

iSCSI testing: What are the test Challenges Under the Hood of a 10 Gb iSCSI Storage Product Certification?

Dr. M. K. Jibbe

Distinguished Engineer

Manager and Technical Lead of Test Architect and Product Certification in India

LSI Corporation (Engenio Storage Group)

iSCSI RAID Storage Testing

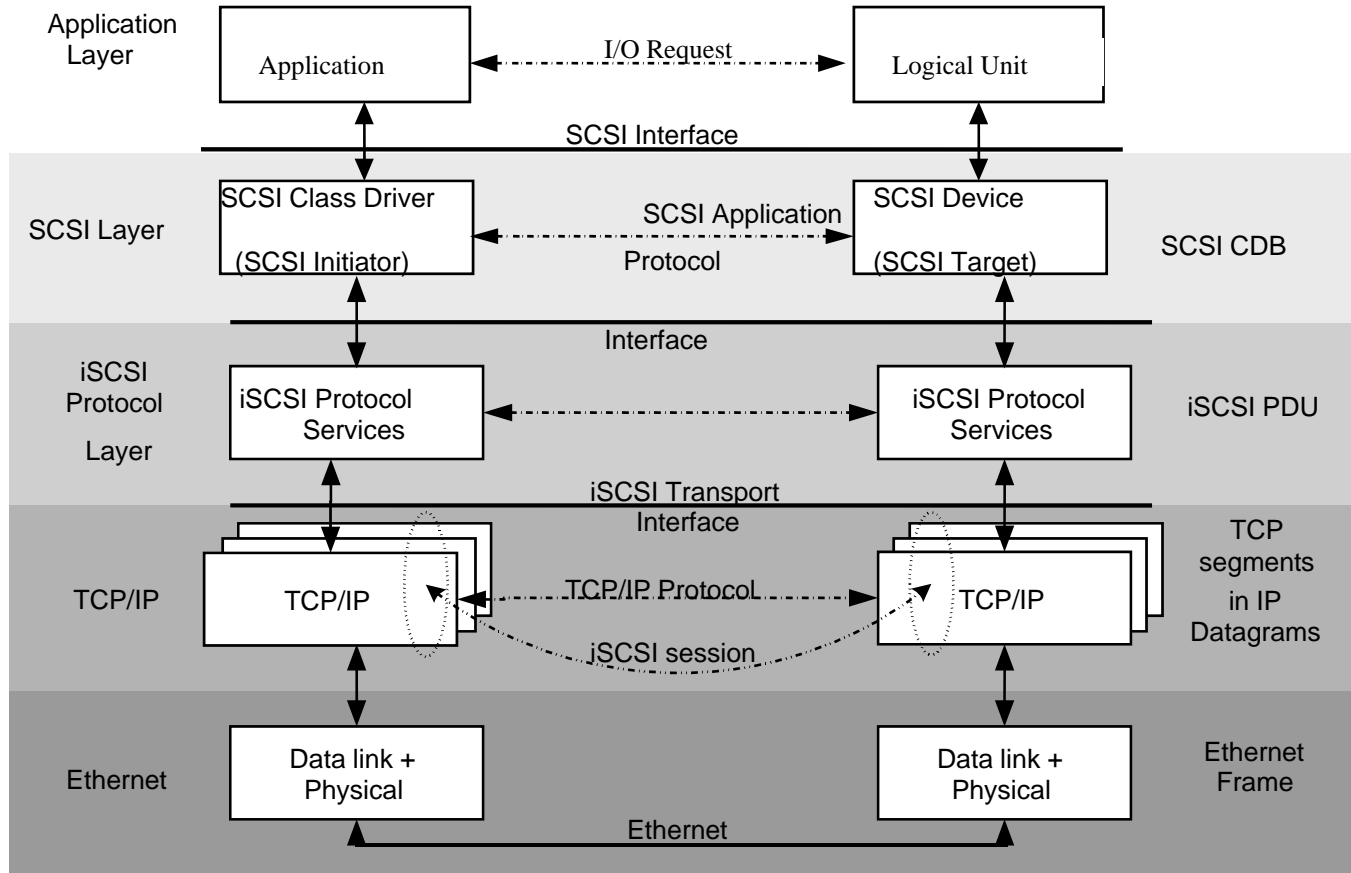
The certification of a 10 Gb iSCSI RAID Storage System elicits a lot of challenges at the development level and the Test / Quality Assurance level. The challenges are due to the fact that a 10 Gb iSCSI is a newly deployed iSCSI host interface in the RAID Storage environment. As a result the size of a development module level test should be designed very carefully to establish a test coverage beyond basic implementation verification, standard RAID testing, or the iSCSI plug fest. These module level tests must tackle the test time windows associated with the following iSCSI characteristics:

- ❑ NIC vs. CNA
- ❑ Device discovery,
- ❑ 10 GB switch traffic control and congestion,
- ❑ Security mechanisms with different Operating systems,
- ❑ Operational parameters associated with I/O retries and recovery
- ❑ Management, Administration, and Integration with Storage products
- ❑ Design For Testability “DFT” mechanisms
- ❑ Diagnostics, problem Isolations
- ❑ IPV4 vs. IPV6

However a number of the module tests above can be leveraged from the certification a 1 Gb iSCSI RAID products.

There are specific features such as backup, snapshot, remote mirroring, and cluster application compatibility that must be supported by the RAID product and must be verified during the testing of the RAID controller host interface.

What is iSCSI?



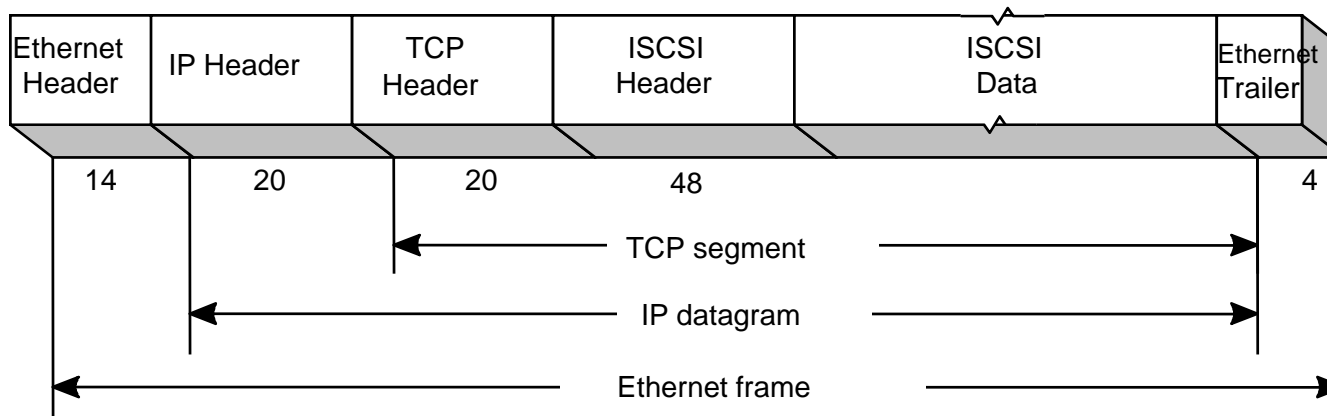
- Replaces shared bus with switched fabric
- Transparently encapsulates SCSI CDBs
- Unlimited target and initiator connectivity

□ **iSCSI: Internet Small Computer Systems Interface**

A TCP/IP based protocol for establishing and managing connections between IP-based storage devices, hosts, and clients.

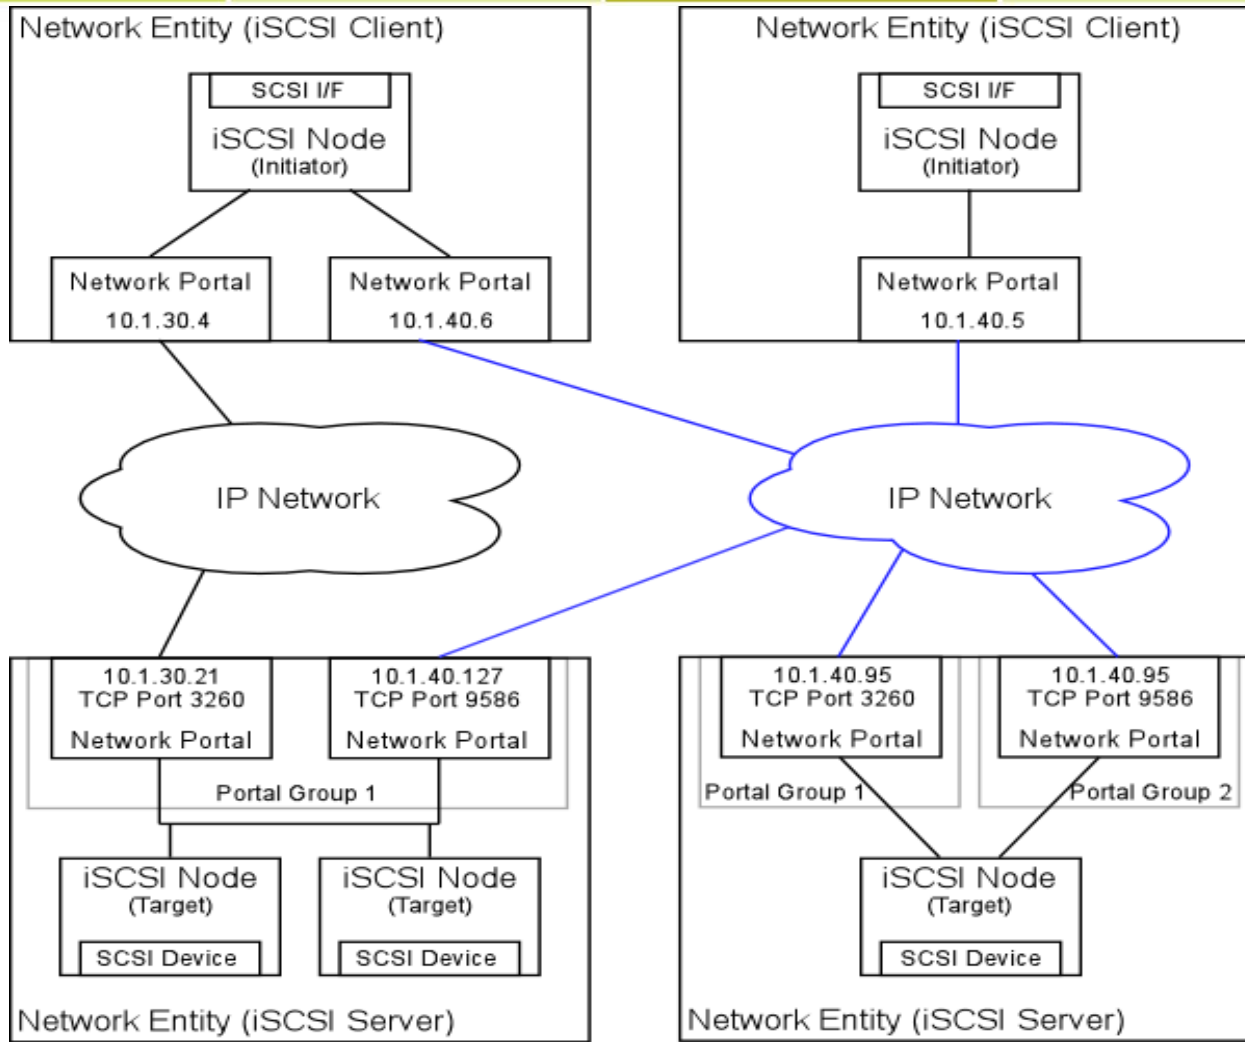
TCP/IP used as a transport for SCSI protocol

