The Fibre Channel Consultant Series

Fibre Channel Arbitrated Loop

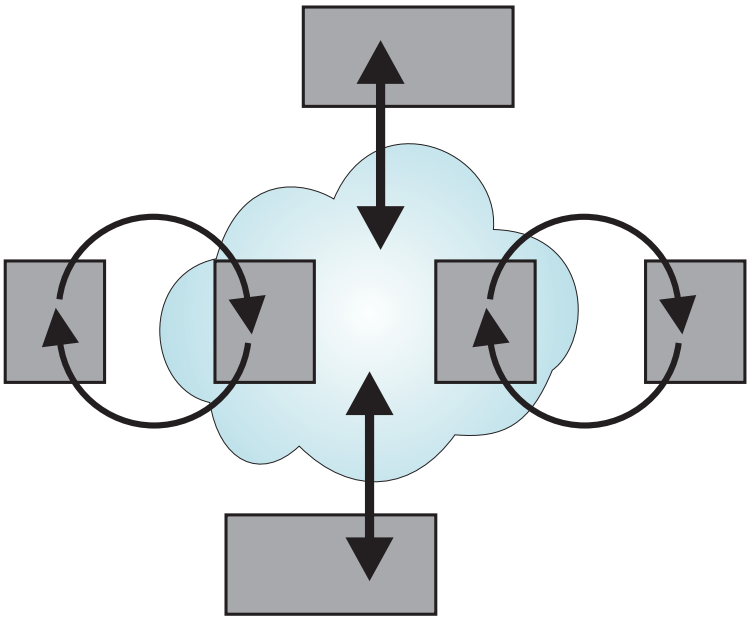

Robert W. Kembel
Horst L. Truestedt

NIA INC. TUCSON

Cover design by David Fischer, Fischer Graphic Services.

ISBN 0-931836-82-4

Published by:

Northwest Learning Associates, Inc.
3061 N. Willow Creek Drive
Tucson, AZ 85712
520-881-0877, Fax: 520-881-0632
www.nlabooks.com

Printed in the United States of America

20 19 18 17 16 15 14 13 12 11

15. Arbitrated Loop Hubs

A Fibre Channel Arbitrated Loop Hub is a device for interconnecting arbitrated loop ports. There are a number of various hub types available today which represent different cost targets and design points. But, before going into these, let's look at what a simple arbitrated loop hub is and how it works.

As has been seen elsewhere in this document, loops may be wired port to port in a cascading configuration. However, if any loop port fails or is not powered on, the loop is not operational.

Hubs are capable of automatically bypassing a loop port (or a loop segment containing more than one loop port) when a loss of signal is detected at the hub. This would occur if no loop port is connected, a connected loop port is powered off, or a loop segment is broken. A hub solves this problem by using port bypass circuits to detect and bypass a non-operational loop segment. The port bypass circuit either opens the loop to insert the active loop ports or closes the loop, ensuring the loop is operational. Hubs provide the ability to hot-plug loop ports in and out of the loop just like the port bypass circuits are used on a backplane for disk and tape device. A block diagram of a simple hub providing this bypass function is shown in Figure 119.

