

SAN Routing



INTRODUCTION

The maturity of the Fibre Channel SAN market is in part affirmed by the growing requirement of enterprises to connect their existing SAN islands together. Whether for disaster recovery, campus storage connectivity, consolidating remote tape backup or content sharing over regional distances, linking SANs together is the next evolutionary step once departmental and application-specific storage networks have been deployed. Unlike connecting Ethernet switches or routers, however, connecting Fibre Channel fabrics is not as simple as cabling ports together. The fabric services that the Fibre Channel architecture provides require special attention to ensure interoperability and stability of the storage network as more Fibre Channel switches are combined into a single network. This paper examines the requirements and consequences of Fibre Channel fabric extension and reviews an alternate and more robust approach to SAN connectivity: SAN Routing.



FIBRE CHANNEL FABRIC BUILDING

By design, Fibre Channel fabrics are self-configuring. When a Fibre Channel end device such as a storage array is attached to a Fibre Channel switch, it must log on to the switch to receive a unique 24-bit network address. This network address is used for routing data through the fabric from one device to another. As shown in the table below, the three-byte address is composed of Domain, Area and Port fields. For the purposes of fabric building, the most significant byte (Byte 0) typically identifies a single fabric switch. After subtracting reserved addresses, there are 239 possible Domain identifiers for individual fabric switches, which is the maximum number of switches that could theoretically be joined in a single Fibre Channel network.

BYTE 0	BYTE 1	BYTE 2
Domain	Area	Port
01-EF	00-EF	00-FF
239	256	256

With a unique Domain identifier, a fabric switch can allocate the remaining two bytes (Area and Port) to provide approximately 64k addresses (256 x 256) for individual devices. Because the fabric switch is responsible for issuing unique network addresses to end devices, the end devices do not require manual address administration. In practice, however, SAN administrators are often forced to verify or assign preferred addresses due to specific requirements of particular host bus adapters or storage arrays.

If each Fibre Channel switch is responsible for automatically assigning unique addresses to its end devices, there must also be some mechanism to ensure that each fabric switch has a unique Domain identifier. By standard, this is accomplished through a process known as principal switch selection. This process determines which switch will act as master in the newly configured SAN and oversee the allocation of unique blocks of addresses to each switch. Without a principal switch selection process, two switches could inadvertently maintain the same Domain identity, which would result in duplicate addressing in the network and misrouting of data.

The principal switch selection is initiated by flooding special frames on all expansion ports (E_Ports) of all attached Fibre Channel switches. Flooding across all expansion ports continues until a pre-established fabric stability timeout value (F_S_TOV) is reached. This value must be high enough to ensure that all attached fabric switches have been duly notified that a principal switch selection process is required. The more switches in the network, the longer this process will be.

During the principal switch selection procedure, one of two types of fabric building will occur. A non-disruptive Build Fabric event will try to preserve existing Domain identities for each switch and thus not require reassignment of address blocks. A disruptive Build Fabric event, by contrast, may cause some switches to acquire a different Domain identifier. This disruptive fabric reconfiguration in turn mandates that each attached end device re-login to the fabric to acquire new network addresses. Any time two or more fabric switches with identical Domain identifiers are connected together, a disruptive fabric reconfiguration is necessary to avoid address duplication.

Ironically, the automatic addressing and self-governing functionality that a Fibre Channel architecture implements to simplify network convergence is making storage networking more problematic. Disruptive fabric reconfiguration has become an issue for customers who are attempting to build large Fibre Channel fabrics or who are extending Fibre Channel fabrics over distance. In some cases, achieving stability of large fabrics is simply not possible, especially when multiple 16-port departmental switches are used instead of more centralized Fibre Channel directors. In other cases, sporadic disruptions on a particular Fibre Channel switch may quickly propagate to all others and trigger additional fabric reconfigurations. This situation is exacerbated when Fibre Channel fabrics are stretched over distance.

