White Paper

Dynamic Zoning for Arbitrated Loop

Tom Clark
Director, Technical Marketing

September 1999

VIXEL
MAKING THE FIBRE CHANNEL CONNECTION

Storage networks based on Fibre Channel arbitrated loop represent the largest installed base of SANs in production environments. Arbitrated loop hubs have been a popular choice for departmental SANs due to ease of installation, lower cost relative to switches and high performance for small and medium configurations.

A fundamental feature of arbitrated loop protocol, however, has excluded loop hubs from certain applications. Since arbitrated loop is a shared transport, the protocol used to assign unique addresses to all loop participants is often disruptive to delay-sensitive applications such as streaming tape backup. This loop initialization routine, or LIP (loop initialization primitive), only takes milliseconds to complete, and yet may cause a tape backup process to abort. Lacking any means to isolate these normal LIP processes, vendors have recommended fabric switches to avoid loop-specific issues such as LIPs. By using a different set of Fibre Channel protocols, fabric switches can be used to create LIP-free environment, but with the accompanying higher cost and complexity of fabrics. Until now, customers have thus been forced into more expensive solutions just to achieve more stable operation. Vixel's new 2100 Zoning Hub is the first Fibre Channel product to resolve this inherent aspect of arbitrated loop, and will allow customers to enjoy the benefits of simplified installation and lower cost while maintaining application integrity.

What's Wrong With LIPs?

Loop initialization is an essential part of life on an arbitrated loop and actually simplifies loop administration. The loop initialization procedure automates the process of address assignment and insures that every device attached to the loop has a unique identity. This relieves the administrator of any manual supervision of addressing whenever new devices are added to a loop.

As a new device is inserted into the loop, it will issue a series of loop initialization primitives, or LIPs, to let the other participants know of its entrance. The new arrival may not have a previously-assigned address (e.g., it if was just powered on), or may have a potentially conflicting address (e.g., if it was moved from one loop hub to another). If no loop initialization process occurred, the new device would either not be able to communicate on the loop or might erroneously receive data intended for another loop node. A newly inserted disk, for example, might have a duplicate address of an existing disk, resulting in data corruption on its inadvertent twin. Loop initialization is therefore essential to guarantee that all loop devices are unique and that servers are talking to their appropriate targets.

The loop initialization routine itself is fairly straight-forward: As shown below, when a device is inserted into the loop (or powered on), it issues a stream of LIPs onto the loop. Codes associated with specific LIP primitives indicate why the device is requesting a loop initialization. A LIP(F7,F7), for example, indicates that the device is entering an active loop and has no valid arbitrated loop physical address, or AL_PA. As the LIPs propagate around the loop, all devices cease their previous activity and enter an initialization state. A temporary loop master is selected to oversee address assignment, and a series of frames are circulated around the loop to allow each participant to select a unique address. Once address selection and identification is complete, the loop devices exit the initialization

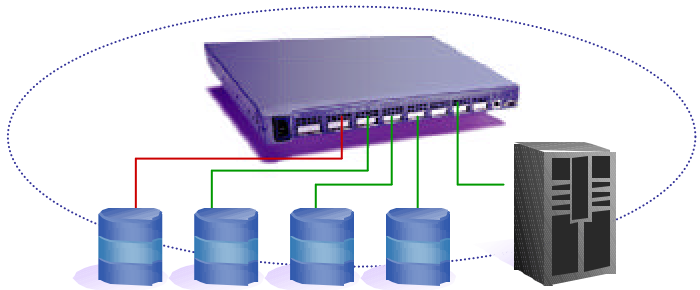

Normal loop operation
A new device is inserted and issues LIPs
- **Suspend all current data transactions**
- **Enter loop initialization mode**
- **Select temporary loop master**
- **Issue addressing frames**
- **Optional address mapping phase**
- **Exit initialization**

Resume previous data transactions
Normal loop operation

routine and resume normal operation. For small and medium loop populations, this entire process lasts only for a few hundred milliseconds or for large loops, a few seconds.

Figure 1. Normal Loop Initialization Routine

When a newly inserted device sends out its LIPs, all previous data movement is suspended until the initialization routine is completed. For most operating systems and applications, the loop initialization process is so brief that the upper layer SCSI transactions will not time out. A write operation from a server to its target, for example, would be momentarily interrupted, and the operation resumed once normal loop activity was reestablished. Thus for the vast majority of storage applications, LIPs do not present a problem and this allows arbitrated loop to accommodate dynamic adds, moves and changes to the SAN without disruption.