Use: Native IP SANs

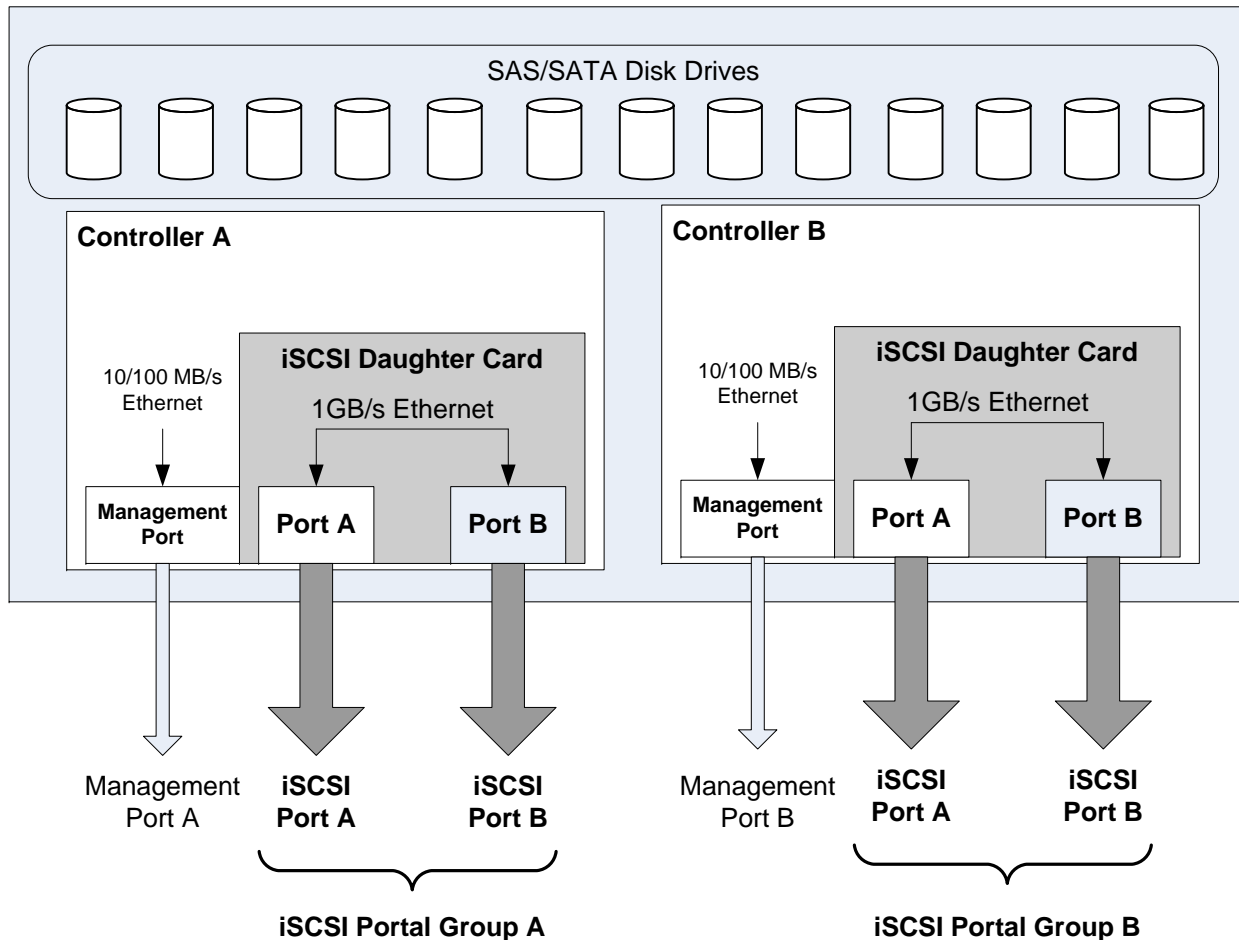


- iSCSI is a **transport protocol for SCSI** that operates on top of TCP through encapsulation of SCSI commands in a TCP/IP stream. Enables the transport of I/O Block data over IP Networks