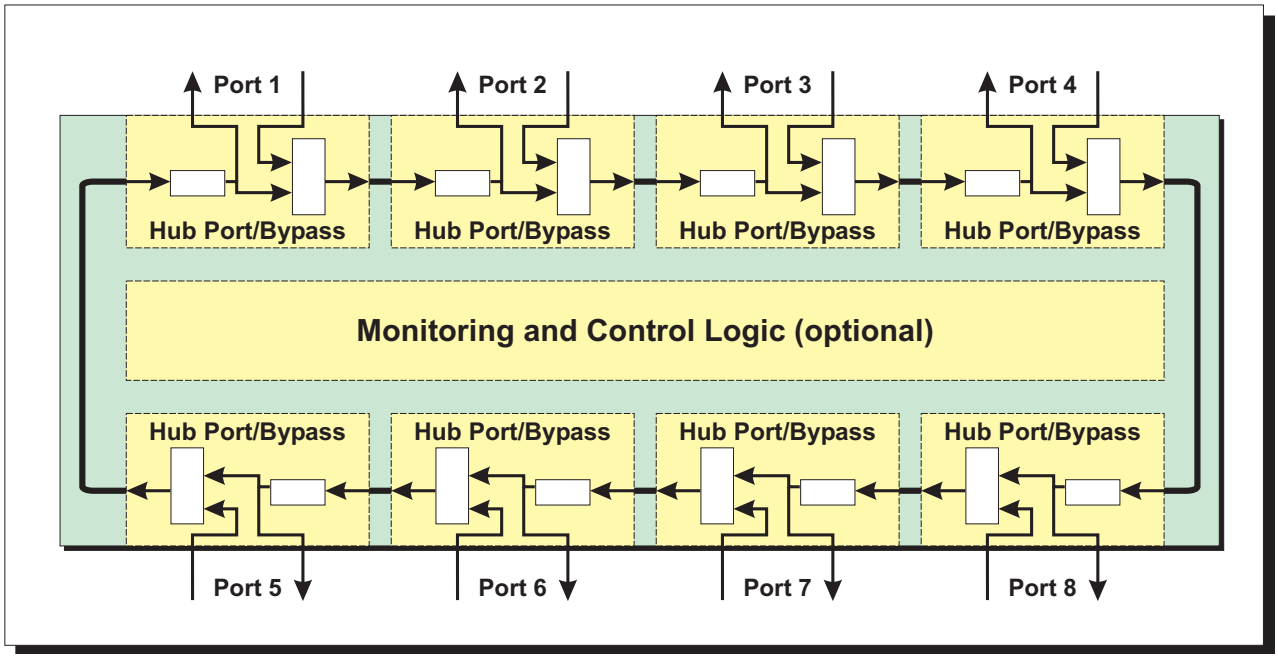

Figure 119. Fibre Channel Hub

The hub provides the bypass function for each of its ports and control logic to determine when an attached loop port (or loop segment) should be bypassed. The means by which a hub decides to bypass the attached loop segment is a function of the design of the hub. The hub may

simply monitor transmitted output of the attached loop segment to determine if the loop segment is working correctly. If no transmitted signal is received, the segment is bypassed. Hubs may also monitor the received signal for a LIP(F8) which indicates that an attached loop segment has detected a failure. As long as this failure persists, the hub could bypass this segment and reactivate the segment when the LIP(F8) is changed to a LIP(F7).

More sophisticated approaches may monitor information sent to the segment and verify that the segment is generating appropriate responses. The hub may also interpret the LPE and LPB primitive sequences to control the bypass function.

The port on a hub is just an attachment point and one or more loop ports may be attached to one hub port. Hubs typically provide between six and 16 ports each and can normally be stacked or cascaded to support larger configurations of up to 127 ports. Normally, the hubs themselves are transparent to the loop ports and provide no additional addressing to an arbitrated loop environment. However, some hubs contain embedded NL_Ports that can be accessed for monitoring and management purposes.

A hub changes the electrical loop into a physical star wiring environment and makes it possible to physically cable a loop. Without a hub, the cables must be separated to go from the transmitter of one loop port to the receiver of the adjacent loop port (although many devices provide dual loop capability for redundancy, it is usually not desirable to contain both loops in the same cable plant). With a hub, the normal bi-directional Fibre Channel cables may be used (in this case, the loop is contained in each direction of the cable).

Many arbitrated loop hubs offer a mix of electrical and optical ports to facilitate mixing the two media types within the same loop. For example, one or more systems could use optical links to connect to the hub, while one or more disk enclosures are connected to the hub using electrical cables. This flexibility may be provided through a fixed set of media interfaces or provided by pluggable GBIC modules.

