Agilent Technologies

Deploying a SAN Extension Network

Technology and Test Considerations

Application Note 1570

Agilent Technologies

Introduction

New software applications, e-business technologies and government regulations for data retention have significantly increased the rate at which organizations create and store information. As such, data now represents the livelihood of many enterprises, government and educational organizations. To ensure the safety and availability of the data against a range of natural and human events, the data needs to be distributed geographically using a robust and redundant storage network.

The first part of this article summarizes distance extension technologies that are available – Wavelength Division Multiplexing (WDM), SONET/SDH, and IP networks. Fibre Channel is the dominant protocol in the Storage Area Network (SAN) industry and most SAN extension devices provide solutions to carry Fibre Channel traffic over the MAN or WAN, thus, this article will focus on the extension of Fibre Channel fabrics.

The technology summary is followed by a review of the challenges in testing and validating a storage extension network. Different test methodologies are proposed to address these challenges, and the Agilent Technologies SAN Tester is used as an example on how to test and qualify the extension network infrastructure.

Technology Summary

WDM

WDM works by multiplexing multiple signals onto the same fiber using different wavelengths. There are two types of WDM technologies – Coarse WDM (CWDM) and Dense WDM (DWDM). The main difference between the two is how many wavelengths can be multiplexed onto the same link – thus the bandwidth of the link available.

One of the key advantages of a WDM solution compared to other extension technologies is the very high bandwidth available. CWDM links can carry up to 16 Gbps of data, while DWDM solutions can carry 320 Gbps of data. (The numbers quoted for CWDM and DWDM bandwidths are approximations. Exact numbers depend on the system configuration). High bandwidths are very important for enterprises that require high speed and low latency connections for synchronous storage applications, such as synchronous replication.

WDM solutions are also very scalable. Enterprises can start with a single wavelength, as data needs grow, the enterprise can add extra bandwidth by adding more wavelengths without changing the physical network. This solution can also be used for service providers looking to provide storage private line services.

WDM systems also have the ability to transport a variety of services, since it is protocol agnostic. So a WDM infrastructure already deployed for IP network traffic can also be use for Fibre Channel storage traffic. In fact, the same link can be used to carry IP/TCP traffic, ESCON/FICON, as well as Fibre Channel.

SONET/SDH

SONET/SDH based networks are widely deployed in the core and metro transport networks. They are well known and trusted to be reliable and cost-effective.

With the introduction of Generic Framing Procedure (GFP) ITU standard G.7041, there is now an even more efficient way of transporting Fibre Channel data traffic through the SONET/SDH infrastructure. GFP allows the mapping of the Fibre Channel frames with very little overhead into SONET/SDH and has minimal impact on data layer latency and throughput.

GFP can also be deployed with Virtual Concatenation (VCAT). VCAT is intended to support the transport of payloads that do not efficiently fit into traditional SONET/SDH payloads. VCAT along with the associated Link Capacity Adjustment Scheme (LCAS) ITU-T standard G.7042 permits flexible extension bandwidth allocation ranging from 50 Mbps to the full rate of Fibre Channel. This means enterprises can match the leased bandwidth to the requirements of the storage application.

IP

Two popular solutions for extending the Fibre Channel over the IP network are FCIP and iFCP.

FCIP

FCIP is currently the most widely supported IP based extension protocol. This is probably due to the fact that it is simple and easy to implement. The basic concept of FCIP is a tunnel that connects two Fibre Channel SAN islands through a IP network. Once connected, the two SAN islands logically merge into a single fabric across the IP tunnel.

An FCIP gateway is required to encapsulate Fibre Channel frames into TCP/IP packets. This is sent through the IP network. On the remote side, another FCIP gateway receives the incoming FCIP traffic, strips off the TCP/IP header before forwarding the native Fibre Channel frames into the SAN (Figure 1). The FCIP gateway can be a separate device, or its functionality can be integrated into the Fibre Channel switch.

The obvious advantage of using FCIP is that existing IP infrastructure can be used to provide the distance extension.

iFCP

The iFCP protocol also provides connectivity using the IP network using a gateway device. To native Fibre Channel devices connected to an iFCP gateway, the gateway mimics a fabric port and provides the associated fabric services. This means existing native Fibre Channel devices can be used with iFCP seamlessly.

From the IP side, each of the Fibre Channel devices connected to the iFCP gateway is given a unique IP address, which is advertised in the IP network. This allows individual Fibre Channel devices to be reached through the IP network via the iFCP gateway. The ability to individually address devices gives iFCP some advantages compared to the FCIP protocol.