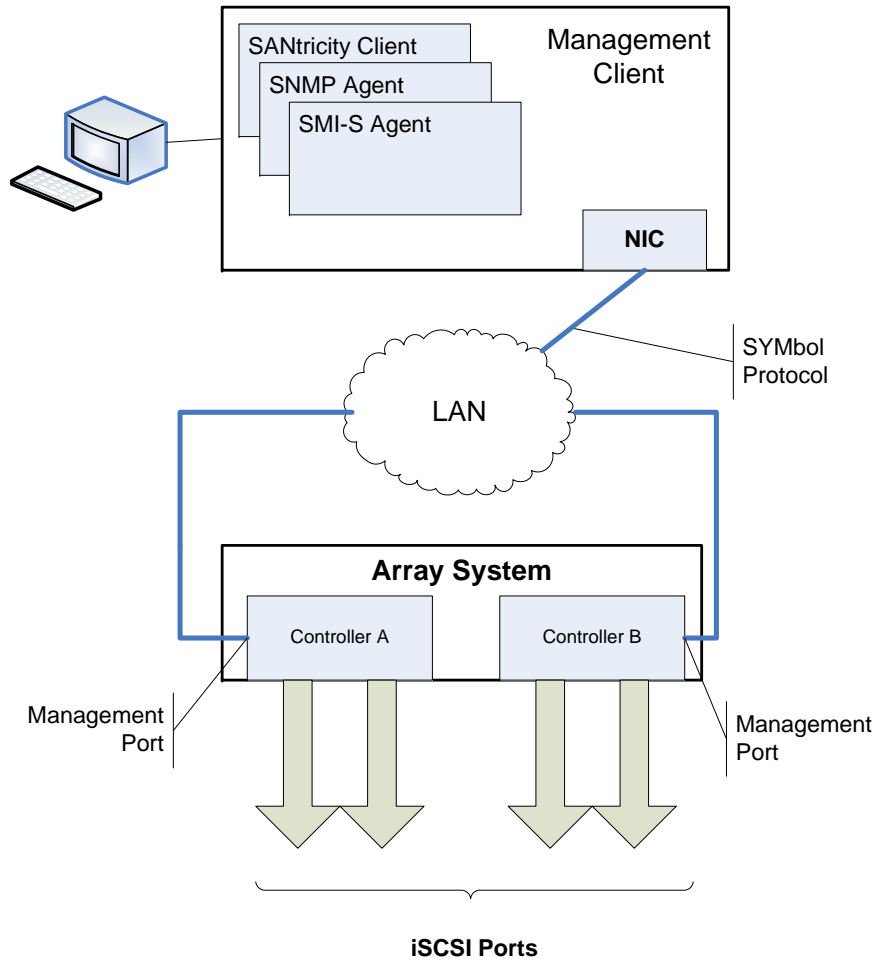
iSCSI Architectural Objects cont'd



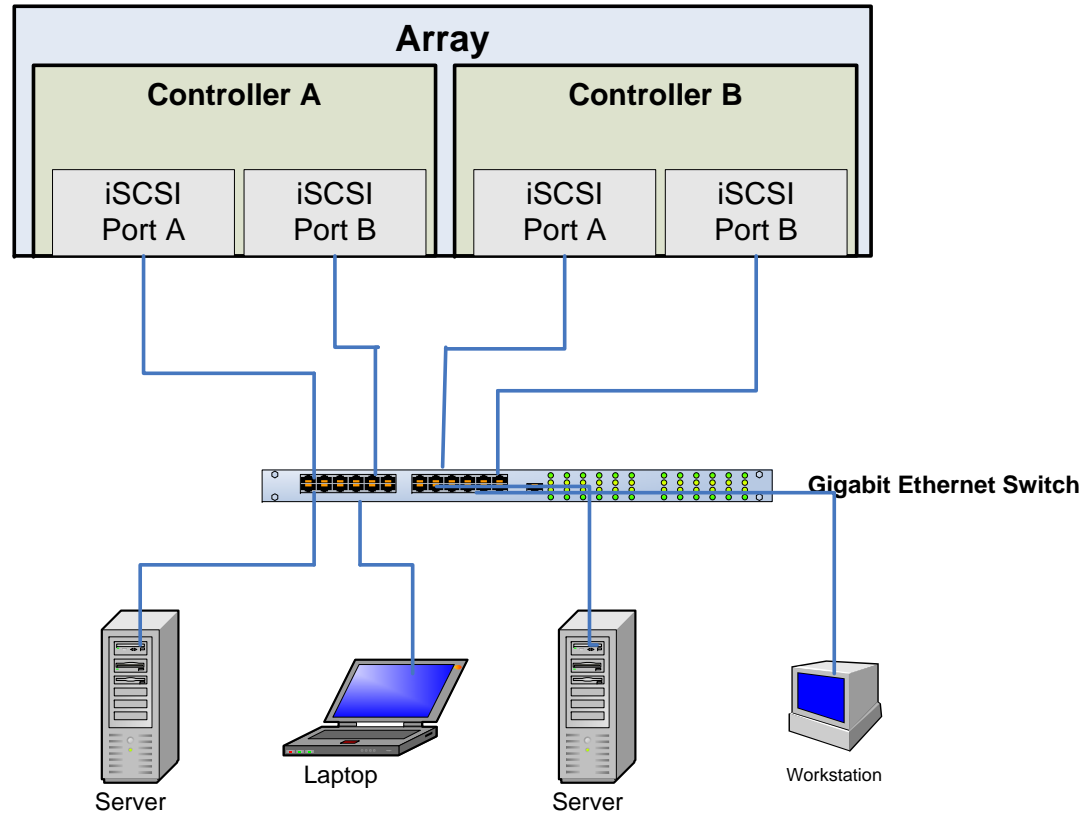
Array Ethernet Port Configuration



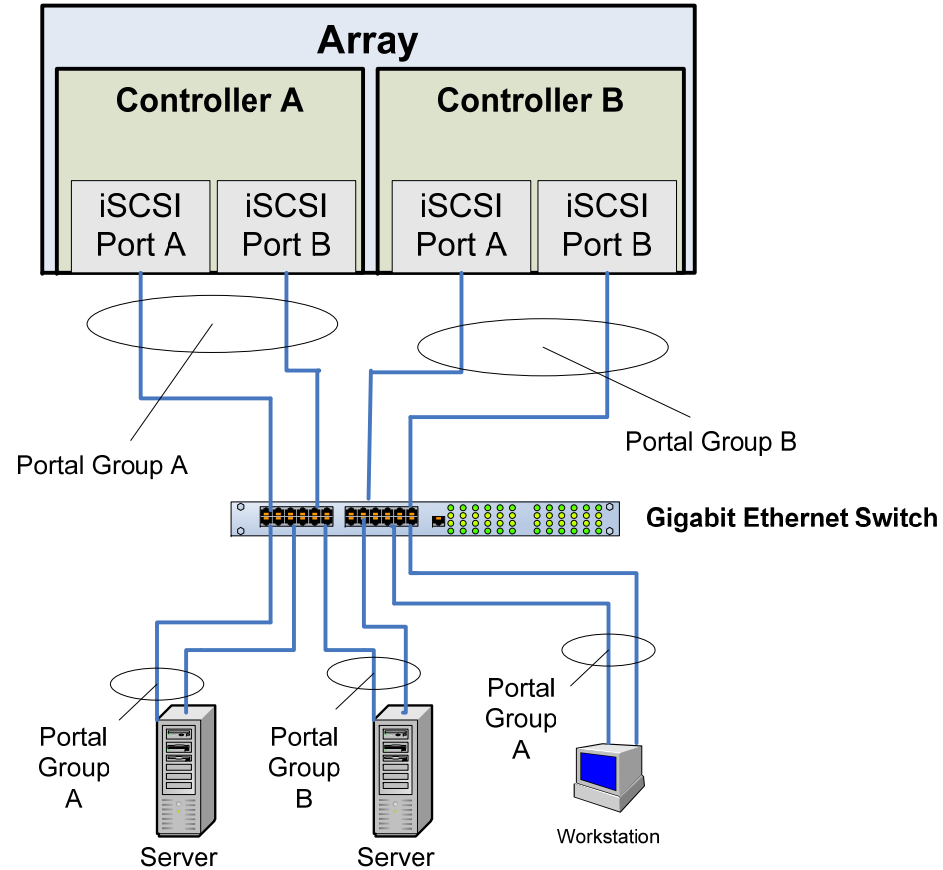
Array Management Configuration



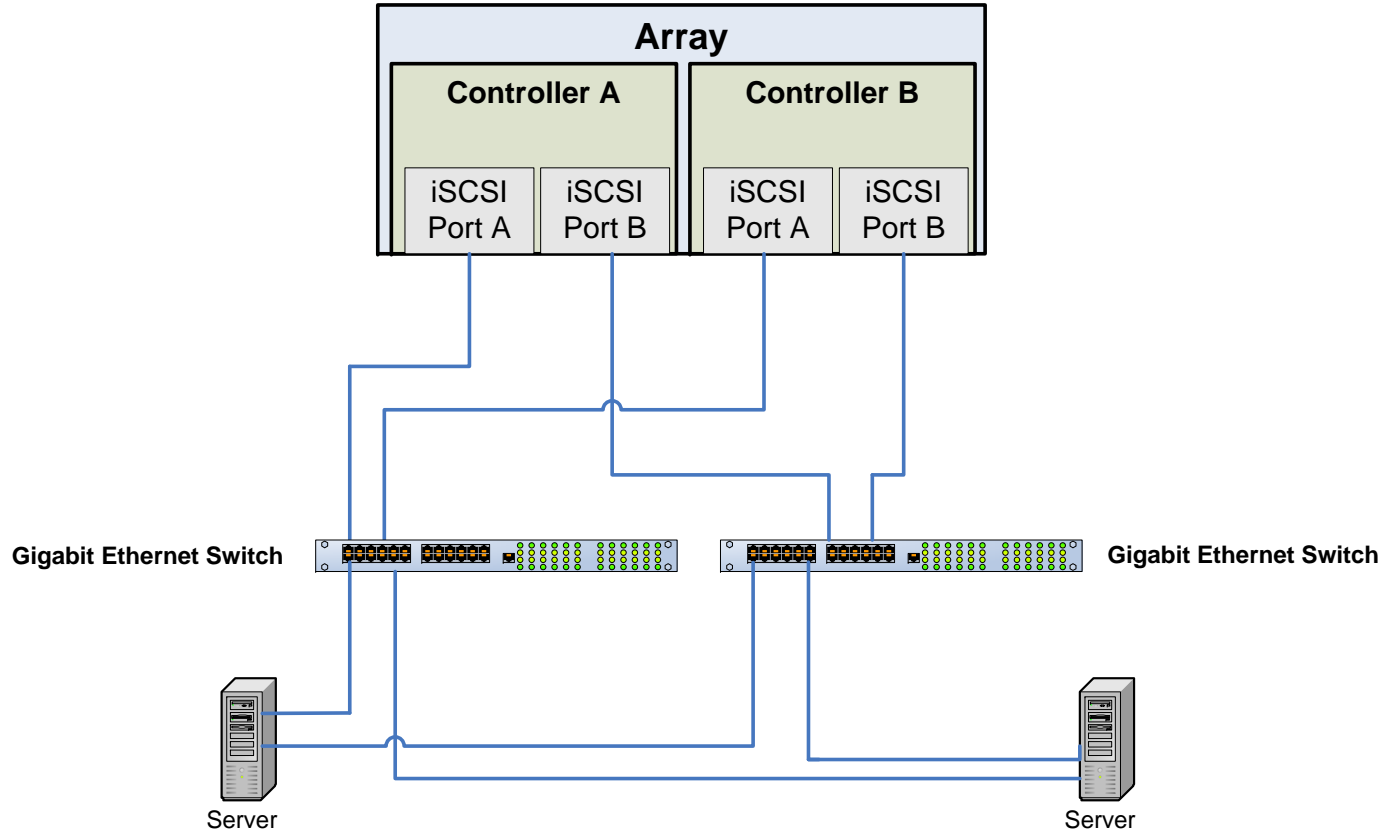
Single Path Topology



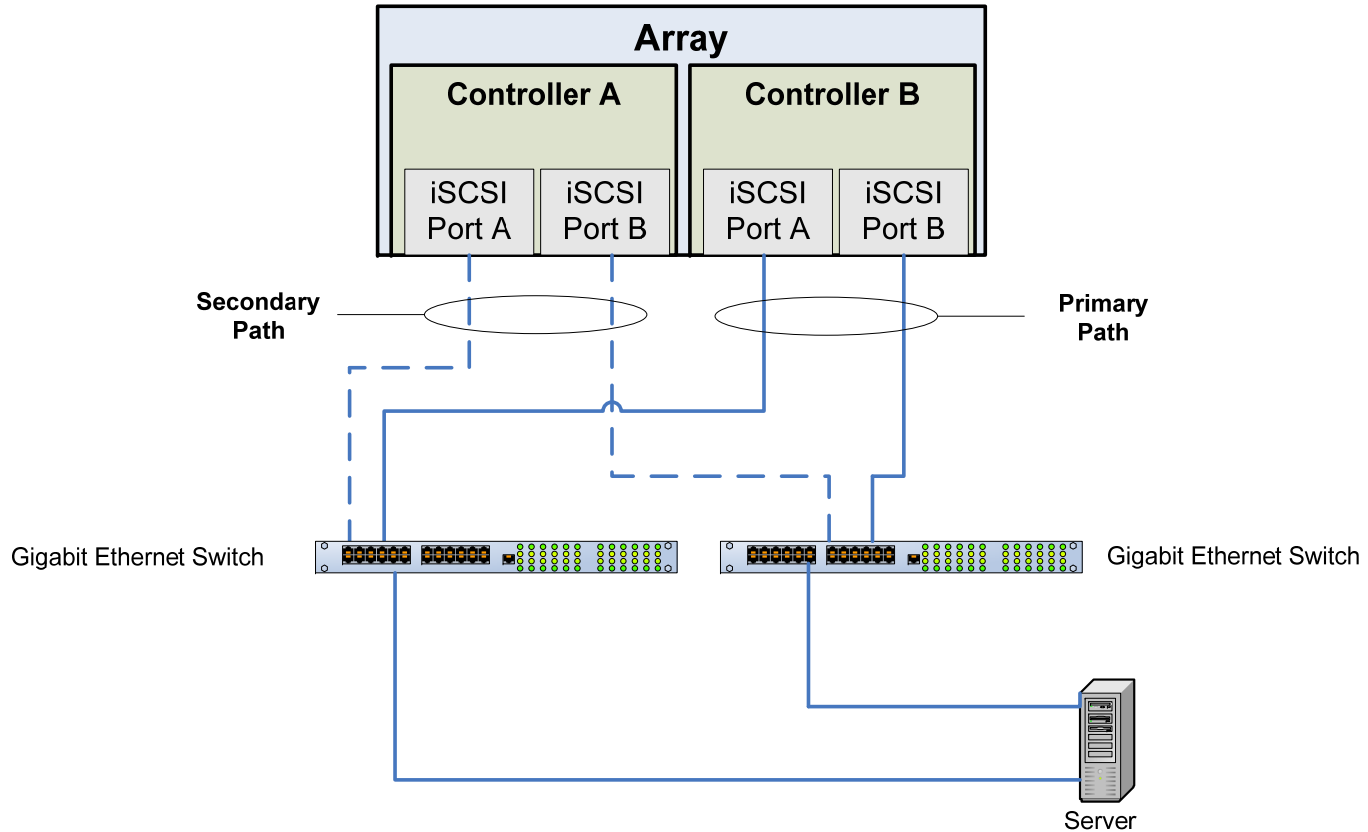
Dual Path Topologies

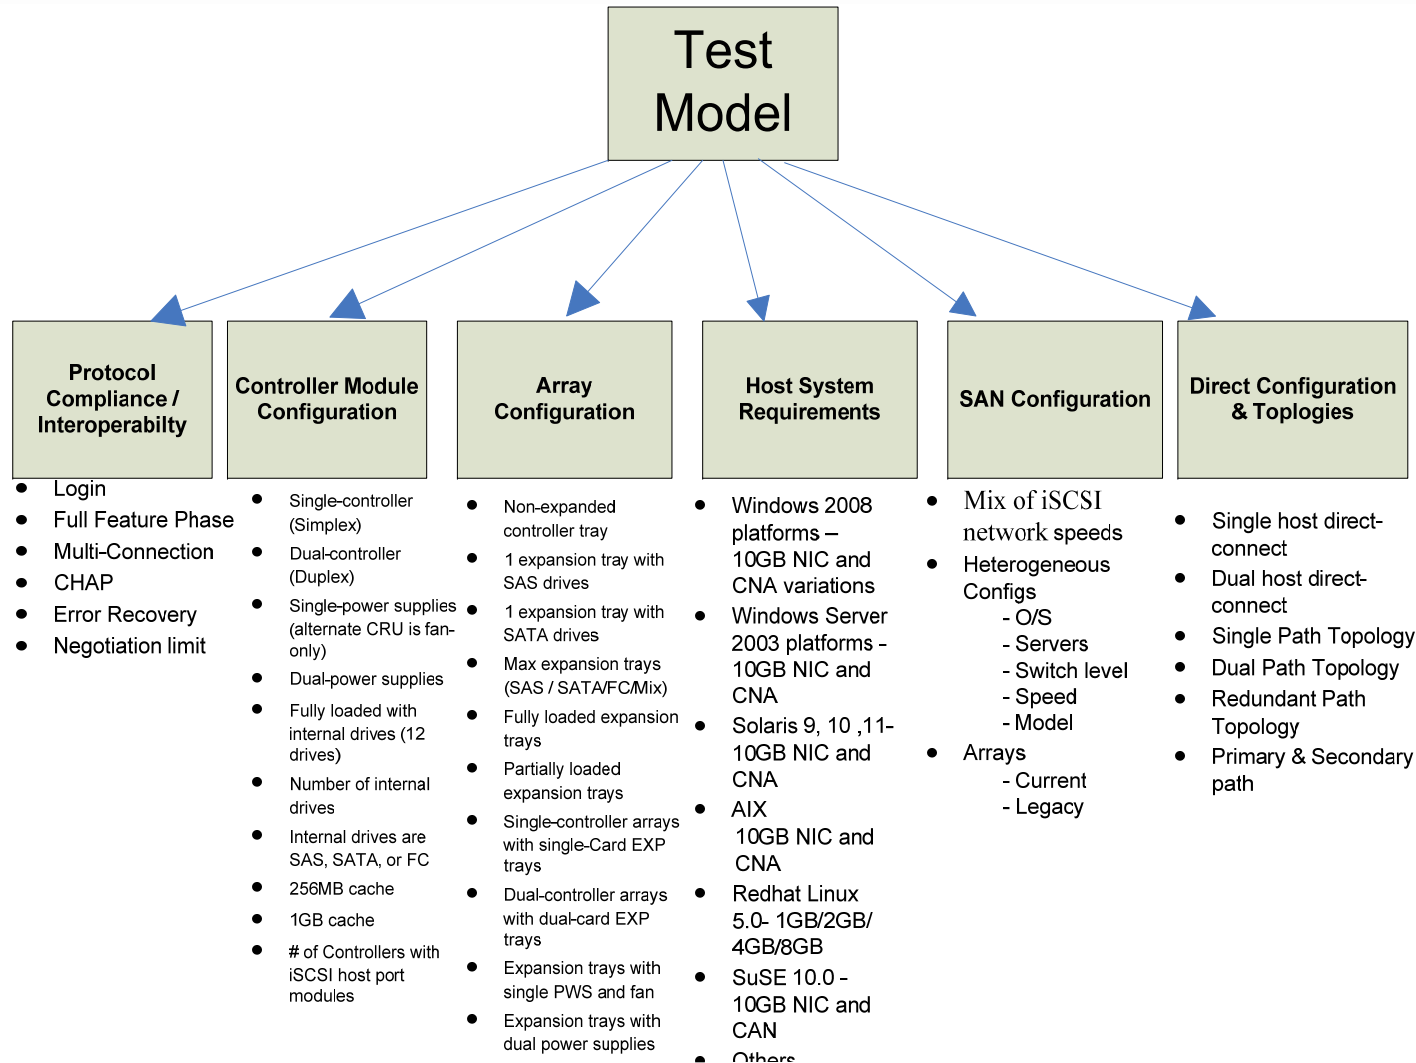


Redundant Dual Path Topologies



Primary and Secondary Paths





❑ Discovery & Initialization Compliance

- ❑ Single CAN /NIC and single array
 - a) Power On
 - b) Enable/ Disable
 - c) Cable pull
 - d) Disconnect / Reconnect target
 - e) Remove / add resources
 - f) Non iSCSI device on the network switch
- ❑ Single Initiator and multiple targets (a, b, c, d, e, & f)
- ❑ Multiple Initiators / Multiple Targets (a, b, c, d, e, & f)

❑ Login Conformance

- ❑ Verify the usage / response of the following parameter during a target login
 - ❑ Standard Login.
 - ❑ CmdSN
 - ❑ Version Active.
 - ❑ T Bit
 - ❑ ExpStatSN
 - ❑ Negotiate Once.
 - ❑ Login Partial Response.
 - ❑ Status Detail
 - ❑ Invalid PDU
 - ❑ Parameter Names
 - ❑ AuthMethod
 - ❑ Header and Data Digest
 - ❑ Header and Data Digest
 - ❑ MaxConnections
 - ❑ TargetAlias
 - ❑ Marker Negotiation
 - ❑ FirstBurstLength
 - ❑ Full Feature Phase Conformance

Reference: UNH IOP

