lower costs, improve efficiency and protect investments. For more information, visit www.storagetek.com, or call 1.800.275.4785 or 01.303.673.2800.

WORLD HEADQUARTERS

Storage Technology Corporation
One StorageTek Drive
Louisville, Colorado 80028 USA
1.800.877.9220 or 01.303.673.5151

NT 0001 A 12/04