White Paper

Data Storage across the MAN by deploying a SAN over Dark Fiber

by Pieter Panis, EuroFiber

1 Introduction

The purpose of this document is to illustrate how a storage area network (SAN) can be designed in a cost-effective way. This document illustrates a case where an enterprise wants to replicate his data in a remote site. Three examples are shown of how this could be resolved. Also other configurations are possible.

2 The Problem

Data Storage at a remote site.

The internet and related activities continue to expand. Enterprises data volumes are growing exponentially. Many enterprises have already deployed effective ways to store and maintain their data. One way of doing this is through a Storage Area Network (SAN) type of network. The idea is to build a dedicated network that connect servers and storage devices and transport storage traffic without burdening the enterprise LAN. Several factors have made SAN an attractive solution: performance, reliability, availability, scalability and ease of management.

Now more and more enterprises are also looking for solutions to replicate or store their data in remote sites. This can have several reasons:

1. Disaster Recovery: maintaining an up-to-the-moment copy of critical data at an alternate site.
2. Content Distribution: data sharing and consolidation, giving multiple servers dynamic update capabilities to a shared pool of data, and copying data from distributed sites to a central site.
3. Cost Savings by making use of shared infrastructure.

This could then result in a SAN, extended over 2 or more sites. The remote site(s) need to be connected to the “primary” enterprise site, in some way or another. The capacity of the connection required could be quite high, especially when there is a lot of data involved or when the primary and remote storage systems need to be synchronized in real time.

These capacity needs could potentially be covered through the leasing of a traditional circuit. However these circuits are not easy to upgrade, lack granularity, and can be quite expensive, especially for higher speeds, and over longer distances. Customers could have the following requirements:

High Bandwidth:

The customer might want to replicate his data in a synchronous way.

Ease/Freedom to Upgrade:

If the customer wants to have if more bandwidth available, he usually does not want to wait several months. He wants to do this whenever he wants.

Transparency, technology control, Security for mission critical applications.

The customer potentially wants to fully determine, understand and control his network design. He wants to understand all the strengths and shortcomings of the implemented technology: redundancy, single points of failure, protection mechanisms, latency, etc.

Security against third parties:

The customer potentially wants to fully control his traffic: he wants nobody to monitor or interrupt his traffic without his permission.

And potentially without extra cost, extra work in terms of network management, or involving extra complexity.

3 SAN Network Solutions

Three solutions are described how you could build a network to replicate data in different buildings:

1. A SAN implemented directly on fiber.
2. A SAN implemented on DWDM over fiber.
3. A SAN implemented on GIGE over fiber.

