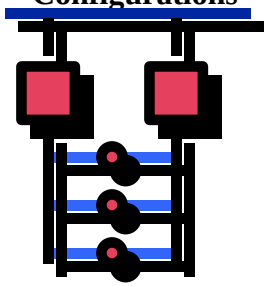

# Proposed Controller Failover Profile

## High Availability Study Group
## X3T10: 95-312r0

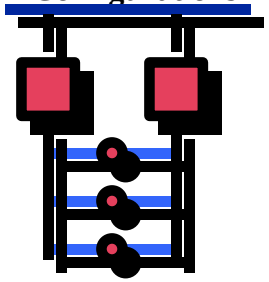

# Proposed Controller Failover Profile

- Overview
  - Purpose
  - Architectural Concepts
  - The Problem
  - Assumptions about Fault Tolerant Controller Configurations & Failover
    - Required Changes for SCC & SCSI-3
  - Functional Description of FT Controller Configuration Usage
  - Additional SCSI-3 Requirements for more Flexibility

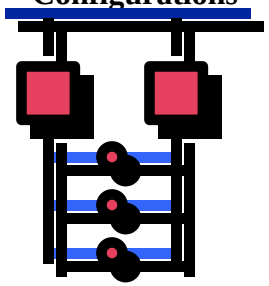

# Purpose for Profile

- Generalize the concept of fault tolerant controller configurations
  - 1 controller logically with many ports
- Generalize usage of SCC to aid OS driver development across all industry platforms for various fault tolerant configuration types
  - standardize setup and/or registration of controllers in FT configuration (with naming independent of serial #'s)
  - standardize reporting of failing controller/returning controller events
  - one port to n-port controller boards
  - any number of controller boards in configuration

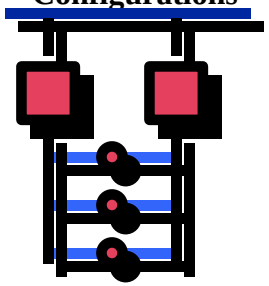

# Architectural Concepts

- FT Controller Configuration definition:
  - Any two or more control units sharing access paths to an arbitrary set of devices/Luns
  - Control Units may be active simultaneously or in some kind of active-standby mode
  - Differing LUN Access models for hosts

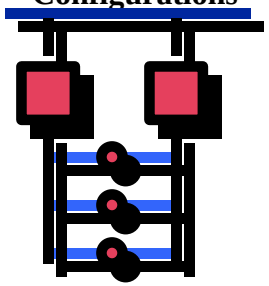

# Architectural Concepts

- FT Controller Configuration Definition:
  - Failover between controllers
  - Failback between controllers
  - Failover/Failback by controllers automatically or under host control
  - Failover Failback notification direct or indirect (message or timeout)

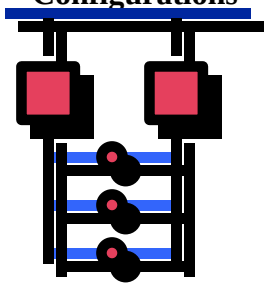

# The Problem

- Non-standard Configuration setup and reporting
- Non-standard failover/failback detections & options for speed/simplicity
- Interoperability with different OS burdened
- Need for standard in open system networked storage environment

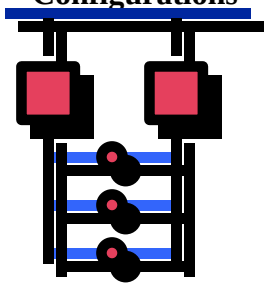

# Assumptions & Changes

- Assumptions are for adherence to SCC models of SACL's
- Assumptions are:
  - Two or more controllers sharing access paths to storage devices
  - The controllers configured with devices logically represent 'one controller with n ports to host'
  - The controllers configured with devices report the same configuration between them

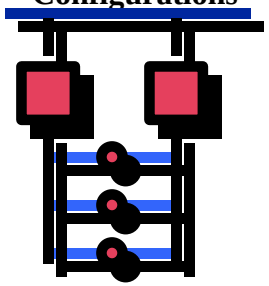

# Assumptions and Changes

- **Assumptions cont;**
  - **Controller communicate with eachother directly (comm) or indirectly (through shared storage) or**
  - **Controller components may have single or multiple host ports & single or multiple device interfaces**
  - **Controllers may be pre-configured or be configured by hosts. Configs verified during controller/host init as well as after initial config**
  - **Any/all surviving controllers within configuration can resume service of storage to host after controller failure.**