Interoperability / Protocol Compliance Multi-Connection Conformance with 10G NIC and 10 G CNA

- Basic I/O Completion(Read and Write)
- Max Connections
- Target Name
- Initial R2T
- Immediate Data
- MaxRecvData SegmentLength
- MaxBurstLength
- FirstBurstLength
- DefaultTime2Wait
- DefaultTime2Retain
- Connection Terminated
- MaxOutstandingR2T
- DataPDUInOrder
- DataSequenceInOrder
- HeaderDigests and DataDigests
- CMDSN
- DataSN
- StatSN
- R2TSN

Reference: UNH IOP

Interoperability / Protocol Compliance

CHAP Conformance with 10G NIC and 10 G CNA

□ GROUP 1: CHAP_A VERIFICATION

- CHAP_A VALID VALUE
- CHAP_A VALID VALUE IN LIST
- CHAP_A INVALID VALUE
- CHAP_A VALID VALUE NOT IN LIST
- CHAP_A OUT OF ORDER

□ GROUP 2: CHAP_I VERIFICATION

- CHAP_I VALID VALUE
- CHAP_I INVALID VALUE
- CHAP_I NO VALUE
- CHAP_I TOO BIG VALUE
- CHAP_I OUT OF ORDER
- CHAP_I REUSED ON SECOND CONNECTION
- CHAP_I DIFFERENT ON SECOND CONNECTION
- CHAP_I REFLECTED
- CHAP_I REFLECTED ON SECOND CONNECTION

□ GROUP 3: CHAP_C VERIFICATION

- CHAP_C REUSED
- CHAP_C BIG VALUE
- CHAP_C SMALL VALUE
- CHAP_C TOO BIG VALUE
- CHAP_C OUT OF ORDER
- CHAP_C RECEIVE REUSED
- CHAP_C REFLECTED
- CHAP_C REFLECTED ON SECOND CONNECTION
- CHAP_C NEW ON SECOND CONNECTION

□ GROUP 4: CHAP_N VERIFICATION

- CHAP_N INVALID
- CHAP_N
- CHAP_N SMALL
- CHAP_N TOO BIG
- CHAP_N OUT OF ORDER
- CHAP_N IDENTICAL
- CHAP_N REFLECT
- CHAP_N DIFFERENT NAME

