



Fault Tolerant  
Controller  
Configurations



# Fault Tolerant Controller Configurations for SCSI

---

**Steve Sicola**  
**Digital Equipment Corporation**




# Fault Tolerant Controller Configurations For SCSI

---

## ■ Overview

- Common Fault Tolerant Controller Configuration for SCSI Today
- Features and Drawbacks of Today's Configurations
- What Metrics of Fault Tolerance are Important to Customers?
- Why Standardize Fault Tolerant Controller Configurations? How?
- Options for Standardization & Conclusions



# Common Fault Tolerant Controller Configurations For SCSI Today


---

- Many System Customers today want high availability with their storage
- Controllers are part of the system, and must be made highly available using redundancy techniques
- Redudancy Techniques include:
  - Active + Hot Standby...
  - Active + Active...


Fault Tolerant  
Controller  
Configurations

# Common Fault Tolerant Controller Configurations For SCSI Today

---




■ A typical Fault Tolerant Controller  
Configuration:



A typical Dual Controller Configuration with two Controllers sharing access to attached devices, in this case with 3 buses being shared between the controllers.

This definition does NOT specify how the controllers are attached to the host or to the devices with parallel SCSI buses




# Common Fault Tolerant Controller Configurations For SCSI Today

---


- **Fault Tolerant Controllers add new details to access of host data by:**
  - **allowing varying attachment to the host computers**
    - **with more than one host port per controller**
    - **with controllers attached on different host buses**

# Common Fault Tolerant Controller Configurations For SCSI Today

---



- **Fault Tolerant Controllers add new details to access of host data by:**
  - **Allowing various types of access to devices**
    - **by the shared topology to the devices between the controllers**
    - **by use of protocol policies such as Reserve & Release in SCSI.**




# Common Fault Tolerant Controller Configurations For SCSI Today

---

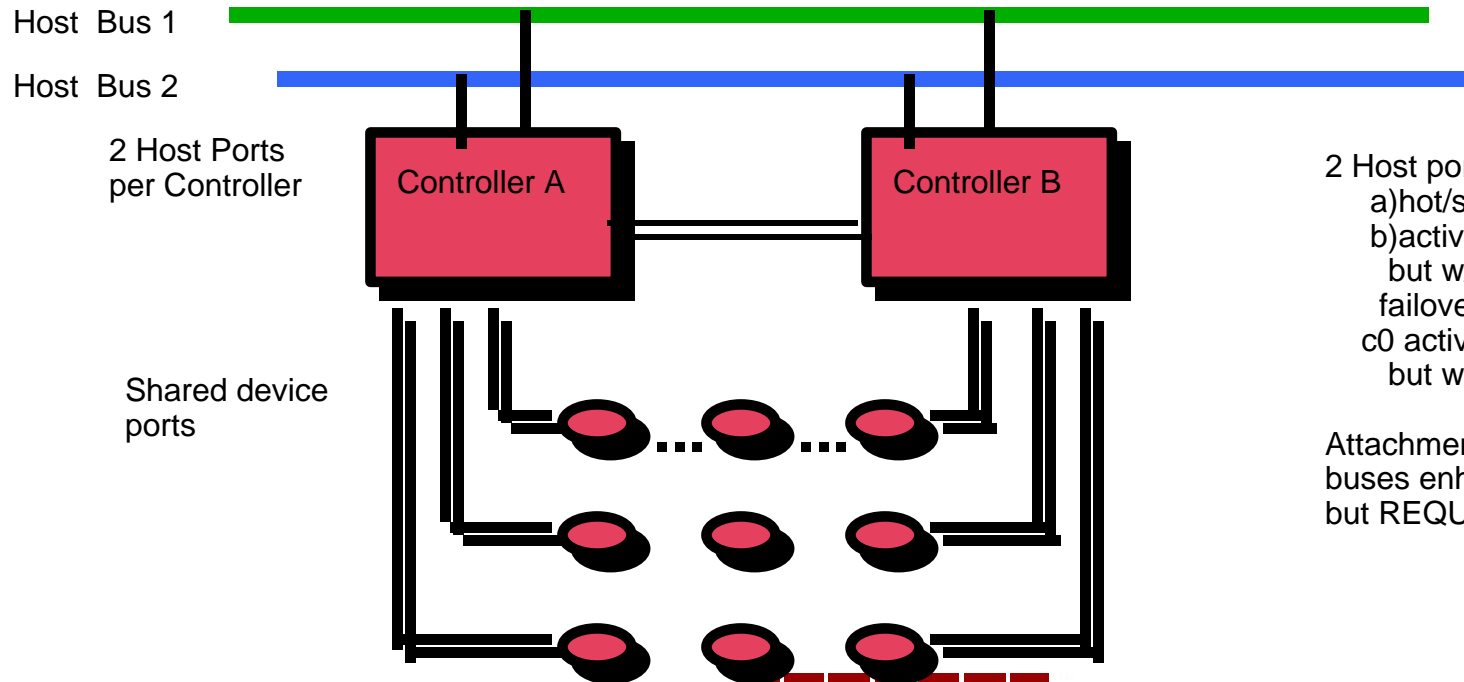
- **Fault Tolerant Controllers add new details to access of host data by:**
  - **Allowing Failover of device control from a failed controller to the survivor**
    - **without host operating system support**
    - **with host operating system support**
  - **Allowing Failback with and without host operating system support.**