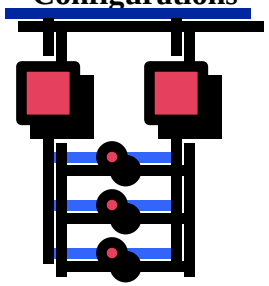

# SCC & Fault Tolerant Controller Assumptions

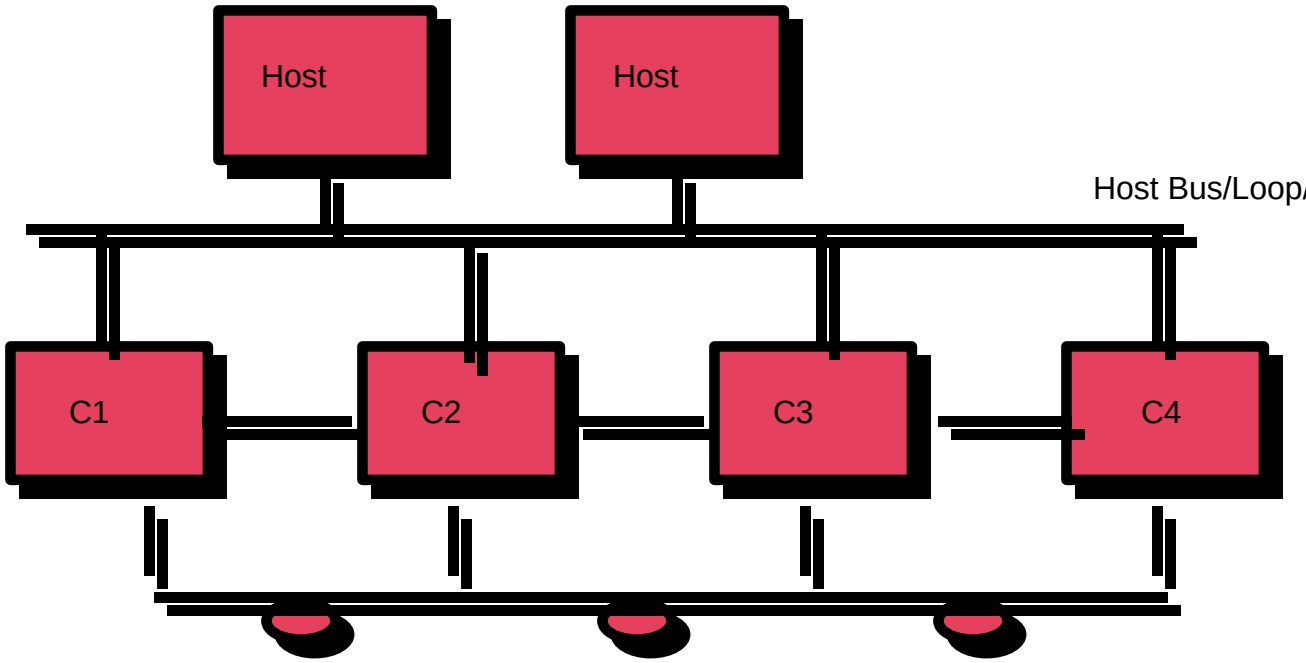


Assumptions
1. All hosts can access all controller
2. All controllers share access paths to storage
3. All controllers can communicate with all other controllers in configuration via shared storage bus(es) or private comm intrfce

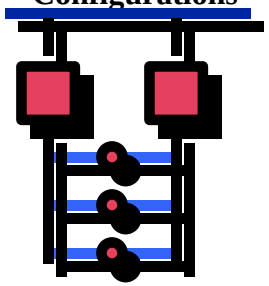

# SCC & SCSI-3 Changes

- SCC changes involve some specific changes to the ATTACH to COMPONENT DEVICE and REPORT COMPONENT DEVICE ATTACHMENTS service actions
  - Changes to Attach involve LUN_C=0 denoting service action for controller attachments
    - Changes to parameter list based on LUN_C=0 for list to refer to controllers to be attached to controller receiving service action.

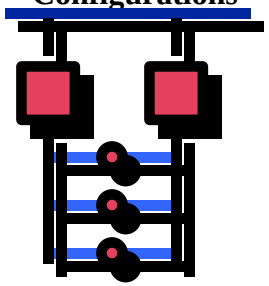

# SCC & SCSI-3 Changes

- Changes to the Report COMPONENT DEVICE ATTACHMENT command
  - LUN_C=0 denotes controller service action to report about controller attachments
  - Response contains information about all current attachments, the name of the attachment, and information controllers eligible to become attached.

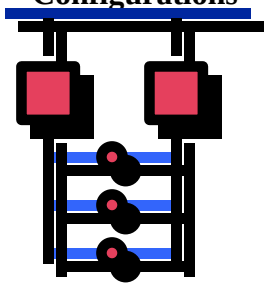

# SCSI-3 Changes

- New ASC/ASCq's for Fault Detection in FT controller configuration:
  - FAILOVER
  - FAILBACK
  - sent to hosts by detecting controller(s) of failed controller. Method determined by SCSI-3 exception handling methods (AEN, Unit Attention, etc..)