Figure 120 on page 277 illustrates an Arbitrated Loop with three servers and three storage devices connected to a hub. While this configuration lacks the appearance of an Arbitrated Loop, it is in fact still a loop. The hub provides a centralized wiring point changing the functional loop into a physical star configuration.

Hubs provide considerable flexibility in configuring Arbitrated Loops while at the same time providing loop robustness. It should be remembered, however, that hubs themselves are electronic devices and may fail. In high-availability applications, the need for redundant loops and dual ported devices is expanded to include redundant hubs.

15.1 Insertion Modes of Hubs

There are several different ways that hubs may decide when to insert loop segments into a loop. The following sections provides several examples.

Insert on Valid Transmission Word. The hub may simply decide to insert a loop segment when any valid transmission word is received.

If a port transmits IDLE, LIP, any other ordered set containing a K28.5 character, or any other valid 8B/10B encoded character during power-up, this type of hub may add the loop segment

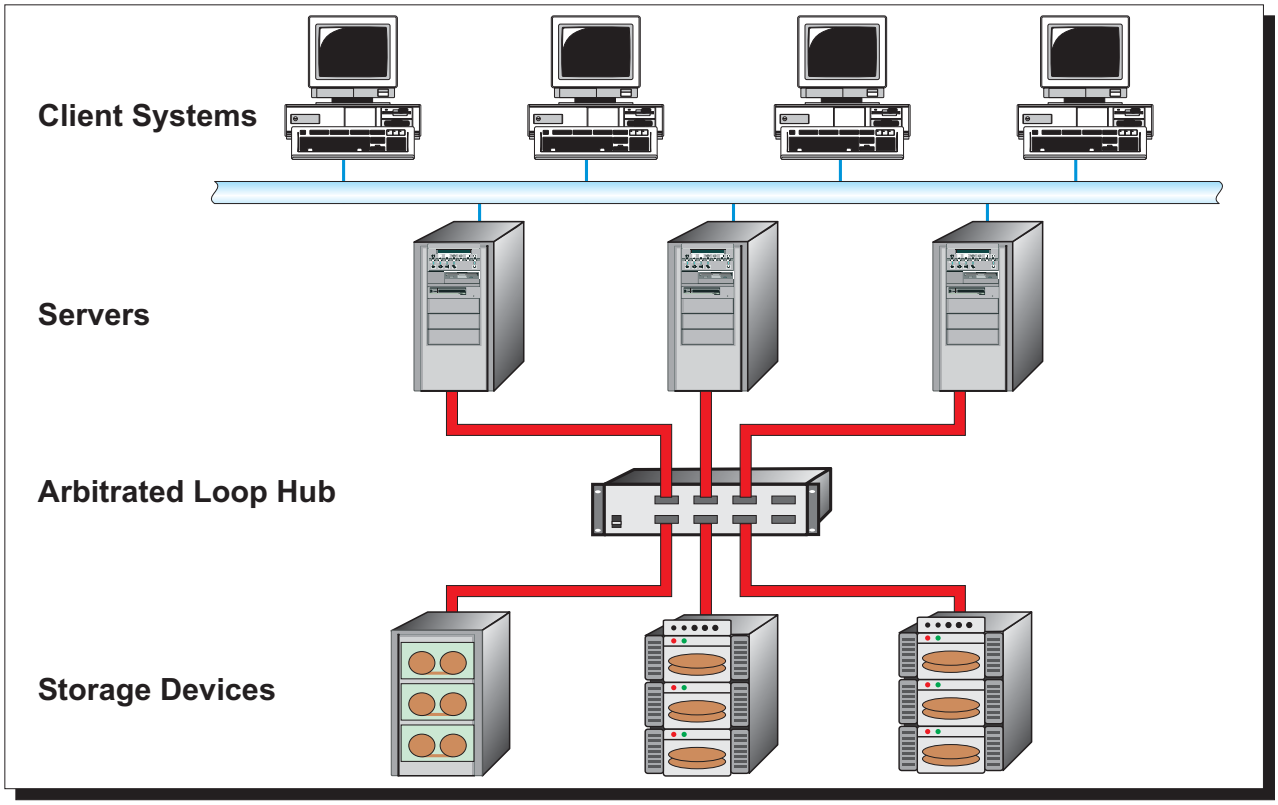

Figure 120. Arbitrated Loop with Hub

into the loop before the port(s) are ready to participate. This would make the loop non-functional while the port(s) are becoming ready.

Simply inserting on any valid transmission word may allow non-loop capable ports to join the loop. When a non-loop capable port attempts initialization, it transmits valid Fibre Channel transmission words.

Insert on Valid Ordered Set. A hub may insert a loop segment when ordered sets have been received for a minimum time period.

However, if a port transmits IDLE, LIP, or any other ordered set containing a K28.5 character during power-up, this type of hub may add the loop segment into the loop before the port(s) are ready to participate. This would make the entire loop non-functional while the port(s) are becoming ready.

This approach may also allow non-loop capable ports to join the loop. When a non-loop capable port attempts initialization, it transmits valid Fibre Channel ordered sets (OLS, NOS).

Insert on LIP. Hubs may be more selective by waiting for a valid, loop-specific ordered set such as the Loop Initialization Primitive (LIP) sequence before inserting an attached loop segment. This ensures that the hub does not insert a non-loop capable port into the loop. Howev-

er, if the hub inserts on any LIP, it is possible the hub may insert a broken segment if LIP(F8) is being received.