The biggest advantage is that of stability. Using FCIP between two Fibre Channel SAN islands will cause the islands to merge into one. This means if there are perturbations in the IP network, it can potentially cause the fabric to rebuild on both sides of the IP tunnel. Using iFCP, the connectivity is between individual devices, and the fabrics stay separate. If perturbations occur in the network, it may affect individual connections, but will not cause fabric rebuilds, thus leading to more stable fabrics on both sides of the IP network.

The disadvantage compared to FCIP is the limited availability of iFCP solutions in the market place. This could be because FCIP is very simple to implement, thus the FCIP solution is widely available and is provided by a number of different manufacturers. In contrast, iFCP is only supported by a limited number of vendors.

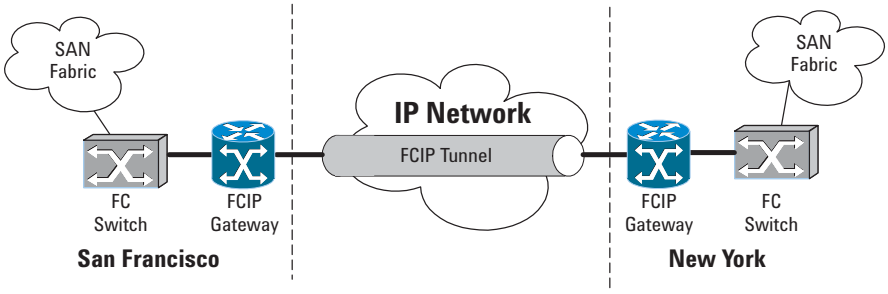

Figure 1. FCIP tunneling between two separate FC SAN islands.

Test Considerations

Selection of the technology to use for the extension network, WDM, SONET/SDH, or IP depends on the storage application and its requirements. For example, small or medium sized businesses' (SMB) remote archive requirements would be very modest in terms of both bandwidth and latency. In that case, it may be sufficient for the SMB to use a leased T3 link to carry the backup data to a remote site using FCIP. The selection of the type of extension network depends on many factors and is out of the scope of this paper. However, regardless of the technology used, the same types of tests must be performed to qualify the network and the devices used. The tests must be able to validate Service Level Agreement (SLA) guarantees, data accuracy and any other requirements of the storage protocols or storage application. These tests may be done by the data center staff, or may be done by service providers who are implementing storage extension to provide a storage private line to enterprise customers.

Network deployment tests

When deploying networks to support storage extension, the deployed networks must be tested to validate the basic performance characteristics that are required by the application. The characterization needs to take into account the network behavior under normal conditions as well as failure or stress conditions.

Network characterization

The application to be deployed using the extension link plays the critical role in determining what kind of network performance is required.

For example, synchronous replication is one type of application that may use an extension link. Synchronous replication is typically used for mission critical data. The basic idea of synchronous replication is to keep a remote disk array completely synchronized with the local disk. Each time a write operation occurs, a write is done on the local disk, then the same write operation occurs on the remote disk. An acknowledgement is sent to the application only when the local and remote write are both successful. This ensures the data on the local copy and the remote copy is kept exactly the same. In the event of a failure in the local disk, the servers can be switched to the remote disk array seamlessly. However, the need to wait for the remote write makes the synchronous replication application very susceptible to delays in the transport layer. Some examples of synchronous replication applications are EMC Symmetrix Remote Data Facility (SRDF), IBM Peer-to-Peer Remote Copy (PPRC), HDS TrueCopy.

To ensure that the extension link can support the synchronous replication application, the storage architect must verify that the delay characteristics meet the requirements of the application, especially under heavy load.

To test that the network satisfies the requirements of the application, the test platform selected must be able to generate I/O and also measure the delay through the network. The testing is usually done in two phases. The first step is to do get an understanding of the network behavior under ideal/non-congested conditions.

The test platform is used to generate I/O through the network, and the network delay behavior is measured to ensure it satisfies the need of the application.

However, there are times when the network might be highly loaded. The second step of the test is to create congestion in the network, and verify network behavior under congested conditions. The test platform in this case is used to generate I/O to fully load the network and see how the application traffic is affected.

The Agilent SAN Tester can be a very helpful tool in understanding the network behavior under both ideal conditions as well as congestion conditions. For example, the I/O generated from the tester can be changed in real time. Thus, by slowly increasing the load and monitoring the real-time statistics, such as delay, a user can clearly identify the traffic level above which the network can no longer effectively support the application (Figure 2).