Fault Tolerant  
Controller  
Configurations



# Common Fault Tolerant Controller Configurations For SCSI Today


- A block diagram of a controller incorporating all possible attributes:



- 2 Host ports allows for:
- a) hot/standby controller w/o host support
  - b) active/active w/o host support but w/one host port idle until failover.
  - c) active/active w/ host support but with all ports active

Attachment two to two different host buses enhances system availability but **REQUIRES** host support






# Common Fault Tolerant Controller Configurations For SCSI Today

---

- Current solutions in industry for the typical configurations shown include:
  - Vendor Unique host operating system support for single port per controller and/or multi-host bus attachment (Failover can be manual or automatic)
  - Target Failover with 2 host ports per controller (2nd idle until failover). Failback can be automatic or manual.
  - Target Failover with one host port, assuming hardware can act as >1 target.



# Features and Drawbacks of Today's Configurations


---

## ■ Features:

- Highly Available Storage achieved
- Single host port controllers with or without host support reduce product costs
- Multi-host port controllers without host OS support achieve redundancy more simply.

# Features and Drawbacks of Today's Configurations

---




## ■ Drawbacks:

- Non-Standard Implementations leading to possible interoperability problems on different platforms in the same system.
- Cost
- Failover Time is not bounded by several techniques used.
- Some techniques used don't adequately abort I/O's on failed controllers.

# What Metrics of Fault Tolerance are Important to Customers?


---




- **Active-Active Configurations to achieve greater availability with their overall system/cluster, as well as performance of two controllers accessing shared data.**
- **Failover/Failback time is critical in production environments.**
- **They want it for free...**

# Why Standardize Fault Tolerant Configurations in SCSI?

---




- To Reduce the cost of Solutions
- Because of the convergence of standards on many different interconnects that will have SCSI RAID controllers (FibreChannel, SSA, FAST 20 SCSI) coupled with customer desires of fault tolerance & interoperability with all their system boxes/software.



# Why Standardize Fault Tolerant Configurations in SCSI? How?

---


- To Generalize the number of controllers in a fault tolerance configuration as well as the type of access allowed by hosts
- Increase speed of failover from current techniques
- Generalize the possible system topologies supported in simple terms



# Why Standardize Fault Tolerant Configurations in SCSI? How?

---

- HOW:
- Support in SCSI for ‘Registration’ of multi-controller configurations sharing device access. This would be contained within ‘conspiring’ controllers and readable by all attached hosts.
- Support in SCSI for ‘Supported Behaviors’ of Controllers during failover/failback to meet the needs of different system topologies and operations



# Options for Standardization


---

- **Addition of Mode Page definition(s) in the Standard Controller Commands document.**
  - To ‘register’ all controllers in a fault tolerant configuration as well as LUNs being served by each controller
  - TO ‘register’ behavior or controllers during failover/failback
- **Additional ASC/ASCQ’s for Failover/Failback Events**




# Options for Standardization

---



- A combination of a new SCC mode page and additions to the Exception handling page(s) in SCSI-3.
- Inquiry data additions for Controllers to give the 'registration' information (yuk).



# Conclusions

---

- **Controllers in Fault Tolerant Configurations are becoming widespread**
- **Many solutions to support for these configurations exist**
- **Standardization would generalize the problem and solution as well as give Operating Systems MUCH more control and visibility of attached storage subsystems.**