Arbitrated Loop Switches

What is a Loop Switch?

- **A "Loop Switch" is a low-cost switching device for arbitrated loop environments**
  - It is often embedded within a disk enclosure (referred to as a "Switched Bunch of Disks" or SBOD
  - The cost per port is around $4
  - A loop switch is controlled by the existing loop protocols and is transparent to the attached devices
    - No changes required to disk drives or initiators due to a loop switch being present
  - Because there are no changes to the Arbitrated Loop protocols you won't find loop switches mentioned in the standards

What is a Loop Switch?

- **A Loop Switch is not a Fabric switch**
  - A Loop Switch does not provide fabric services
    - No Name Server function
    - No Management Server
    - No State Change Registration or Notification
    - No FLOGI or address assignement
  - A Loop Switch does not buffer frames
    - It does not participate in flow control
    - It does not do speed matching (all ports must operate at the same speed)
  - It is somewhat analogous to a Serial Attached SCSI (SAS) expander

Disk Enclosure with Loop Switch

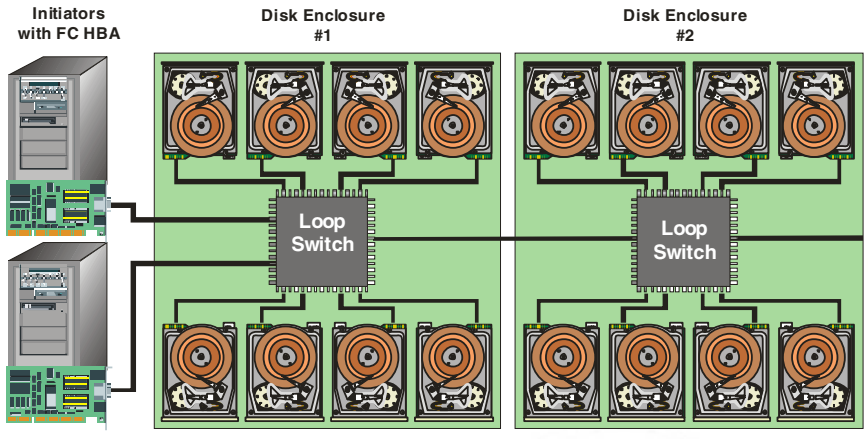

Advantages of Loop Switch

- **Loop switches provide a number of benefits:**
  - Traffic does not flow through intermediate devices
    - Avoids the accumulated latency inherent in large loops
    - Avoids the potential for traffic corruption in intermediate devices
  - Loop switches improve performance
    - All communications is effectively as if it were on a 2-ported loop
    - Multiple devices can communicate simultaneously
    - For example, two initiators can communicate with two different disks at the same time)
  - Loop switches improve availability and serviceability
    - If a device fails, it is automatically isolated from the loop by the switch
    - When a device does fail, it is easier to determine which device has failed
    - Loop switches may minimize disruptions due to loop initialization

Loop Switch Architecture

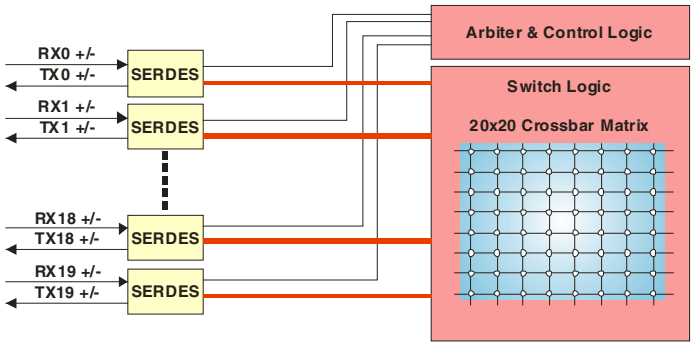

- **Loop Switch consists of:**
  - Serializer/Deserializer (SERDES) for each port
  - Arbiter and Control Logic
  - Crossbar switch matrix
  - Management and control functions