The exceptions to this general rule are presented by delay-sensitive applications like full motion video and streaming tape backup. Streaming video cannot tolerate any disruptions, and a spontaneous LIP procedure may cause the video image to momentarily fragment. Likewise, streaming tape backup expects a continuous flow of data from server to tape. A disruption of this flow typically causes the tape backup to abort, since the tape controller cannot reposition the tape quickly enough to resume operation once the stream is reestablished. After some initial experimentation with arbitrated loop, most video applications for SANs are now implemented with switches. Tape backup configurations may still use arbitrated loop hubs, but typically for small, dedicated configurations that are not subject to changes to the topology. For more complex, multi-server backup applications, switches are normally deployed.

LIP Isolation

To leverage the lower cost and greater simplicity of arbitrated loop for delay-sensitive applications, it is necessary to restrict the propagation of LIPs around the loop. LIP isolation presents several engineering challenges, since controlling LIPs may potentially result in duplicate addressing. Vixel pioneered LIP isolation on its first switch product, the Vixel 4000, and has migrated this feature to its next-generation, Vixel 8100 fabric switch. Since the Vixel 8100 controls the address allocation for arbitrated loop devices, it can insure that restricting LIPs does not result in duplicate loop addresses. Through Vixel's SAN InSite management application, a customer can select which ports should receive LIPs if a new loop device is attached. This value-added feature has made the Vixel 8100 highly attractive for video applications that require the bandwidth of switching and the ability to support private loop protocols. Although the Vixel 8100 is the lowest cost fabric switch in the SAN market, opportunities still exist for hub-only solutions that can restrict the propagation of LIPs.

The Vixel 2100 managed hub is the first Fibre Channel product to introduce LIP isolation specifically for loop hub applications. This enhancement to Vixel's product portfolio provides customers with a cost-effective means to implement arbitrated loop for tape backup and other applications, and provides greater flexibility for the design of both departmental and enterprise SANs. As shown in Figure 2 below, when the Vixel 2100 is zoned for two or more loops, loop initializations on one loop do not effect the operation of the other loop zones.

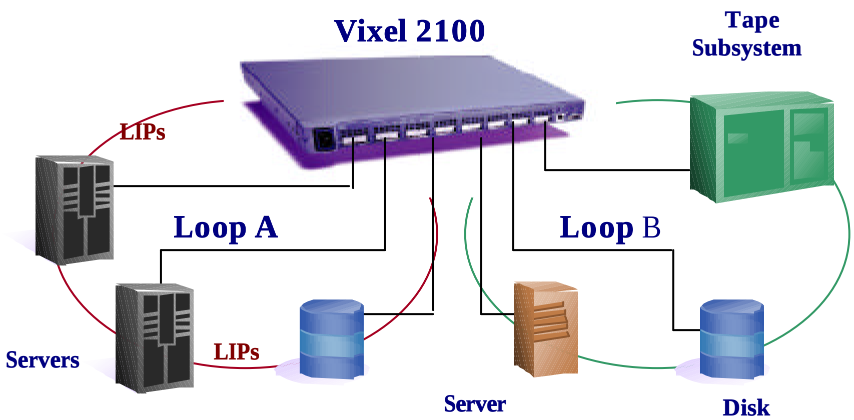

Figure 2. LIP Isolation Between Loop Zones

LIP Isolation and Dynamic Zoning

The Vixel 2100 accomplishes LIP isolation by a technique called *dynamic zoning*. Zoning allows any of the 2100's ports to be associated in an exclusive loop segment or mini-loop for a specific transaction, while program control of these port groupings allow the zones to be dynamically created and reconfigured at will. Since the members of a particular zone are not subject to the activity of other hub ports outside the defined zone, LIPs issued on the other hub ports are automatically restricted.

The 8 ports of the Vixel 2100 can be configured as many as 4 distinct zones. Unlike some segmenting hub products which simply divide a single hub into two loops (e.g., ports 1-4 for Loop A, and ports 5-8 for Loop B), Vixel's dynamic zoning allows any ports to be configured into a common zone. In the diagram below, ports 2 and 5 are occupied by servers, ports 1,3,4 and 6 have attached disk arrays, and port 8 supports a tape subsystem.