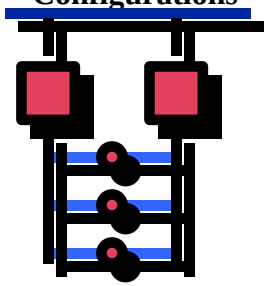

# FT Controller Configurations

- Using these changes, hosts can
  - Configure fault tolerant controller configurations
  - Efficient configuration checks by hosts (top level controller checks, followed by One scan down through a controller to verify LUN/device configuration
  - Failover/Failback much more quickly since controllers can detect partner failures faster
  - Identification of Load balancing opportunities
  - Consistent and Complimentary to Persistent Reserve & Global device/LUN IDs

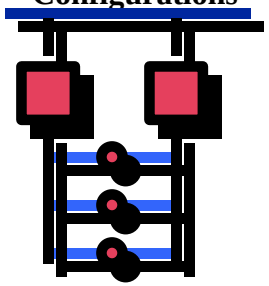

# Additional SCC Requirements for Configuration Flexibility

- For more Flexible Configuration, controller configuration capabilities to support multiple configurations of LUNs between attached controllers is possible/desireable
  - With networked storage
  - With serial storage & high connectivity
- Capabilities should be reported and controllable

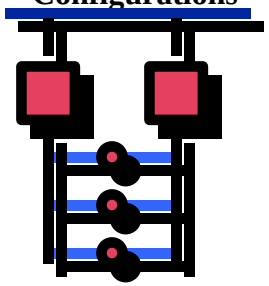

# Fault Tolerant Configurations

- ‘N’ Controller component configurations may want/need different LUN access models.
  - Total sharing of all LUNs configured between all controllers attached to eachother.
    - This is represented by the profile as it stands today with the proposed SCC changes
    - This ties a set of controllers to all LUNs configured from any/all controllers in the attached configuration. Class 1 configuration
    - Other devices may share access but comprise LUNs for different controller attachments.

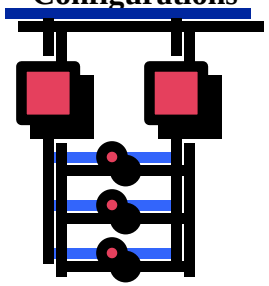

# Fault Tolerant Configurations

- The Controller Attachment and Controller to LUN attachments allow for
  - Ease of configuration: less OS polling of all targets and LUNs to build configuration maps with one or two level controller configurations
    - class 1 configuration (one level)
    - class 2 configuration (two level)

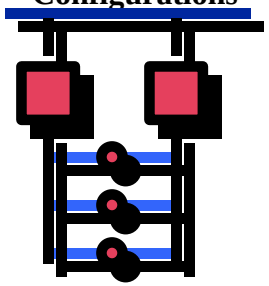

# Fault Tolerant Configurations

- Controller Attachments and Controller to LUN Attachments also provide for
  - Easier use and management of Global IDs for devices and LUNs within a controller configuration.

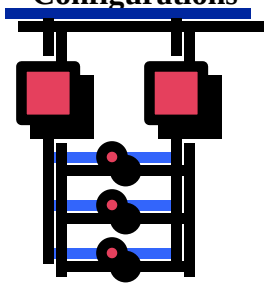

# Fault Tolerant Configurations

- LUNs attached to specific sets of controllers within an attached controller configuration
  - This requires an Attachment of Controller to specific LUN. It also implies a set of flags for reporting of and control of the LUN access method to be employed by the controller(s) attached to the LUN.

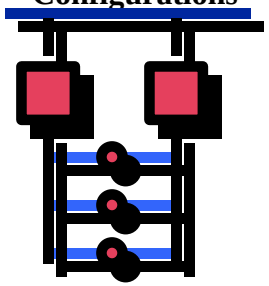

# Fault Tolerant Configuations

- The LUN attachment method is basically a ‘Sub-component attachment between controllers to sets of LUNs on shared access paths.  Class 2 configuration.
  - A Create Controller-LUN Attachment command would need to be added
    - Global ID Assignment
    - Setup Class of LUN Service

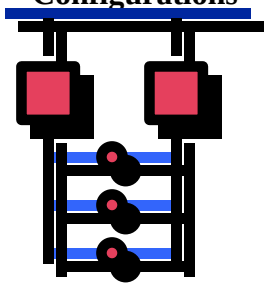

# LUN Access/Service Methods

- LUN access with controllers may be:
  - Allowed by only one controller, that controller only responds to read/write commands
  - Allowed by one controller at a time, but requires interlock commands to bind/unbind from a controller (i.e. Reserve/Release)
  - Allowed by both controller simultaneously, assuming a high level of interlock on LUN accesses

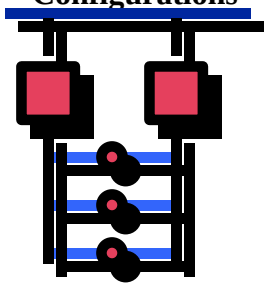

# Fault Tolerant Configurations

- A Report Controller-LUN Attachment Command will also be required
  - To Report Attached LUNs
  - To Report Eligible LUNs
- The Report Component Device Attachment (for Controllers, LUN_C=0) needs
  - To Report Class of LUN access/service Allowed
  - Controller Configuration Type (1,2,other)