Loop Switch Connections

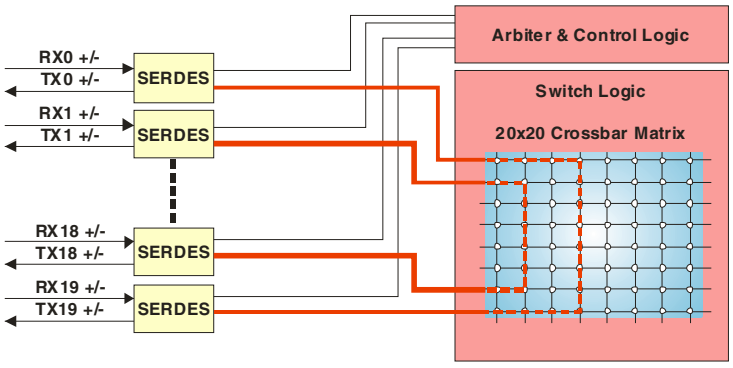

- **The switching function connects pairs of ports**
  - Driven by the existing loop protocols
  - Provides a switched connection between the loop ports
  - Multiple concurrent connections can exist

Loop Switch Operations

- **When a loop port needs access to the loop, it must arbitrate**
  - Transmits its own ARB(x) in place of received IDLEs and lower priority ARBs
    - Loop switch simply returns the received ARB(x) allowing the port to win arbitration
  - Loop port transmits an OPN ordered set to open a loop circuit
    - Loop Switch traps the open and looks at the destination AL_PA
    - If the destination is available, the switch matrix is configured to connect the source and destination ports and the OPN is forwarded to the destination creating a normal loop circuit
    - If the destination is not available, the loop switch can wait or respond with CLS (an indication that the destination is not ready to receive frames)
  - Normal frame transmission and flow control
    - The connection provides a full-duplex circuit between the two ports
  - When a CLS is detected in both directions, the connection is ended

Loop Switch Operation

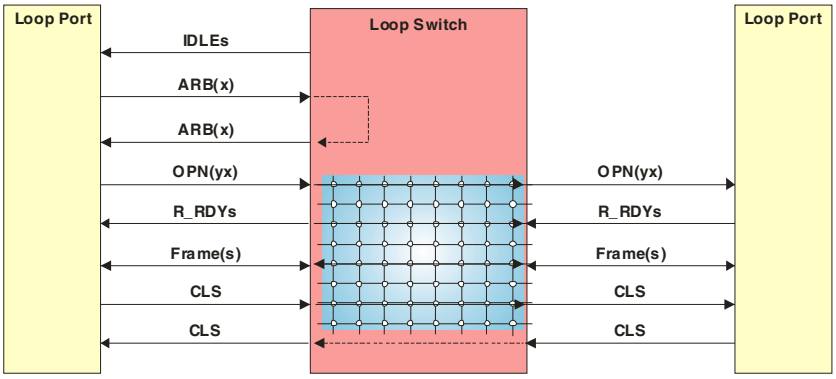

- **OPN creates the connection**
- **CLS in both direction ends the connection**

Managing Loop Initializations

- **Loop switches can minimize disruption due to loop initialization**
- **Normally, when a new device is added to the loop, it starts initialization**
  - Needs to perform initialization to acquire an address (AL_PA)
  - Loop port transmits the Loop Initialization Primitive (LIP)
- **Initialization suspends all other activity on the loop while it is being processed**
  - Could require several milliseconds on a large loop
  - When LIP is transmitted, it can corrupt one or more frames requiring recovery after the initialization
- **A loop switch can isolate the LIP and initialization to just the one port**
  - Eliminating the potentially disruptive effects of initialization

Loop Switch Summary

- **Loop switches provide significant enhancements to an arbitrated loop**
  - Minimize latencies on large loop
  - Improve performance
  - Enable multiple concurrent operations
  - Minimize disruptive effect of Loop Initialization
  - Provides better fault isolation and serviceability
- **While adding minimal cost to the loop solution**
  - Loop switches are about $4 per port
  - Will usually be embedded in disk enclosures and replace the Port Bypass Circuit (PBC)
- **Can expect that almost all loop configurations will be switched in the future**

End of “Arbitrated Loop Switches”