During normal operation, it may be desirable to have all hubs ports configured as a single loop. During a tape backup of the server on port 2, however, the tape backup process should be shielded against any LIP disruption from other ports. If a new device were inserted into port 7, for example, a LIP would likely abort the backup procedure. Prior to launching the tape backup routine, the Vixel 2100 can be programmed to create a zone which would include port 2 (the designated server), port 4 (the target to be backed up), and port 8 (the tape subsystem). Creating this zone would automatically cause a new loop initialization for the zoned members, as well as the default, secondary zone (ports 1, 3, 5, 6 and 7). This ensures that the devices in the newly zoned segments have unique addresses, and since the tape backup process has not yet been started, this initialization causes no disruption. Once the zone is created, the tape backup can proceed.

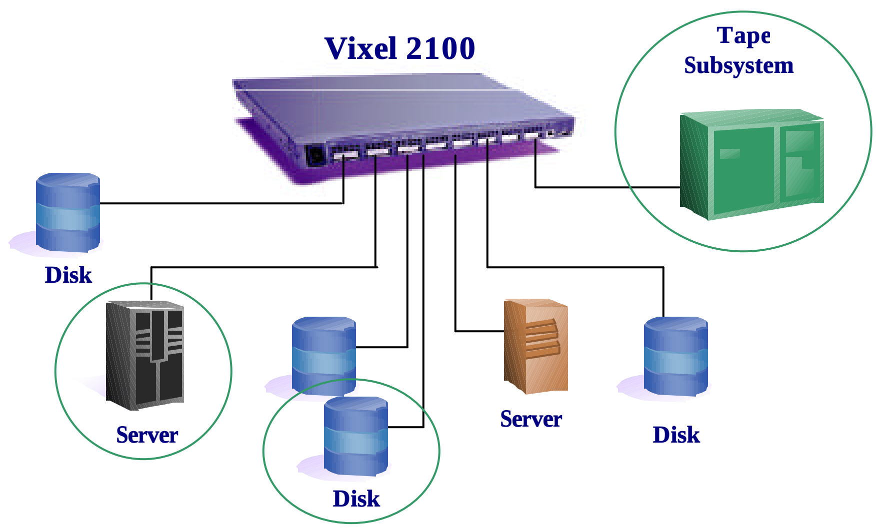

Figure 3 Backup Zone for Server, Disk and Tape Subsystem

Since the server, target disk and tape have been isolated from the rest of the hub, the server can initiate its backup without disturbance. If the server on port 5, for example, is powered off and on, its LIPs will only propagate on ports 1, 3, 5, 6 and 7. Ports 2, 4 and 8, and the tape backup underway, would not be effected. In addition, once the server on 2 port is finished backing up the disk array on port 4, a new zone can be dynamically created to backup the disk on port 6. This would allow the array on port 4 to be accessed by the other server, while giving the server on port 2 exclusive access to the disk on port 6 and the tape subsystem. Likewise, if the server on port 5 is scheduled to backup its array on port 3, the previous zoning can be dissolved and a new zone created to support ports 3, 5, and 8. In each instance, as new zones are created, an initialization is triggered to insure unique addressing, and then the zone-specific application can be launched.

Vixel's dynamic zoning provides a number of methods for user control of zone definitions. Zones may be created using Telnet scripting via the 2100's console interface, through a web browser interface, by setting SNMP MIB variables or through Vixel's SAN InSite management platform. Future implementations will provide an API for embedded control from third-party backup applications. For Unix environments, simple shell scripting can be used to automatically create zones, perform backups, and dynamically allocate new zones.