FIBRE CHANNEL SAN EXTENSION

The principal switch selection process is invoked whenever a potential addressing conflict occurs between two or more joined fabric switches. Whether two switches are separated by a few feet

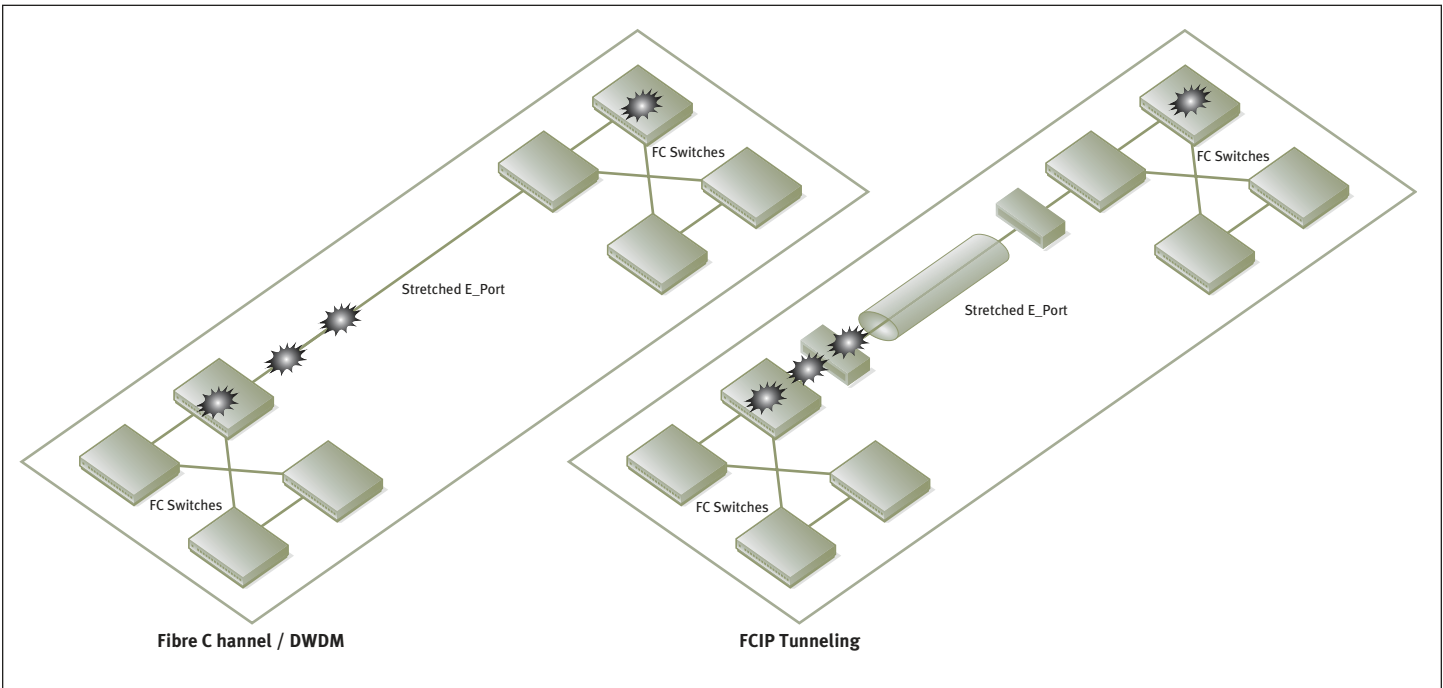


Figure 1: SAN extension using DWDM or FCIP tunneling

within a data center or by hundreds of miles between remote sites, establishing a connection between their expansion ports may trigger a fabric reconfiguration. A stretched fabric connection such as E_Port connectivity over a wide area network, however, is vulnerable to both disruptive fabric reconfigurations initiated by local site events and by those initiated by disruptions in the WAN transport.

Traditional SAN extension builds a single Fibre Channel fabric between two remote locations. As shown in Figure 1-A, the E_Port connection may be directly through dedicated fiber for metropolitan distances or through dense wave division multiplexing (DWDM) over a fiber optic cable plant. Native Fibre Channel extension over fiber may support metropolitan distances to a radius of 50 or more miles. To extend further, the Fibre Channel over IP (FCIP) tunneling protocol and IP network services may be used, as shown in Figure 1-B. FCIP simply wraps each Fibre Channel frame in an IP datagram for transport over an IP network. At the receiving end, the IP datagram

is removed and native Fibre Channel frames are delivered to the remote SAN.