3.1 Solution 1

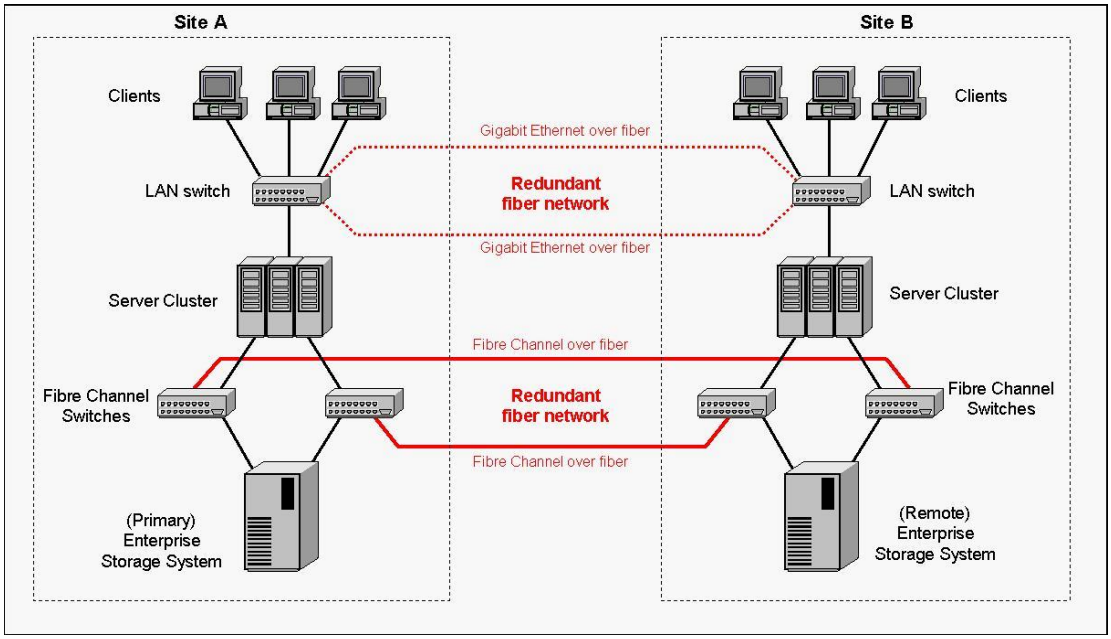

SAN over fiber

Network Architecture:

Typically the servers have Ethernet and Fibre Channel connections to communicate with both the Ethernet LAN and the Fibre Channel SAN. The Fibre Channel switches can be interconnected directly with dark fiber, implemented on a metropolitan network. The distance would however be limited to 10 km.

In the case above two fibre channel switches are used for redundancy purposes. To make the solution completely redundant, the two fiber connections can be implemented redundantly, making a fiber ring.

Also the LAN switches can be interconnected with an additional dark fiber connection, following the same fiber route as the fibers used for transporting the Fibre Channel protocol. The interconnection can simply be done by equipping the LAN switch with the proper interface cards, or GBICs.

How does it work?

The Fibre Channel switches are equipped with GBICs. Short wave lasers can go up to 500 meters, Long Wave lasers up to 10 km.

Advantages:

- Cost effective equipment solution: no major extra equipment costs, only investment in the right interface card needed.
- High bandwidth available per fiber pair: 1 Gbps or even more.

Disadvantages:

- One fiber pair needed per link.
- Distance is limited to 10km.
- Protection switching, if necessary, has to be performed by LAN or SAN equipment itself.

3.2 Solution 2

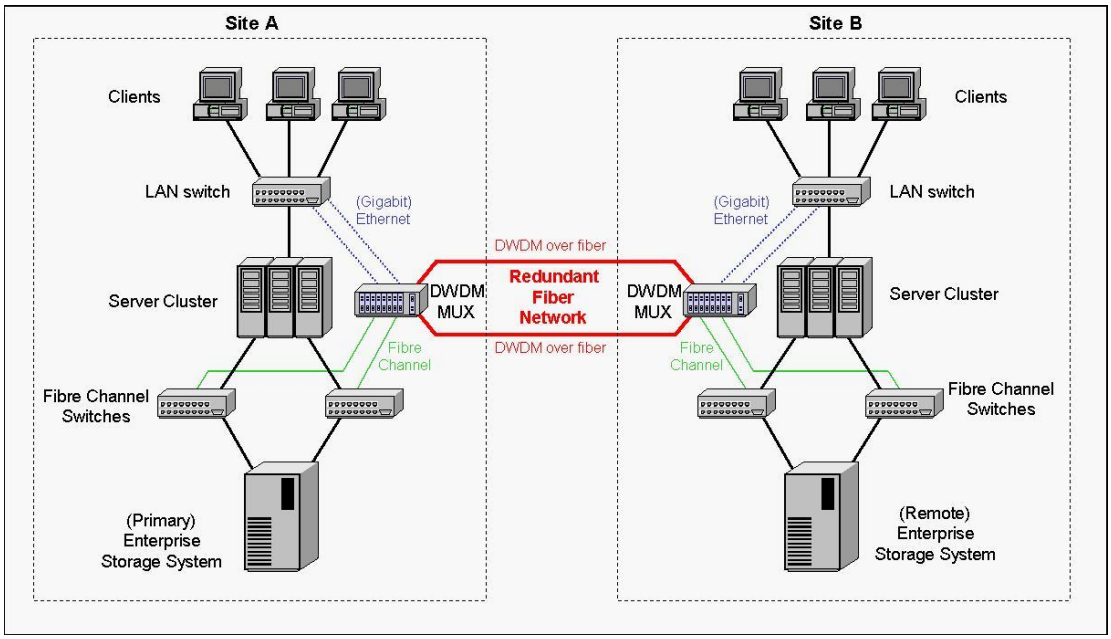

SAN over DWDM (over fiber)

Network Architecture:

Typically the servers have Ethernet and Fibre Channel connections to communicate with both the Ethernet LAN and the Fibre Channel SAN.