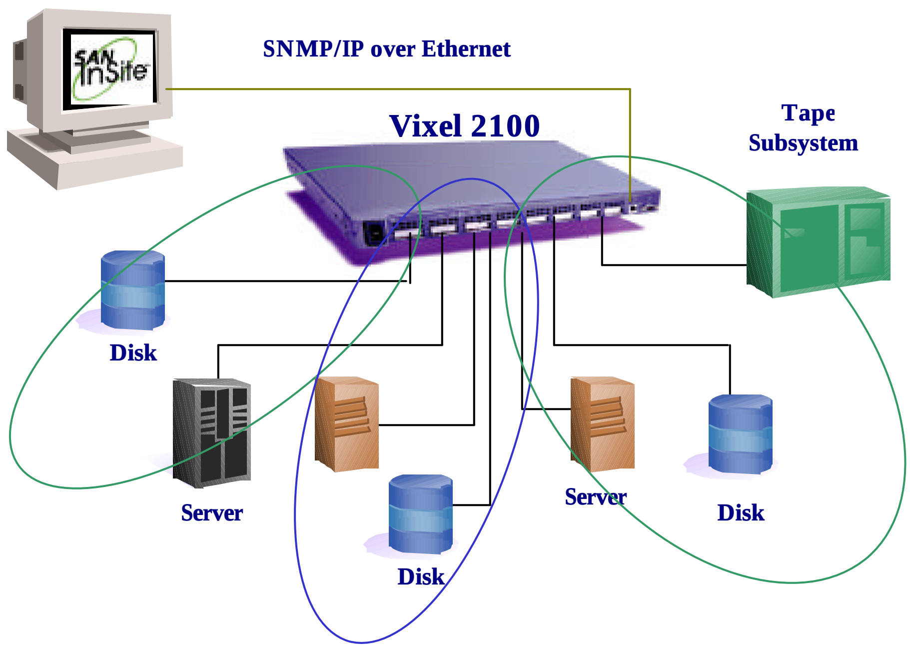

Figure 4. Dynamic Zoning via Vixel SAN InSite or Telnet Scripting

Other Benefits of Dynamic Zoning for Arbitrated Loop

The Vixel 2100's ability to isolate LIPs from delay-sensitive applications makes it an ideal choice for tape backup and other solutions for the mid and low ranges of the SAN market. A departmental SAN can maximize use of an expensive tape resource between multiple servers, but at a much lower cost for the Fibre Channel interconnect.

In addition, dynamic zoning enables a single hub to be used for multiple, discrete loops. Database migration, for example, could exploit the 2100's zoning capability to isolate data conversion traffic from other loop participants. Or, two or more small SAN configurations could be built using a single 2100, allowing several departments to leverage their SAN investment. As shown below, Engineering, Human Resources and Finance could share a single hub, without risk that users in one department would inadvertently access data in another's.

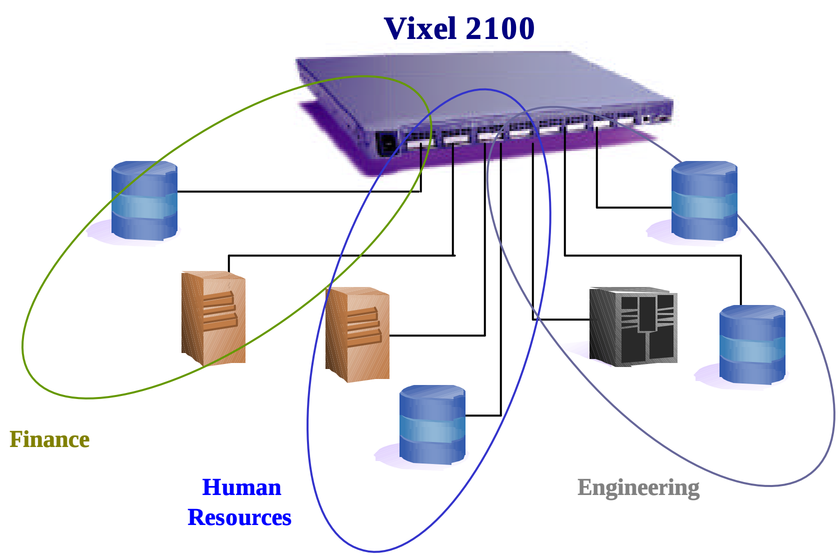

Figure 5. Semi-permanent Zones for Multi-departmental Configuration

The Vixel 2100's dynamic zoning with LIP isolation is a continuation of Vixel's commitment to provide value-added, low-cost solutions for the SAN market. Along with Vixel fabric switches, managed hubs, entry-level hubs, transceivers and comprehensive SAN InSite management software, the Vixel 2100 offers customers all the essential building blocks for constructing the SAN and the means to maximize the benefits SANs offer.

About the author:

Tom Clark is Technical Marketing Director for Vixel Corporation and the author of *Designing Storage Area Networks: A Practical Reference for Implementing Fibre Channel SANs*, Addison Wesley Longman.

For more information on the above subject, contact your reseller:

Unylogix Technologies Inc.
Tel.: (514) 253-5200
email: info@unylogix.com
web site: www.unylogix.com

Copyright © Vixel Corporation. All Rights Reserved. Contents subject to change. Part #00041069-001 Rev. A