Whether via direct Fibre Channel connection, DWDM or use of the FCIP protocol, standard Fibre Channel fabric building occurs. The stretched E_Port connection between Fibre Channel switches at either end triggers traditional principal switch protocols between them. From the standpoint of fabric building, the only difference between the stretched E_Port connection and a local data center E_Port connection is the additional latency introduced by the WAN link and the intervening FCIP or DWDM equipment.

As illustrated in the Figure 1 diagrams, a fabric reconfiguration event at one local site will necessarily propagate to the other site. If, for example, an operational Fibre Channel switch with a conflicting Domain identity is inserted into one site, the entire stretched fabric must undergo fabric reconfiguration. Likewise, a disruption in the wide area link may cause the single extended SAN to devolve into

separate SAN islands until stable service is restored. In either case, storage transactions are disrupted for the duration. For mission-critical storage over distance applications such as disaster recovery, an extended SAN may defeat the purpose of the installation itself. Created to ensure data availability, the extended SAN may inadvertently introduce instabilities that subvert highly reliable data access.

AVOIDING FABRIC RECONFIGURATION WITH SAN ROUTING

Because traditional SAN extension creates one large, flat Fibre Channel network, it has been difficult to create stable configurations for large deployments within data centers and between geographically remote sites. Analogous to Layer 2 (link layer) bridged LANs of the 1980s, extended SANs are vulnerable to flooding or broadcast storms that interfere with normal data traffic.

What is needed is a Layer 3 routing function that can provide connectivity between SAN islands while preserving the autonomy of each local SAN. This solution is provided by SAN Routing and use of the Internet Fibre Channel Protocol (iFCP) to eliminate fabric reconfiguration issues.

First introduced by Nishan Systems in its IP storage switches, SAN Routing provides interoperable E_Port connectivity to each local SAN fabric. Unlike conventional SAN extension, however, SAN Routing terminates the E_Port connection at each site. Fabric building is thus restricted to each location and Fibre Channel switch-to-switch protocols are not passed across the IP network. As shown in Figure 2, if two or more sites are connected by

McDATA IP storage switches, each site remains a separate SAN. Only authorized (zoned) connections between storage devices and servers are allowed across the IP network.

By preserving the autonomy of each local SAN, a SAN Routing solution ensures that disruptions at one site will be isolated and not allowed to propagate to other locations. As shown in Figure 2, a fabric reconfiguration event at SAN C would not affect storage transactions occurring between SAN A and SAN B. This provides the greatest stability for the connected SANs and promotes high availability for disaster recovery, consolidated tape backup and other storage applications.

Analogous to Layer 3 network routing, SAN Routing facilitates high performance transactions for each SAN segment, but does not