In the case above the Fibre Channel switches and the LAN switches need to be interconnected using a metropolitan fiber connection, or metropolitan fiber ring in case of redundancy.

This is done via a DWDM transmission layer, connecting the LAN switches and Fibre Channel switches to a DWDM Multiplexer, which transmits al of these signals across a metropolitan network to the other site.

How does it work?

A DWDM unit can support different types of modules/protocols: typical SDH protocols such as STM-1 to STM-16, but also Fibre Channel and Ethernet protocols, such as Fast Ethernet and Gigabit Ethernet. The Fibre Channel switches are patched to the respective modules in the DWDM chassis supporting Fibre Channel protocol. The LAN switches are patched to the respective modules in the DWDM chassis supporting Fast Ethernet or Gigabit Ethernet.

The DWDM Multiplexer unit maps different protocols/services on different wavelengths. These wavelengths are then multiplexed in single light bundle, which can be transported on a single fiber pair, over long distances.

Advantages:

- More than one service over a single fiber pair, mapping different services on different wavelengths.

- You can easily connect additional services, just by adding another module in the DWDM chassis.
- Virtually unlimited capacity available: there is enterprise equipment available from multiple vendors who can multiplex 32 wavelengths of up to 2.5 Gbps per wavelength. But there is also more expensive DWDM equipment available multiplexing even more wavelengths, and supporting 10 Gbps per wavelength.
- Investment in the DWDM Equipment can be done gradually: per chassis, per module.
- Protection mechanisms can potentially be implemented.
- Longer distance scan be spanned, up to 60km or more.

Disadvantages:

- DWDM equipment investment cost: chassis and initial modules.

Remark:

- The maximum allowable distance between the sites is determined by the delay that can be afforded between the two sites. This can be critical in cases where real time or synchronous replication is needed.

3.3 Solution 3

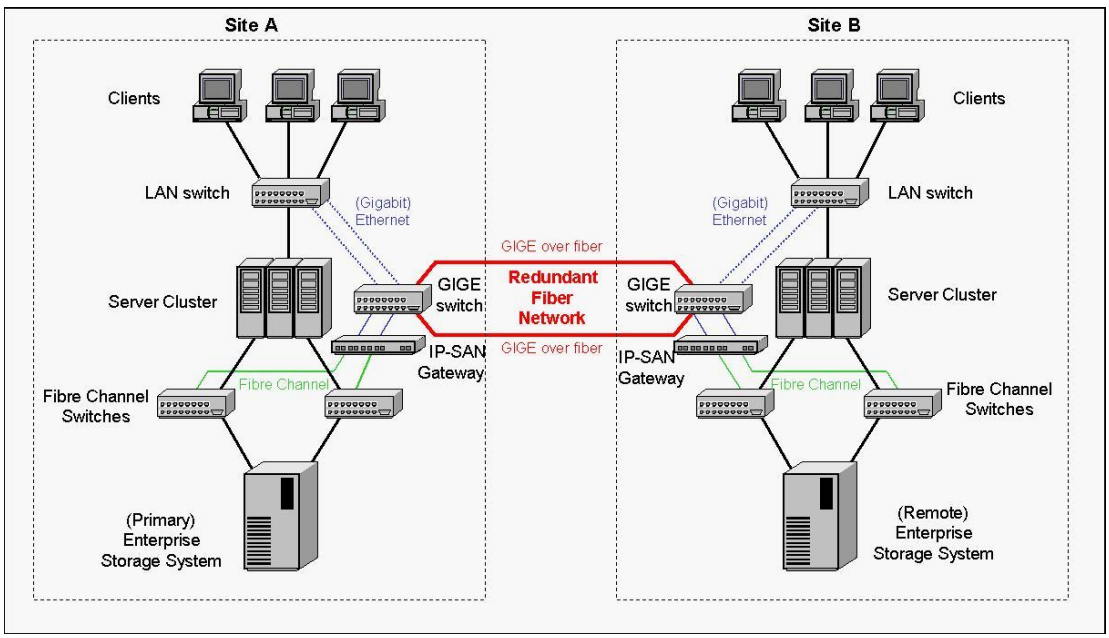

SAN over IP over Gigabit Ethernet (over fiber)

Network Architecture:

Typically the servers have Ethernet and Fibre Channel connections to communicate with both the Ethernet LAN and the Fibre Channel SAN.