□ GROUP 5: CHAP_R VERIFICATION

- CHAP_R INVALID VALUE
- CHAP_R TOO BIG
- CHAP_R TOO SMALL
- CHAP_R OUT OF ORDER

Reference: UNH IOP

Interoperability / Protocol Compliance Error Recovery Conformance with I0G NIC and I0 G CNA

- Retry Advertent
- Retry After Digest Error
- Allegiance Reassignment
- R2T Snack Support
- Data Snack Support
- Status Snack Support
- Resegmentation SNACK Support
- Usage of Reject CMD PDU
- Termination of tasks
- Format Errors
- Header Digest Error
- Out of order DataSN
- Protocol Error
- Drop Immediate CMD
- Drop Non-Immediate CMD
- Drop Solicited Data-out
- Drop Data-In
- Drop Text Response, Request
- Drop NOP-In & Out
- Data Digest Error on non Immediate Data
- Data Digest Error on Immediate Data
- Data Digest Error on Unsolicited Data F=0
- Data Digest Error on Unsolicited Data F=I
- Data Digest Error on solicited Data F=0
- Data Digest Error on solicited Data F=I
- Data Digest Error on Data-In F=0
- Data Digest Error on Data-In F=I
- Data Digest Error on NOP-In
- Data Digest Error on Immediate NOP-In
- Data Digest Error on NOP-Out
- Data Digest Error on Immediate NOP-Out
- Data Digest Error on Text Request
- Data Digest Error on Immediate Test Request
- Data Digest Error on Text Response
- Connection Reinstatement

Reference: UNH IOP

- ❑ Recognition of drive pulls/pushes, internal/expansion, FC/SAS/SATA
 - ❑ unassigned drives
 - ❑ assigned drives
 - ❑ GHS spare drives
 - ❑ GHS in-use drives
- ❑ Proper handling of drive failures, internal/expansion, FC/SAS/SATA
 - ❑ unassigned drives
 - ❑ assigned drives
 - ❑ GHS spare drives
 - ❑ GHS in-use drives
- ❑ Diagnostic reporting
 - ❑ host-side
 - ❑ drive-side
- ❑ Error Handling
 - ❑ Recovery guru spot-check
 - ❑ Recovery guru in-depth analysis of reports which previously were FC-centric
- ❑ For dual back-end systems, back-end failure of a single FC/SATA/SAS channel
- ❑ Volume rebalancing
- ❑ I/O shipping
 - ❑ proper operation
 - ❑ rebalancing based on volume geometry
- ❑ Proper handling of iSCSI host-side link interruptions
 - ❑ P2P or Link failures, including LIPs, pathblock, ..
 - ❑ Host interface pulls/pushes during I/O
 - ❑ HBA pulls/pushes during I/O
 - ❑ Switch failures
- ❑ Spot-check volume/host mapping, esp. limitations
- ❑ Different host speeds on the various host ports
- ❑ Spot-check snapshot
- ❑ Spot-check volume copy
- ❑ LED proper behavior, including any new ready-to-remove indicators
- ❑ Proper operation of, detection of, etc.
 - ❑ Fans
 - ❑ Power Supplies
 - ❑ Other managed components (interfaces, whatever)
- ❑ Minimal Major Event Logs behavior checking
- ❑ Tray ID Behavior, conflict checking/reporting
- ❑ Upgrade/Downgrade between simplex/duplex
 - ❑ Implies conversion between single/dual tray card expansion trays
- ❑ Controller firmware download, staged download
- ❑ Drive firmware download
- ❑ Tray firmware download
- ❑ CLI support - primarily to ensure nothing was broken by introduction of FC/SATA/SAS support at the controller level

RAID System / Stress tests with 10G SDC and 10G CNA

STORAGE DEVELOPER CONFERENCE

SNIA ■ SANTA CLARA, 2009

- ❑ System/stress testing (all with media scan enabled)
- ❑ I/O with controller Reboot/failure/drive failure & iSCSI Switch Reboot/ Port Disable/Enable
 - ❑ vanilla
 - ❑ degraded volumes
 - ❑ during reconstruction
 - ❑ during copy-back
 - ❑ during CFW download
 - ❑ during volume configurations
 - ❑ with snapshots in play
 - ❑ with volume copy in play
 - ❑ with RVM in play
 - ❑ with short-run reconfiguration in play
 - ❑ Short/medium/long distances (Delay Simulator)
- ❑ Volume migration
- ❑ Tray migration
- ❑ Excessive reconfigurations
- ❑ Excessive stress tests
- ❑ Large configurations
- ❑ Host Management software
- ❑ Host Context Agent
- ❑ Support bundle limitations

iSCSI Impacts on RAID Functionalities with **SDC**

STORAGE DEVELOPER CONFERENCE

SNIA ■ SANTA CLARA, 2009

- ❑ RAID Resource Discovery
- ❑ RAID Login
- ❑ RAID Configuration, Management, and SAN
- ❑ CMD Exchange
- ❑ Data Exchange
- ❑ Status Exchange
- ❑ Exception Handling
 - ❑ Time out
 - ❑ Failover
 - ❑ Recovery
- ❑ RAID Premium Features
- ❑ RAID Performance
- ❑ RAID Applications

iSCSI RAID System Connections and Sessions Testing

- ❑ iSCSI Connection with 10G CNA:
 - ❑ Verify a TCP connection over which the initiator and target communicate via iSCSI PDUs
 - ❑ Verify uniquely identified in a session by an initiator defined connection ID (CID)
 - ❑ Verify the response and any data associated with an iSCSI command must be returned on the same connection
- ❑ iSCSI Session with 10G CNA:
 - ❑ Verify a set of iSCSI connections that link an iSCSI initiator and target
 - ❑ Verify uniquely identified by a 64 bit Session ID (SID) built from a 48 bit initiator defined Initiator Session ID (ISID) and a 16 bit target defined Target Session Identifying Handle (TSIH)
 - ❑ Verify resources of a target (i.e., LUNs) must be identical across all connections that make up a session
 - ❑ Verify commands can be alternated across all connections in a session for bandwidth aggregation
 - ❑ Verify error recovery connections can be created on the same network portal as a failed connection

iSCSI RAID System Names Testing with ISDC

STORAGE DEVELOPER CONFERENCE

SNIA ■ SANTA CLARA, 2009

- iSCSI Name:
 - Verify that the Identified iSCSI node and its encapsulated SCSI device
 - Verify the Usage in authentication of targets to initiators
 - Verify it is world wide unique
 - Verify the Utilization of existing naming authorities
 - Verify human readable 233 character name
- iSCSI Alias:
 - Verify user assigned name
 - Verify that it does not need to be unique
 - Verify its exchanged during login but not used by iSCSI protocol
 - Verify human readable 255 character name

iSCSI RAID System Sequence Numbers Testing with 10G CNA

- ❑ Verify iSCSI uses three sequence numbers:
Command, Status, and Data
- ❑ Verify sequence numbers support:
 - ❑ Detection of lost packets
 - ❑ Ordered command and data delivery
 - ❑ iSCSI layer flow control
- ❑ Verify sequence number format:
 - ❑ Sequence numbers are implemented as 32 bit fields in the iSCSI PDU headers
 - ❑ They form a continually increasing sequence
 - ❑ Computations on sequence numbers are defined in RFC1982 – Serial Number Arithmetic
- ❑ Verify exchange Locations:
 - ❑ Sequence numbers are exchanged in iSCSI PDU headers:
 - ❑ Examples:
 - ❑ SCSI Command PDU: CmdSN and ExpStatSN
 - ❑ SCSI Response PDU: StatSN, ExpCmdSN, and MaxCmdSN
 - ❑ Data-In PDU: StatSN, ExpCmdSN, MaxCmdSN, and DataSN

- Verify Command sequence number variables:
 - CmdSN – current command sequence number
 - ExpCmdSN – next command sequence number expected by target
 - MaxCmdSN – maximum command sequence number that can be accepted by the target
- Verify Command SN characteristics:
 - Session wide scope
 - Target's iSCSI command queuing capacity is:
 - MaxCmdSN – CmdSN
 - Initiator detects lost commands by examining ExpCmdSN in iSCSI response and NOP-In command PDUs received from target
 - All commands except those marked for immediate delivery increment the CmdSN

iSCSI RAID System Status Sequence Numbers Testing with 10G CNA

- ❑ Verify Status Sequence Number Variables:
 - StatSN – current response sequence number
 - ExpStatSN – next response sequence number expected by the initiator
- ❑ Verify Status Sequence Number Characteristics
 - ❑ Connection scope
 - ❑ No queuing capacity is required – there is 1 response for each command and no more
 - ❑ Target detects lost responses by examining ExpStatSN in iSCSI commands PDUs received from target
 - ❑ All responses increment the StatSN

iSCSI RAID System Data Sequence Numbers Testing with 10G CNA

- ❑ Verify Data Sequence Number Variables:
 - DataSN – current data PDU in data transfer
 - R2TSN – current ready to transfer PDU in data transfer
- ❑ Verify Data Sequence Number Characteristics
 - ❑ Data transfer scope
 - ❑ DataSN is generated by initiator for the Data-Out PDU and by target for the Data-In PDU
 - ❑ R2TSN is generated by target for R2T (ready to transmit) PDU
 - ❑ Data sequence numbers are implicitly acknowledged by receipt of the associated response StatusSN
 - ❑ DataSN sequence numbers are explicitly acknowledged on request by a special form of the SNACK PDU
 - ❑ Data sequence numbers acknowledgement is requested by a bit in the Data-In and Data-Out PDU headers
 - ❑ The receiver can detect missing Data PDUs by checking DataSN
 - ❑ Retransmission of missing Data PDUs are requested by SNACK commands

iSCSI RAID System Command Queue Testing with 10G CNA\

- ❑ Verify commands may arrive out of order at target:
 - ❑ Lost TCP packets pending retransmission
 - ❑ Different transmission delays of connections in a session
 - ❑ Commands rejected due to data digest errors
 - ❑ Commands set for immediate delivery
- ❑ Verify the command queue:
 - ❑ Holds commands received out of order until preceding commands arrive
 - ❑ Command queue is session wide
 - ❑ Task management commands may affect commands resident in the command queue or already passed to the SCSI layer