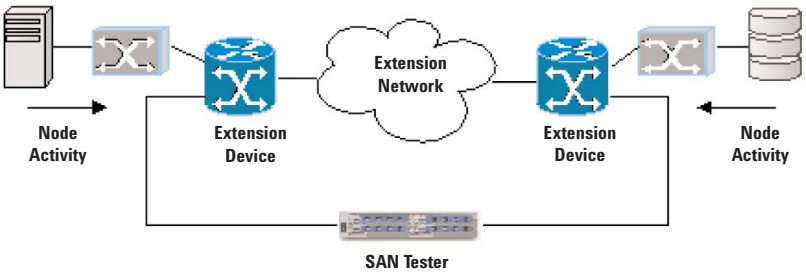

Figure 2. SAN Tester creating stress load on extension system.

Failure recovery

Regardless of the type of network used for extension, failures can always occur. If FCIP is used, this could be in the form of one of the IP routers being overloaded and causing IP encapsulated FC frames to be dropped. In the case of SONET/SDH or DWDM, one of the links on the ring could be broken. These failures in the extension link can be the cause of longer latency delays or application timeouts. Thus, understanding implications for the storage application when failure occurs is another test that should be performed. Typical measurements for the failure are the failure recovery time (how long it takes for the network to recover), and how many frames are lost.

Each of the three networks (WDM, SONET, and IP) handle failures differently, and the expected network "down time" may be very different as well. For example, SONET/SDH has an Automatic Protection System (APS) that guarantees recovery within 50 ms, whereas IP is a best effort service and by itself does not provide any guarantees.

Failure recovery time is typically difficult to measure, and may require some specialized equipment.

The Agilent SAN Tester provides a quick and easy way for the user to measure the failover characteristics of the extension link, including failure recovery time, and how many frames are lost.

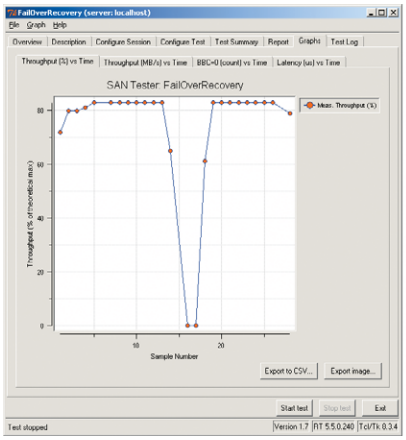

Figure 3 - Failover recovery test

Service Level Agreement testing

Today's trends show that Service Level Agreements (SLAs) are a mandatory part of the relationship between storage private line providers, IT departments and their end users. Thus for both service providers and enterprise IT group, testing needs to be done to ensure that SLAs can be guaranteed.

It is possible to guarantee SLAs by significantly over-specifying the network. For example, if the guaranteed latency is 20 µs, then a network designed with a 5 µs latency will satisfy the SLA. If the bandwidth guarantee is 50 MB/s, then a dedicated link of 100 MB/s will satisfy the guarantee. However, this is a more expensive approach, as most of the resources will be idle for a majority of the time.

The other way to ensure that SLAs are satisfied is by using a combination of Quality of Service (QoS) policies and robust testing. By implementing the QoS policy and then testing under the different stress conditions, it gives confidence in the network's ability to satisfy the SLAs.

There are many types of SLAs. Some SLAs specify a fixed bandwidth. For example, if a company pays for a fixed amount of bandwidth, the storage service provider should make sure that the bandwidth is enforced. When this company tries to send more data than the allocated bandwidth, it is throttled back, so the data from this company doesn't cause congestion and adversely affect other data carried on the same extension link. Another type of SLA guarantees maximum network latency. This type of SLA can be enforced on all the traffic from one source (company), or can be implemented only on the traffic from one specific application.

Companies may have multiple SLAs implemented to addresses the needs of different applications. Therefore the test platform must be able to simultaneously validate all the different types of SLAs.

The Agilent SAN Tester platform can be used to test the different types of SLAs. By providing more I/O than the allocated bandwidth and monitoring the real-time statistics, it is possible to validate the QoS policy enforcement of throttling the traffic.

To test the policy enforcement for latency, the Agilent SAN Tester can be used to generate two streams of I/O marked with different QoS settings. By increasing both types of I/O to overload the link, the QoS policy should be enforced, to guarantee the required bandwidth and latency for the higher priority I/O, and throttling the lower priority I/O stream.

The Agilent SAN Tester shows the SLA enforcement clearly in a graphical format in the measurement view. Figure 4 shows an example of the results from the measurement system. As the lower priority traffic is slowly increased, the network throttles it to 50 MB/s to protect the required bandwidth of the higher priority traffic.