In the case above the Fibre Channel switches and the LAN switches need to be interconnected using a metropolitan fiber connection, or metropolitan fiber ring in case of redundancy.

This is done via a Gigabit Ethernet (GIGE) transmission layer, connecting the LAN switches and Fibre Channel switches to a GIGE switch, which transmits the signals across a metropolitan network to the other site, using the Ethernet protocol.

How does it work?

A GIGE switch is designed to transport Ethernet and IP packets. Interconnecting LAN switches using GIGE is obvious. To transport Fibre Channel over an Ethernet network

however is not that obvious: you need a gateway in between to map the Fibre Channel packets in to Ethernet frames.
This gateway is than connected straight to the GIGE switch.

The Gigabit Ethernet network can transport the signals over long distances across the metropolitan network.

Advantages:

- Shared network for storage and LAN: more than one service over a single fiber pair, mapping different services in different Ethernet frames.
- GIGE can be upgraded to 10Gigabit Ethernet, when necessary and when available.
- Protection mechanisms can potentially be implemented.
- Longer distance scan be spanned, up to 60km or more.
- Equipment investment not as high as DWDM investment.

Disadvantages:

- All services need to be mapped on Ethernet frames before being transported over a GIGE network.
- Fiber capacity is limited to the transmission technology implemented on it.

Remark:

- The maximum allowable distance between the sites is determined by the delay that can be afforded between the two sites. This can be critical in cases where real time or synchronous replication is needed.

4 Conclusion

The customer has now the ability to replicate his data across a metropolitan area. A solution based on dark fiber has the following advantages:

- Virtually unlimited bandwidth, only limited by the transmission layer (DWDM or GIGE) or by the interfaces on the LAN or SAN switches.

- Freedom to upgrade his bandwidth, for instance by plugging in additional modules in the DWDM chassis. No upgrade of the backbone, i.e. the fiber infrastructure, is needed. He can do these upgrades whenever needed.

- A transparent solution with redundancy, if required.

- Security against third parties: the only way for a third party to interrupt or monitor the traffic would be by cutting the fiber, and put equipment in between: this cannot easily be done without being noticed both by the customer or EuroFiber.

- No extra, or very limited, extra network management needed.

He can build such a solution by directly plugging the fibers in to his Fibre Channel switches but also in his LAN switches. However if he wants to make more efficient use of his fiber, or if he wants to span longer distances, he might have to implement a transmission layer between the fiber and the storage equipment. This can for instance be DWDM or Gigabit Ethernet. Factors that have to considered before making a choice in transmission layer are: does the replication have to be asynchronous or synchronous (or how much delay can I afford?), can I afford the service being down for a short while (what level of redundancy do I need?), how many services do I want to multiplex on one fiber pair, etc.

Headquarter Brussels

EuroFiber
Boulevard de la Woluwe 2
1150 Brussels
Belgium
Tel.: +32 (0) 2 307 1200
Fax: +32 (0) 2 307 1290

The Netherlands

EuroFiber Nederland
Paasheuvelweg 20
1105 BJ Amsterdam
The Netherlands
Tel.: +31 (0) 20 562 6000
Fax: +31 (0) 20 562 6001

EuroFiber is a pioneer and leader in metropolitan fiber connectivity solutions in metropolitan areas in Belgium and the Netherlands. Our core business is the provisioning of dark fiber networks in those areas.

Visit **www.eurofiber.com** to learn more about our company, network, products, and more.

Notice: This document is for information purposes only, and does not imply any guarantee from EuroFiber concerning any configuration described. Contact our Headquarters in Brussels for more information on products and preferred vendors to build a total solution.