Testing for iSCSI RAID System iSCSI Protocol Optimizations with 10G CNA

- ❑ Verify Immediate Data:
 - ❑ A data transfer sent in the SCSI command PDU
 - ❑ Eliminates need for target to send an R2T PDU and the initiator to send a Data-Out PDU to transfer the data
 - ❑ Session wide setting - negotiated during first session login
 - ❑ Maximum immediate data transfer limited by PDU size
- ❑ Verify Unsolicited Data:
 - ❑ A sequence of one or more Data-Out PDUs immediately sent after the corresponding SCSI command PDU containing a SCSI write command
 - ❑ Eliminates need for target to send R2T before starting transfer
 - ❑ Session wide setting - negotiated during first session login
 - ❑ Maximum unsolicited data length negotiated during login.
- ❑ Verify Status Phase Collapse:
 - ❑ Status returned in final Data-In PDU of a data transfer
 - ❑ Eliminates separate SCSI Response PDU
 - ❑ Always enabled, target may use if desired

Testing for iSCSI RAID System Useful Hardware Accelerations (If Applicable) with 10G CNA

- ❑ Check iSCSI:
 - ❑ Computes and consumes iSCSI Digests
 - ❑ Detects iSCSI Header and Data corruption
 - ❑ Verify that negotiation is fixed
- ❑ Check IPsec:
 - ❑ Encryption algorithm
 - ❑ Done as a bump in the wire
 - ❑ Re-keying and Security Association setup may be done in interface firmware
 - ❑ Security association setup before 1st login command reaches iSCSI layer
- ❑ Check No TOE is applicable:

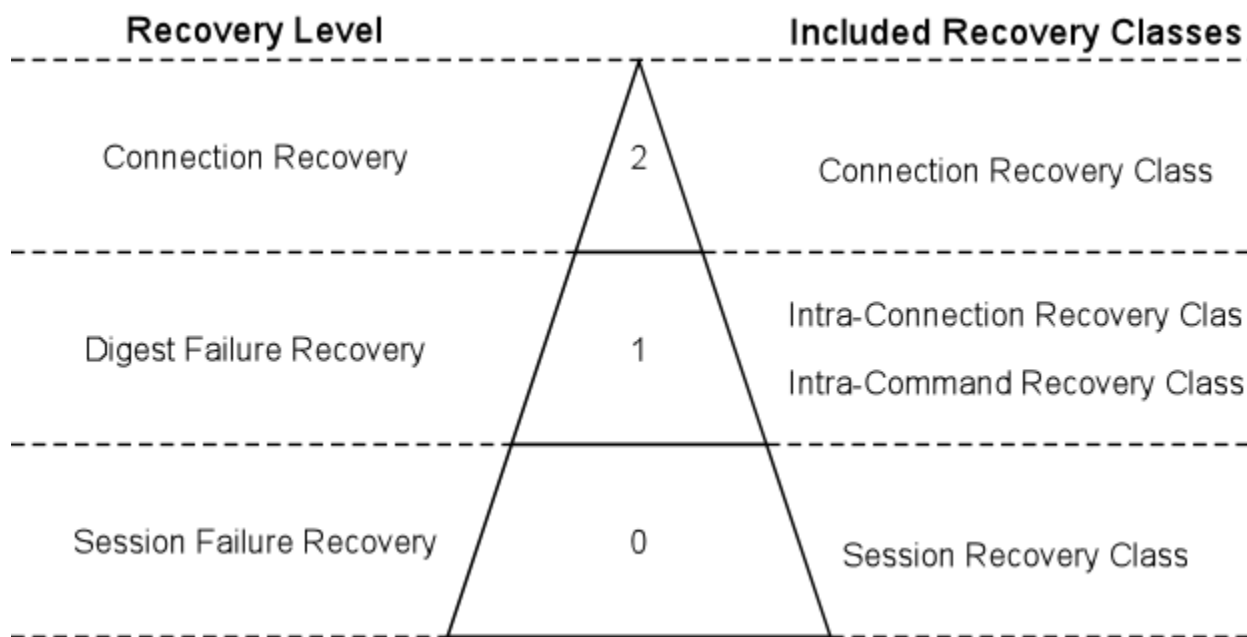
Testing for iSCSI RAID System iSCSI Error Recovery with 10G CNA

- ❑ Check the inherited Problems in iSCSI:
 - ❑ Issue I/O stress because IP protocol is less reliable than parallel SCSI or FC
 - ❑ Verify longer distances over possibly less reliable connections
 - ❑ Verify transport errors
 - (Traditional solution for transport errors is to reset the transport
 - ❑ Where is that **Internet Reset Button/Command?**)
- ❑ Negative tests
 - ❑ Inject the following errors
 - ❑ Types of errors:
 - ❑ Lost PDUs
 - ❑ Corrupted PDU header, data, CRC, and trailer via a jammer
 - ❑ Swallow / addition iSCSI control PDUs
 - ❑ Protocol Errors (Refer to standard)
 - ❑ Lost TCP connection
 - ❑ Verifications:
 - ❑ Handle error at lowest level in protocol possible to limit effect on SCSI layer
 - ❑ Preserve the connection if possible to contain the error in the iSCSI layer of the protocol
 - ❑ Defines error classes and error recovery level hierarchy
 - ❑ Rely on TCP layer for packet level error detection and retry

Testing for iSCSI RAID System Recovery Classes with 10G CNA

- ❑ Verify within a Command
 - ❑ Retry command without SCSI layer intervention
 - ❑ Errors Handled:
 - Data Digest Errors, Lost Data PDUs, and Header Digest Errors
 - ❑ When iSCSI receive markers are not enabled Header Digest Errors cause a connection failure because PDU framing is lost.
- ❑ verify within a Connection
 - ❑ Connection is preserved but a command may need to be restarted
 - ❑ Errors handled:
 - Lost Command PDUs, Lost Response PDUs, status and responses not acknowledged
- ❑ Verify Connection Recovery
 - ❑ Connection is rebuilt and commands are restarted
 - ❑ Allegiance of commands on the failed connection must be reassigned
 - ❑ Errors handled:
 - lost TCP Connection, asynchronous message from target reporting connection termination, Header Digest Errors if iSCSI markers are not used
- ❑ Verify Session Recovery
 - ❑ Session and all its connections are terminated
 - ❑ Commands must be retried by SCSI layer – Big hit on performance
 - ❑ Errors handled:
 - Any error that cannot be handled in a lower recovery class

Recovery Hierarchy with I0G CNA



- ❑ Error recovery level 0 is required
- ❑ Error recovery levels 1 and 2 are optional.
- ❑ Error recover level is negotiated at login.
- ❑ Error recovery setting is Session Wide.

Testing for iSCSI RAID System IP Security with 10G CNA

- ❑ Verify IPsec protocol provides
 - ❑ Set of security services in the IP layer
 - ❑ Established before and separate from the iSCSI session
 - ❑ Designed for peer to peer security and secure tunneling
- ❑ Verify Capabilities:
 - ❑ Integrity – detecting modified, inserted, or deleted data
 - ❑ Confidentiality – protects data from examination by encryption
 - ❑ Anti-replay – detects retransmission of PDUs
 - ❑ Authentication – per packet verification of sender
- ❑ Verify two levels of security provided:
 - ❑ IPsec AH – Addition of authentication header to packets
 - ❑ Provides integrity, anti-replay, and confidentiality
 - ❑ IPsec ESP – Encrypts and encapsulates data payload
 - ❑ Adds confidentiality to security capabilities
 - ❑ Computationally intensive – requires HW assist
 - ❑ iSCSI requires use of 3DES and AES algorithms
 - ❑ IPsec ESP includes per packet authentication

Testing for iSCSI RAID System Login with 10G CNA

- ❑ Verify the process by which a TCP connection and iSCSI session is established between an initiator and target
- ❑ Verify operations performed during login
 - ❑ Security Negotiation:
 - ❑ Security association of the connection is established
 - ❑ Authentication:
 - ❑ Establishing end to end trust between the initiator and target
 - ❑ Identification:
 - ❑ Exchange of target and initiator iSCSI names and aliases
 - ❑ Definition of iSCSI session and connection Ids
 - ❑ Declaration of session type
 - ❑ Operational Parameter Negotiation:
 - ❑ Agreement on session and connection operational parameters
 - ❑ Standard Features:
 - ❑ Queue depth, maximum number of connections in session, error recovery level, data segment lengths, data burst length, data PDU ordering
 - ❑ Optional Features:
 - ❑ Fixed interval markers, iSCSI digests, immediate and unsolicited data

Testing for iSCSI RAID System Login with 10G CNA (cont)

- Verify login Process:
 - A sequence of Login Request PDUs from initiator and Login Response PDU's from target
 - Authentication and operational parameter data is passed between initiator and target in named key/value pairs in the PDU data segments:
 - Example Data Segment from a leading iSCSI Login Request
 - InitiatorName=eui.I234567890ABCDEF
 - InitiatorAlias=bobspc
 - AuthMethod=None,CHAP
 - TargetName=eui.FEDCBA0987654321
 - TargetAddress=storagearray:3270:3
 - SessionType=Normal
 - Example Reply from the storage array
 - TargetAlias=honkinBigArray
 - AuthMethod=None
 - TargetPortalGroupTag=3
 - During login, only the Login Request, Logout Request, and Reject PDUs are allowed

Testing for iSCSI RAID System Login with 10G CNA (cont)

- ❑ Verify Four” Login/Session Phases:
 - ❑ Security Association:
 - IPsec Security Association is established on TCP connection
 - Occurs before first Login Request PDU is received by iSCSI layer
 - ❑ Authentication:
 - Optional login phase where secrets and challenges are exchanged between the initiator and target in Login Request and Login Response PDUs
 - Must be the first explicit login phase
 - If target or initiator agree to the “None” authentication method, the Authentication phase can be skipped
 - ❑ Operational Parameter Negotiation:
 - Negotiation of required connection and session operational settings
 - Negotiation of optional feature operation
 - If target or initiator agree to the enter full feature phase the operational parameter negotiation can be skipped and all operational parameters are assigned default settings
 - Most vendors are doing explicit negotiation of operational parameters
 - ❑ Full Feature Phase Operation:
 - Non-login commands can now be sent to the iSCSI target