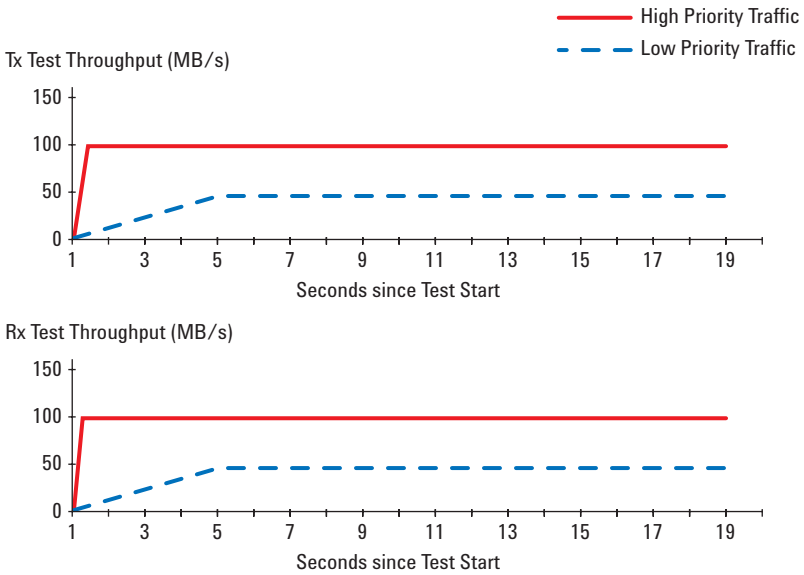

Figure 4. Using San Tester Real Time Measurement to validate SLA.

Summary

Distance extension is a key strategy used by organizations to protect the data that is central to their business and survival. There are many reliable technologies that can be used to implement distance extension. What has been shown in this paper is that testing is a key consideration in addition to the choice of the underlying technology. The network must be tested to ensure that it does support the target storage application under both normal and adverse network conditions and to ensure that the extension network is indeed robust and reliable.

Simulating different network conditions is one of the key network test challenges. Normal network behavior can be easily tested, however specialized test platforms may be required to reproduce stress conditions and error scenarios.

Agilent SAN Tester is a platform that helps address these test challenges. It offers extended device simulation, error injection and stress capabilities to simulate real network scenarios. Also, with automated test scripts, the testing can be simplified and test time significantly reduced. Testing with the Agilent SAN Tester helps ensure that critical data is well protected through a robust and reliable geographically distributed network.

Agilent Technologies' Test and Measurement Support, Services, and Assistance
Agilent Technologies aims to maximize the value you receive, while minimizing your risk and problems. We strive to ensure that you get the test and measurement capabilities you paid for and obtain the support you need. Our extensive support resources and services can help you choose the right Agilent products for your applications and apply them successfully. Every instrument and system we sell has a global warranty. Two concepts underlie Agilent's overall support policy: "Our Promise" and "Your Advantage."

Our Promise
Our Promise means your Agilent test and measurement equipment will meet its advertised performance and functionality. When you are choosing new equipment, we will help you with product information, including realistic performance specifications and practical recommendations from experienced test engineers. When you receive your new Agilent equipment, we can help verify that it works properly and help with initial product operation.

Your Advantage
Your Advantage means that Agilent offers a wide range of additional expert test and measurement services, which you can purchase according to your unique technical and business needs. Solve problems efficiently and gain a competitive edge by contracting with us for calibration, extra-cost upgrades, out-of-warranty repairs, and on-site education and training, as well as design, system integration, project management, and other professional engineering services. Experienced Agilent engineers and technicians worldwide can help you maximize your productivity, optimize the return on investment of your Agilent instruments and systems, and obtain dependable measurement accuracy for the life of those products.

Agilent Email Updates

www.agilent.com/find/emailupdates
Get the latest information on the products and applications you select.

Agilent Direct

www.agilent.com/find/agilentdirect
Quickly choose and use your test equipment solutions with confidence.

For more information on Agilent Technologies' products, applications or services, please contact your local Agilent office. The complete list is available at:

www.agilent.com/find/contactus

Phone or Fax
United States:
(tel) 800 829 4444
(fax) 800 829 4433

Canada:
(tel) 877 894 4414
(fax) 800 746 4866

China:
(tel) 800 810 0189
(fax) 800 820 2816

Europe:
(tel) 31 20 547 2111

Japan:
(tel) (81) 426 56 7832
(fax) (81) 426 56 7840

Korea:
(tel) (080) 769 0800
(fax) (080) 769 0900

Latin America:
(tel) (305) 269 7500

Taiwan:
(tel) 0800 047 866
(fax) 0800 286 331

Other Asia Pacific Countries:
(tel) (65) 6375 8100
(fax) (65) 6755 0042
Email: tm_ap@agilent.com
Contacts revised: 05/27/05

Printed in USA, January 19, 2006
5989-4653EN

www.agilent.com/find/santester

Agilent Technologies