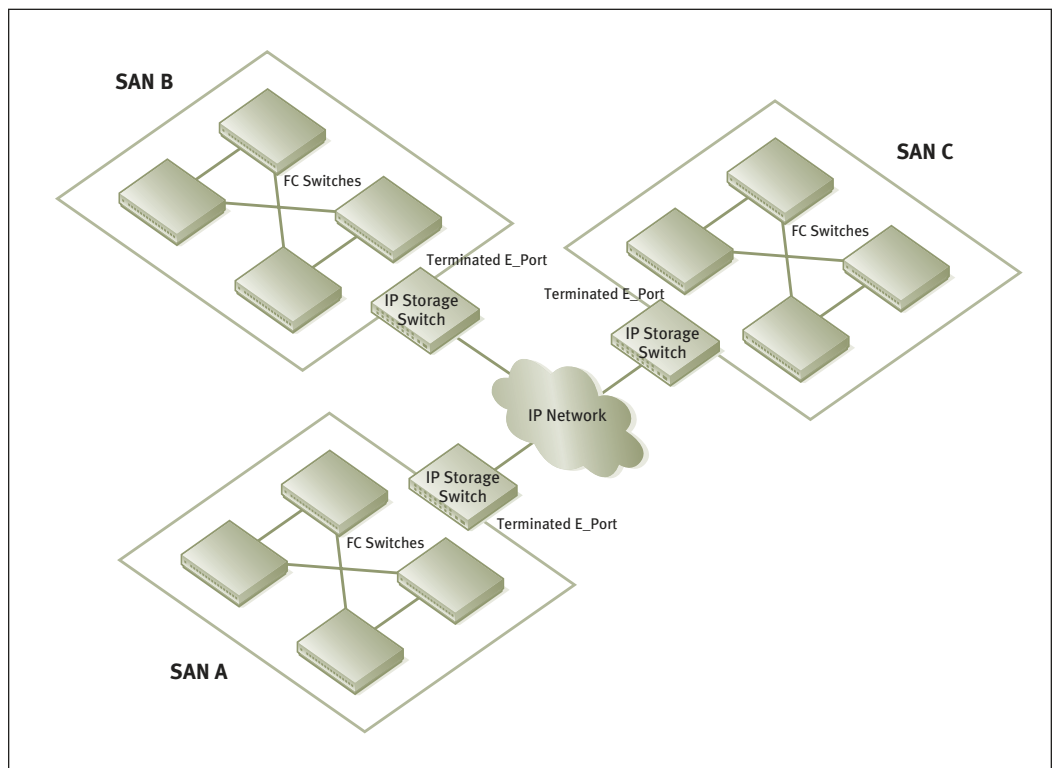


Figure 2: SAN routing using iFCP preserves the autonomy of each site



expose the storage network as a whole to potential disruption. Just a Layer 3 routers eliminated the broadcast storms common to bridged LANs, McDATA's IP storage switches and SAN Routing provide fault isolation against disruptive Fibre Channel fabric reconfigurations. This enables customers to deploy complex, multi-SAN storage solutions while leveraging the cost benefits of available and affordable IP network services.

In addition to fault isolation, a SAN Routing solution streamlines SAN connectivity by eliminating address overlap issues. McDATA's SAN Routing permits use of duplicate Domain addressing on separate Fibre Channel switches without incurring routing errors. In the Figure 2 diagram, for example, a fabric switch in SAN A could have the same Domain identity as a fabric switch in SAN C. Since the switches at both sites remain in separate SANs, address conflicts between the SANs do not occur. The SAN Routing engine provides the appropriate address translation only for the devices that have been authorized to communicate across the network. The administrator, therefore, does not have to constantly monitor or manipulate switch addresses to prevent Domain ID conflicts.

SAN ROUTING AND INTEROPERABILITY

One of the persistent problems associated with large Fibre Channel fabrics has been the issue of multi-vendor interoperability. Due to delays in establishing common standards for switch-to-switch communications and the Fabric Shortest Path First (FSPF) protocol, Fibre Channel switch vendors must now support multiple interoperability modes. In addition to standards-compliant E_Port mode, a vendor may also support one or more proprietary modes. Furthermore, ongoing enhancements to both proprietary and open systems switch-to-switch protocols result in various microcode versions for a particular switch model. A switch resold through an OEM is typically one microcode version level behind the same switch sold directly by the manufacturer. This makes it difficult for customers to directly connect fabric switches to build larger SANs, even if the switches are made by the same vendor. Simply upgrading to a common microcode version may not be possible, either due to disruption of product traffic or fear of violating the OEM's warranty.

To address this long-standing interoperability issue, McDATA IP storage switches provide concurrent support for multiple E_Port compatibility modes, including vendor-specific and open systems versions. This makes it possible to implement SAN Routing for customers who have both OEM and direct-marketed versions of a vendor's switches, as well as switches from different vendors. Thus in addition to the benefits of fault isolation and stability provided by SAN Routing, customers can leverage their existing Fibre Channel switch assets that otherwise would be a source of perpetual interoperability conflicts.

SAN ROUTING IN THE DATA CENTER

SAN Routing provides the most robust solution for connecting existing SAN islands together. By maintaining the autonomy of each SAN site and restricting propagation of potential faults, SAN Routing is especially suited to storage over distance applications such as disaster recovery and remote tape backup. In some data center situations, however, it is desirable to link separate SANs to share tape subsystem or other storage assets. To avoid building one large Fibre Channel fabric with its accompanying fabric reconfiguration issues, SAN Routing can be used to provide both connectivity and fault isolation.