Testing for iSCSI RAID System Authentication with 10G CNA

- ❑ Verify process by which the initiator and target establish end to end trust.
 - ❑ Occurs only during login.
CHAP re-challenge is not supported by iSCSI
 - ❑ Authentication Methods:
 - None – No authentication is performed
 - CHAP – Challenge Handshake Authentication Protocol
 - KRB5 – Kerberos V1
 - SPKMI and SPKM2 – Simple Public-Key GSS-API Mechanism
 - SRP – Secure Remote Password
 - ❑ Only None and CHAP must be implemented

Note: iSCSI standard specifies which optional features and strength levels are required for each supported Authentication Method

Testing for iSCSI RAID System Session Types with 10G CNA

- ❑ Verify two types of sessions supported: **Normal and Discovery**
- ❑ Verify declared by initiator in leading Login Request PDU in the SessionType operational parameter
- ❑ Verify Discovery Session
 - ❑ Commands limited to:
 - Text Request with “SendTargets” key/value pair
 - Logout Request with reason “Close the Session”
 - ❑ No access to storage resources
- ❑ Verify Normal Session
 - ❑ All commands accepted
 - ❑ Access to storage resources determined by internal storage access control lists

Testing for iSCSI RAID System SendTargets – Discovery with IOG CNA

- Verify the following Characteristics:
 - Defined by iSCSI specification
 - Done with Text Request and Text Response PDUs
 - SendTargets available in full feature phase both normal and discovery sessions
 - Discovery sessions exist only for SendTargets discovery
 - Available on all network portals of a target device

We are considering an option to disable “Discovery Sessions” on an array
- Verify SendTargets response depends on the session type and the SendTargets request’s argument
- Verify initiator must be pre-configured with at least one address of each of its targets to use SendTargets discovery

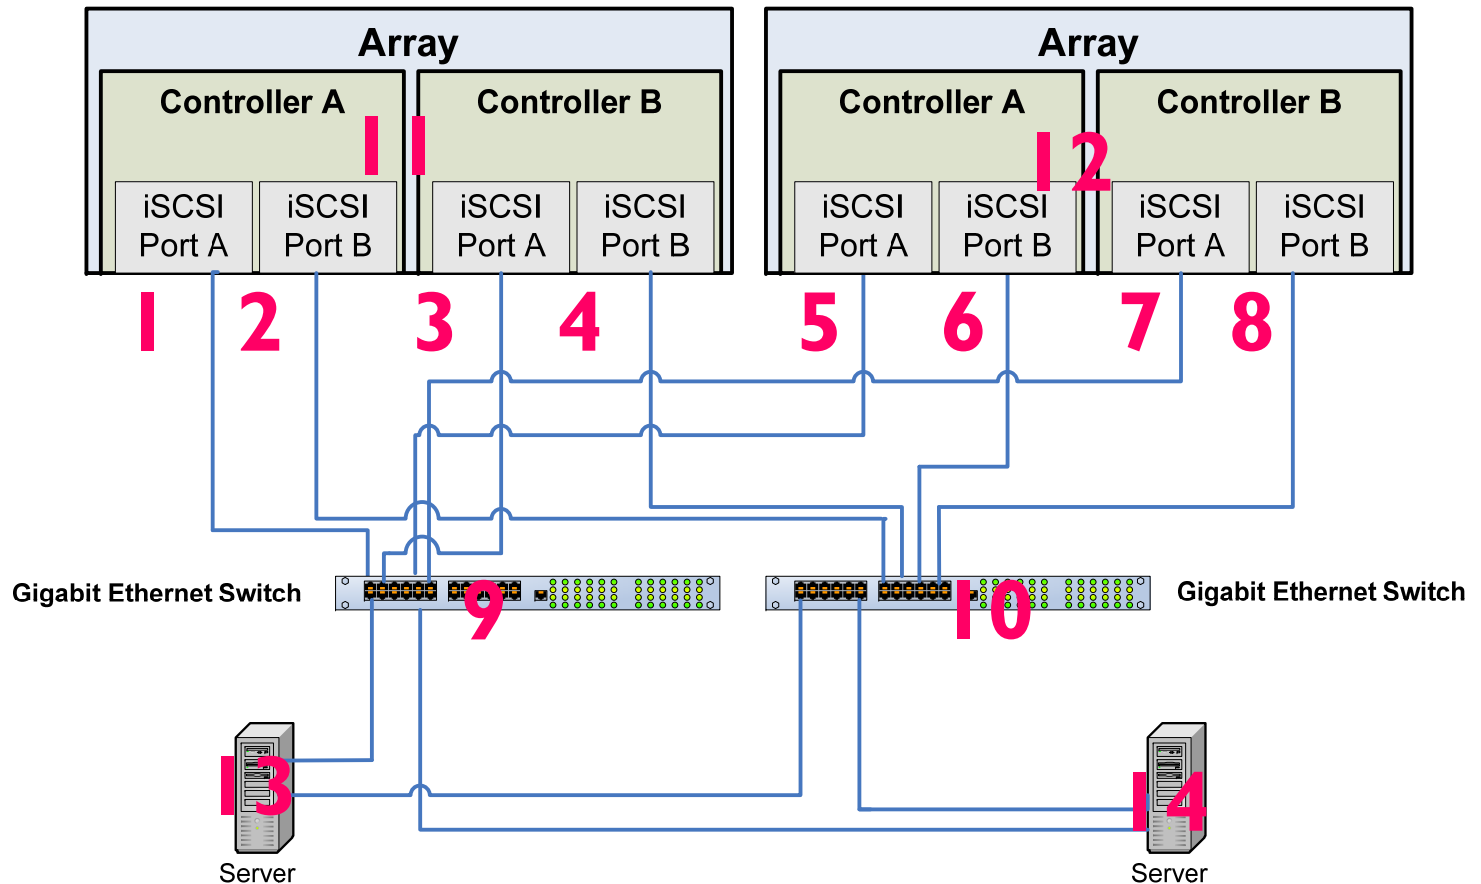
Testing for iSCSI RAID System SLP with 10G CNA

- ❑ Verify that discovery is performed using “Service Location Protocol”
- ❑ iSCSI use of SLP is defined in:
 - Finding iSCSI Targets and Name Servers Using SLP
- ❑ Verify the following characteristics:
 - ❑ Network portals are formatted as SLP service advertisements
 - ❑ SLP service advertisements are URLs
 - ❑ Protocol is not routable because it uses broadcasts
 - ❑ Does not support discovery domains
- ❑ Verify SLP Service Agent
 - ❑ Runs on iSCSI device with targets
 - ❑ Locates the SLP server by broadcast, domain name, or address
 - ❑ Exports service advertisements to SLP servers
 - ❑ Returns service advertisements to SLP user agents
- ❑ Verify SLP User Agent
 - ❑ Runs on iSCSI device with an initiator
 - ❑ Locates the SLP sever by broadcast, domain name, or address
 - ❑ Queries the SLP server for target service advertisements
 - ❑ Broadcasts to locate SLP service agents to get service advertisements if an SLP server is not found

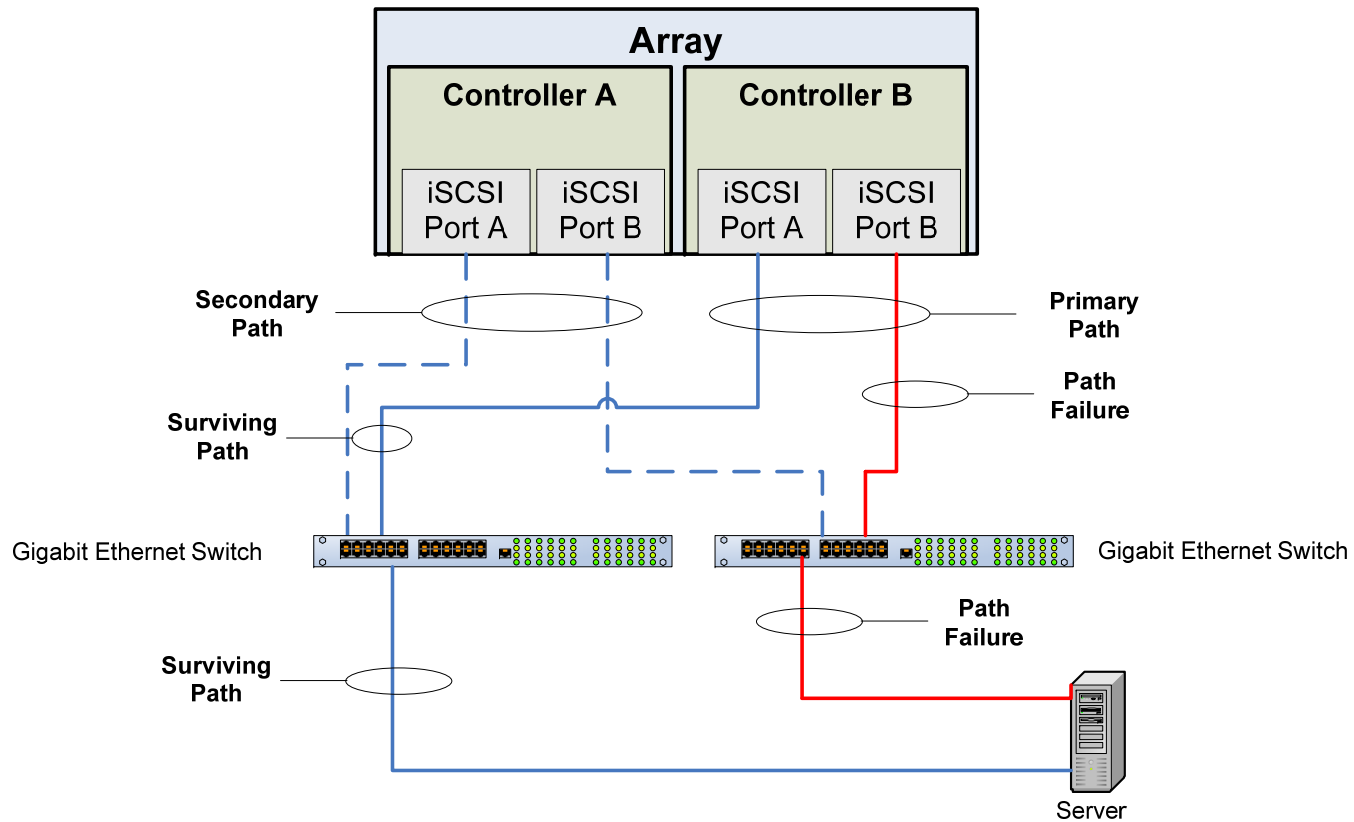
Testing for iSCSI RAID System Ancillary Protocols with IOG CNA