Insert after Diagnostic Verification. Hubs may use even more sophisticated methods to determine when to insert a loop segment.

A hub may be able to determine that the loop segment is operational by transmitting an ordered set that is not used by an loop port, but which would be forwarded by all loop ports (e.g., transmit a LIP(F0,F0)). Once this LIP is received by the hub, it could add the loop segment since all loop ports on the segment are now ready to participate on the loop.

Insert Under Management Control. Hubs may also wait for an external input (i.e., operator, SCSI Enclosure Services (SES), or SNMP request) to insert a loop segment into the Loop.

15.2 Removal Modes of Hubs

Hubs may automatically bypass a loop segment if it detects the following conditions:

1. If loss of signal or synchronization (>R_T_TOV), the loop segment is isolated.
2. If LIP(F8) is received, the loop segment is isolated since at least one loop port on this loop segment is reporting a loop failure.

15.3 Sophisticated Hubs

Considerable sophistication may be built into a hub, usually at additional cost. One of the purposes of a hub is to provide a number of functions, but for less cost than a Fibre Channel switch. A comparison between hubs and switches can be found in Figure 26.

ITEM	HUBS	SWITCH
Power supplies	Usually two for redundancy	Usually two for redundancy
Memory	Minimal (possible speed matching buffers if using re-clocker)	Considerable memory for each FL_Port to support a number of frames
Media adapters	GBICs, GLM, Electrical	GBICs, GLM, Electrical
Enclosure	Sheet metal	Sheet metal
Fabric Services	None required (could provide an F/NL_Port once private loop ports disappear)	Required to support large address space
Maximum number of N*_Ports	126	Approximately 14 million
Throughput	Maximum full duplex circuit (e.g., at 1Gbs, 200MB/s)	Number of ports times data rate
Error Isolation	Remove failing loop segment	Remove failing loop or N_Port
Media Converters	Optical and electrical supported	Optical and electrical supported

Table 26. HUB and Switch Comparison

Simple hubs may not re-time the signal; sophisticated hubs usually re-time, eliminating signal jitter and providing automatic clock speed matching on every port. As a result, these hubs can be cascaded, providing increased scalability and flexibility in network design. It was already indicated that sophisticated hubs may provide considerable loop management capabilities (e.g., configuration, error, and performance statistics); automatic LIP(F7) insertion, LIP(F8) bypass, automatic port bypass and character integrity checks.