As shown in Figure 3, a single McDATA IP storage switch can be used to connect multiple SANs within a data center, building or campus. Each SAN continues to function as a separate fabric, and yet each has access to authorized storage resources through the IP storage switch. This provides the benefits of shared access while eliminating interoperability and fabric building issues. A disruption in the Development SAN, for example, would be restricted to the Development group, while Production, Finance and Engineering could continue normal storage operations. This solution is also highly scalable and can accommodate additional SANs and IP storage switches to support a large population of devices over time.

SUMMARY

McDATA's implementation of SAN Routing offers an innovative, interoperable and standards-compliant solution for customers who



want to leverage their existing Fibre Channel assets for disaster recovery and other storage applications and yet avoid the inherent vulnerabilities of Fibre Channel fabric building and SAN extension. By providing a SAN routing function for Fibre Channel fabrics, SAN Routing enables customers to build very large storage networks and

extend SAN traffic to virtually any distance and any number of sites. This provides new opportunities for customers to implement enterprise-wide SAN strategies that are stable, interoperable and scalable.

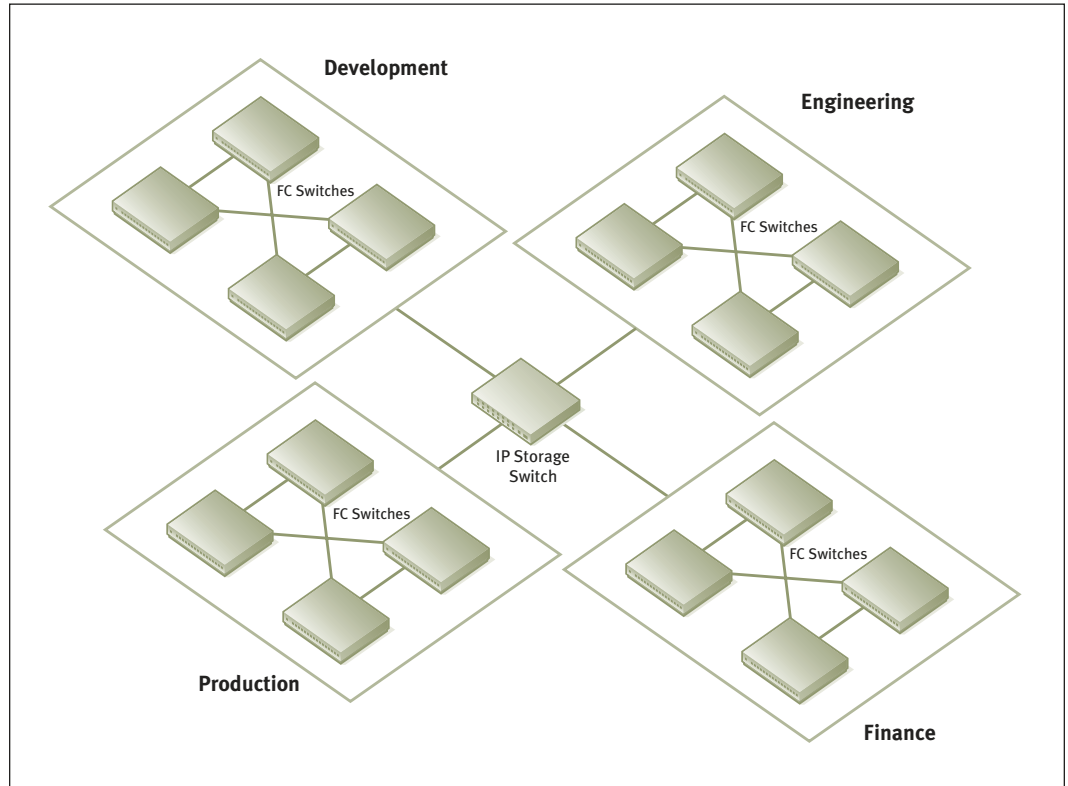


Figure 3: SAN routing in the data center or campus