- ❑ Verify DHCP – Dynamic Host Configuration Protocol
 - ❑ Lookup Network Portal IP addresses
 - ❑ Lookup iSNS server address
- ❑ Verify DNS – Domain Name Service
 - ❑ Domain name to address translations
- ❑ Verify SNMP – Simple Network Management Protocol
 - ❑ Export MIB information for PHY, MAC, IP, TCP, and iSCSI

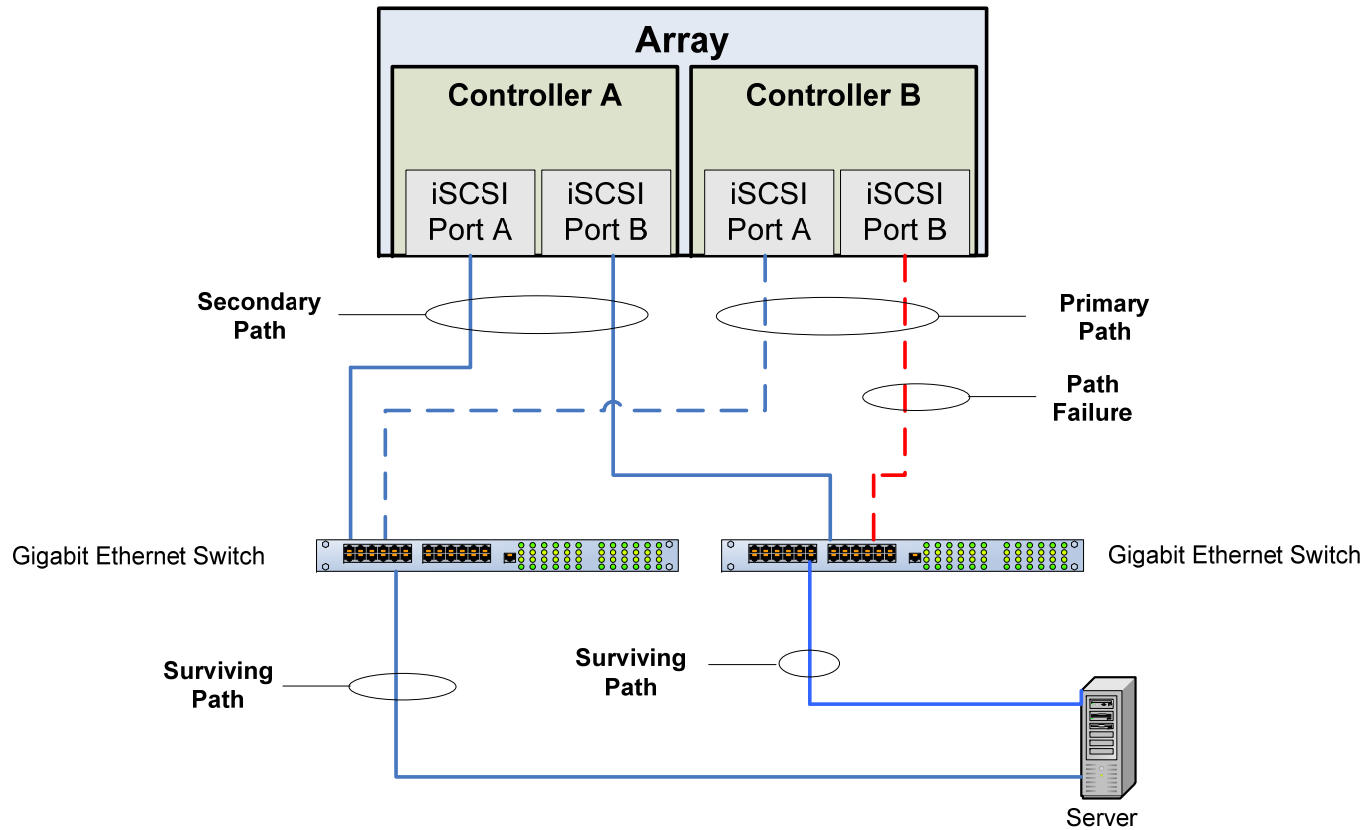
Exception Handling Test Points with 10G CNA and 10G NIC



Degraded Paths with 10G CNA and 10G NIC



Surviving Paths with 10G CNA and 10G NIC



Design for Testability “DFT”

- ❑ Emulating multiple host
 - ❑ Hardware at the 10 G CNA level
 - ❑ Software with 10G NIC
- ❑ Emulate multiple login
 - ❑ Single / Multiple sessions
 - ❑ Single / Multiple connections
- ❑ Emulate soft reset and Hard reset
- ❑ Emulate flaky ports
- ❑ Emulate flaky cables
- ❑ Emulate components' removal and reinsertion (Cable, Controller, Drive, etc)
- ❑ Unreadable sectors
- ❑ Injection of parity errors / ECC at different components
 - ❑ Cache
 - ❑ NIC and CNA
 - ❑ Memory
 - ❑ Others
- ❑ Inject syntax errors at the CMD, Data, and Status PDUs

- ❑ State capture
- ❑ Debug Queue log
- ❑ Diagnostic bundle (Host interface)
- ❑ VKI_EDIT_OPTIONS
- ❑ Major Event Logs
- ❑ Statistics on port basis
- ❑ Statistic on a connection basis
- ❑ Statistic on a Session basis

IPV4 vs. IPV6

- ❑ Verify the controllers support both IPv4 and IPv6 protocols simultaneously.
- ❑ Verify IPv4 Changes
 - ❑ As a result of having multiple protocols, the IPv4 capabilities of each port may be enabled or disabled.
- ❑ Verify IPv6 Address Notation
 - ❑ The IPv6 address space is 128 bits or 16 bytes, and is represented by eight sixteen-bit hexadecimal blocks separated by colons.
- ❑ Verify IPv6 Address Configuration
- ❑ Verify Multiple Routable Addresses
- ❑ Verify IPv6 ICMP ECHO
- ❑ Verify IPv6 VLANs
- ❑ Verify Ethernet Priority for IPv6
- ❑ Verify IPv6 MTU Size
- ❑ Verify IPv6 Optional Header Extensions

Note: Those test are CAN and NIC independent

Testing Snapshot in an ISCSI Environment with 10G CNA and 10G NIC

- ❑ Verify snapshot with Software and hardware initiators
- ❑ Verify management software control of Snapshot with different RAID levels
 - ❑ Creation Single / Multiple
 - ❑ Removal w/o deletion of base volume/ w/o I/Os / Mirror Change
- ❑ Verify volume states (Optimal, Degraded, reconstruct, etc) with snapshot
- ❑ Verify concurrent configuration with snapshot
- ❑ Verify component swap with snapshot
- ❑ Verify firmware download with snapshot
- ❑ Verify volume migration with snapshot
- ❑ Verify the functionality of the Syn / Asyn Cache command with snapshots.
- ❑ Verify NetBackup with Snapshot
- ❑ Verify Remote mirroring with Snapshot

- ❑ Verify RVM with Software and hardware initiators
- ❑ Verify that activation and deactivation of RVM
 - ❑ Firmware download
 - ❑ Component swap
 - ❑ Different volume states and RAID levels
 - ❑ Minimum and Maximum configurations
 - ❑ Small, medium, and large I/Os (Raw / FS)
 - ❑ Synchronous and Asynchronous modes
 - ❑ GUI and CLI
 - ❑ Persistent reservation W/O Cluster
 - ❑ Failing and Unfailing components

Testing Cluster in an iSCSI Environment with mixed 10G CNA and 10G NIC

- ❑ Verify cluster installation with Software and hardware initiators across different O/Ses
- ❑ Verify Client and Agent management of a single and multi-node clusters
- ❑ Verify Component firmware and driver installations
- ❑ Verify Firmware and Host Software up / down grades
- ❑ Verify snapshot and RVM in a cluster environment
- ❑ Verify Components' and Nodes' failover (Failover driver functionality)
- ❑ Verify cluster W/O Persistent reservations
- ❑ Verify system reconfiguration in a Cluster environment
- ❑ Verify Resources migration between different arrays
- ❑ Verify volume states (Optimal, Degraded, reconstruct, etc)
- ❑ Verify components swap

- Development and test teams have to verify the compatibility of 10 G iSCSI and 1 G iSCSI the followings functionalities before releasing Host software and Controller firmware to a customer
 - Basic protocol handshake / compliance
 - RAID Resource Discovery with 10 CNA and NIC
 - RAID Login with 10 CNA and NIC
 - RAID Configuration, Management, and SAN
 - CMD Exchange over a 10 G Network
 - Data Exchange over a 10 G Network
 - Status Exchange over a 10 G Network
 - Exception Handling over a 10 G Network
 - Time out
 - Failover
 - Recovery
 - Connection and component loss
 - RAID Premium Features over a 10 G Network
 - RAID Performance over a 10 G Network
 - RAID Applications over a 10 G Network
 - 10 G iSCSI and 1 G iSCSI can coexist on the same network
- Please do not focus on just testing the protocol!