15.3.1 Switching Hubs

Some hubs are capable of providing more performance for a loop environment. As shown in this book, only two loop ports are allowed to communicate with each other at any given time (excluding the replicate mode). With a “switching” hub, multiple loop segments may be created to allow more than two loop ports to simultaneously communicate. These hubs provide the following operational characteristics (note: for this to function properly, since there are no hub port frame buffers, a login BB_Credit of zero must be used):

1. During the initialization process, all loop segments are combined into one large loop of up to 126 NL_Ports and 1 FL_Port.

2. During the initialization process, the hub port identifies the AL_PAs that are contained in each loop segment (the difference in the loop initialization AL_PA bits of what the hub port transmits and what it receives back).

3. After the initialization process, the hub ports become repeaters on each loop segment. Any port that arbitrates on a loop segment only interferes with the other ports on the same loop segment. Once a port wins arbitration, it transmits the usual OPN(yx). If the open recipient is on the same loop segment, the OPN(yx) is delivered directly and a loop circuit is created; if the open recipient is on another loop segment, a connection is made between the two loop segments through the hub and the OPN(yx) is transferred to the other hub port where the open recipient resides. The hub port arbitrates on this loop segment using the x value of the OPN(yx) (i.e., the open originator) and when it wins, forwards the OPN(yx) to the destination port. Since a login BB_Credit of zero is required, the open originator waits for the R_RDY which provides credit to send frames.

Switching hubs such as this need to avoid possible deadlock conditions where a port one loop segment is attempting to open a port on a different loop segment where some other port has already won arbitration and want to open yet another port.

15.4 Hub Summary

Fibre Channel Arbitrated Loop hubs facilitate loop implementation by aggregating loop ports via a physical star configuration. Loop hubs typically provide 6 to 12 ports, and can be used to build larger loops via cascading hubs. hubs provide greater control and reliability and are usually found in a wiring closet (i.e., cable concentrator) to facilitate management of a Fibre Channel network. Hubs employ port bypass circuits at each port to keep dysfunctional L_Ports from disrupting loop traffic. Most hubs provide status and diagnostic LED's at each port.

15.5 Chapter Summary

Arbitrated Loop Hubs

- Provide significant enhancements to the loop
- Physical cabling flexibility
  - Allows use of standard cables
  - May provide media flexibility through use of GBICs
- Automatic bypassing of faulty loop segments
  - Hub makes decision to bypass loop segment
  - Criteria based on received signal, or
  - Hub may validate loop segment before insertion into loop
- May provide signal reconditioning
  - Signal amplitude regeneration
  - Signal timing restored

Hub Insertion Modes

- Insert on valid transmission word or valid ordered set
  - May admit non-loop ports
  - May admit ports prematurely
- Insert on loop-specific ordered set such as LIP
  - May admit broken loop segments - LIP(F8)
- Insert after validation
  - Validate loop segment using LIP(F0,F0) or similar ordered set
- Insert on management directive
  - Wait for external input to insert loop segment

Hub Deletion Modes

- Hub may bypass a loop segment on
  - Loss of signal or loss of synchronization
  - Receipt of LIP(F8)
  - Receipt of LPB for loop port/segment
  - If a link error threshold is exceeded
  - Under management directive
- Most hubs provide visual indication of bypassed segments
  - Light on hub panel
  - Via management interface

Switching Hubs

- Switching hubs may allow for configuration changes
  - Controlled by external inputs
  - Ethernet control port, for example
  - Allows loop to be configured into different segments
- Switching hubs may dynamically connect loop segments
  - Dynamically segments and joins loop segments
  - Driven by arbitrated loop protocols directly
  - May impose constraints on ports (such as use of zero login credit)
  - Treated as one large loop during initialization