

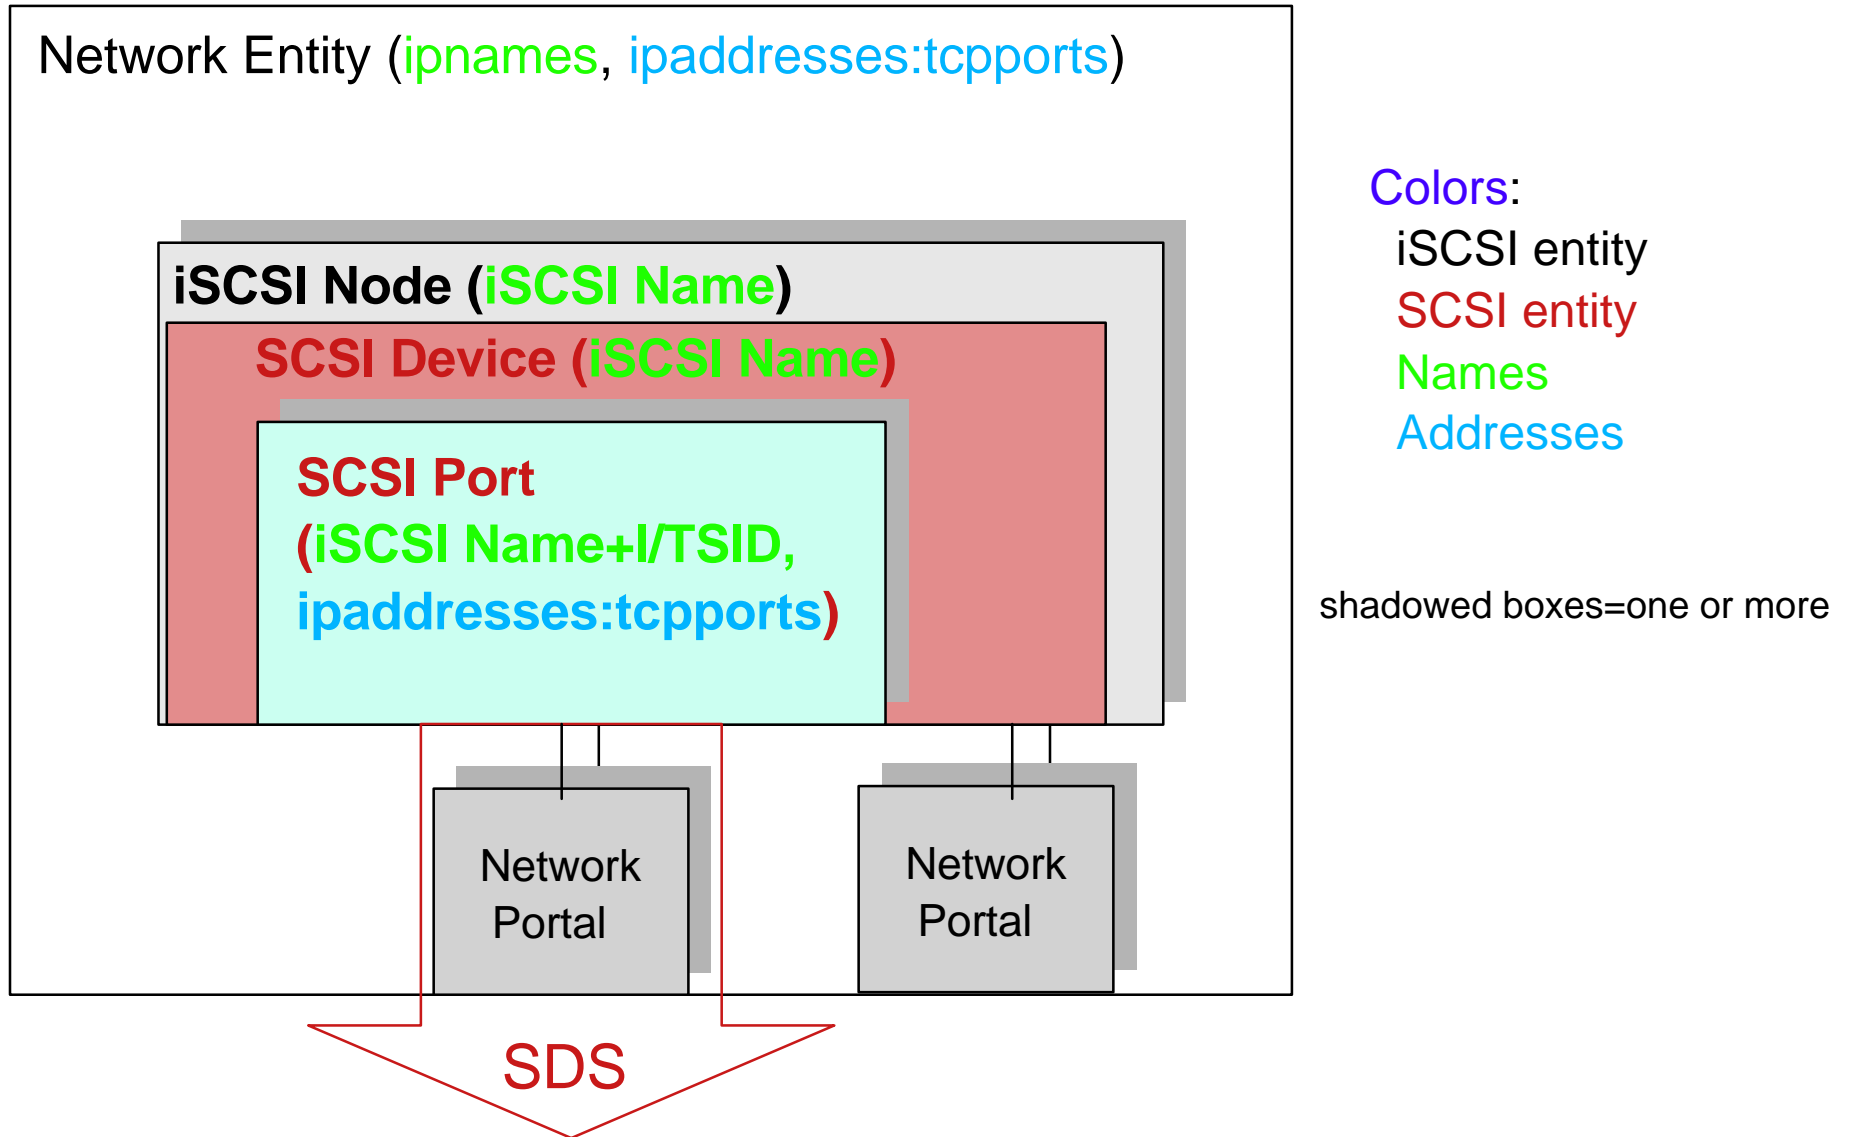
SAM-2 and iSCSI

N&DT

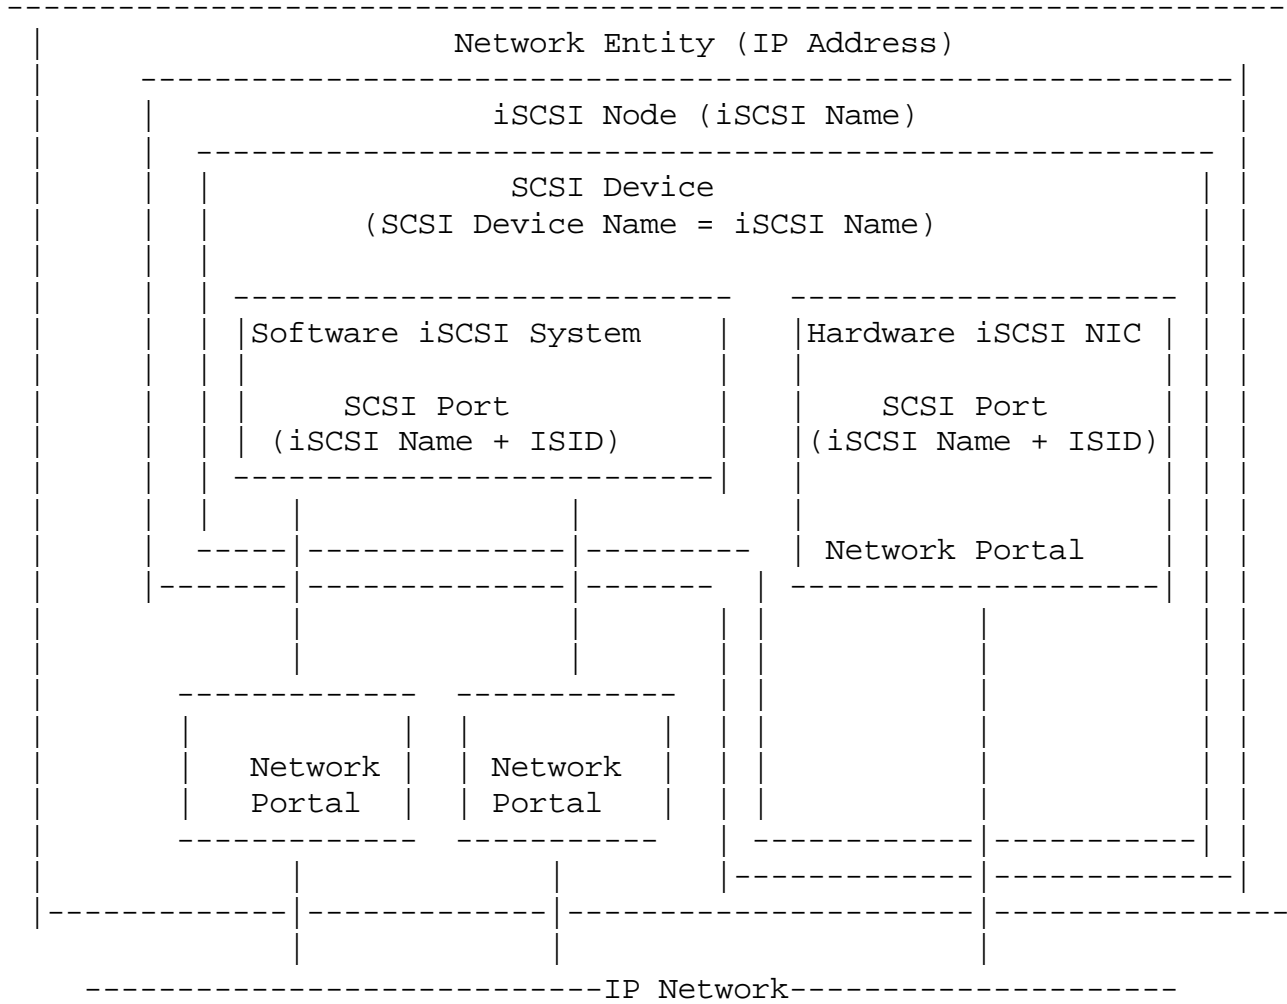
IPS Interim meeting, Nashua

May 2001

SAM-2/iSCSI model



iSCSI HW/SW types



Entity Mapping

- SCSI Device = component of an iSCSI Node
 - ▶ at most one per iSCSI Node
 - ▶ some iSCSI Node may have no SCSI devices (e.g., canonical target might serve only for "SendTargets")
 - ▶ SCSI Device Name = iSCSI Name (of iSCSI Node)
- SCSI Port = "end point of a session"
 - ▶ dynamically created when sessions get created
 - ▶ SCSI Port Name = iSCSI Name + (I/T) SID of session
 - ▶ SCSI Port Address is TBD
 - could equal name
 - could equal ipaddresses:tcpports of the session
 - "address=identifier"?
- SCSI I_T Nexus = session
 - ▶ nexus joins two SCSI Ports (not SCSI Devices)

Consequences of Mapping

- Initiator can't reuse an ISID for an independent session (no parallel nexus)
- If the session goes down, the nexus can be restored only if
 - ▶ the initiator reuses the old ISID
 - ▶ the target reuses the old TSID
 - ▶ restored state of the nexus depends on what the target can/should maintain through tear-down/rebuild phase

Nexus State (partial)

- Some nexus state
 - ▶ persistent reservations
 - ▶ task set
 - ▶ mode page settings
 - ▶ unit attention conditions
 - ▶ access control enrollment state (per spec)
 - ▶ alias lists (yet to be approved)
- Not all survive nexus tear-down/rebuild
 - ▶ persistent reservations even through target power cycles
 - ▶ depends on the cause of tear-down
- We need to build a table of events that describe tear-down and rebuild scenarios and the effects on these states

Other SCSI minutia

- SCSI Access Controls TransportID for iSCSI
 - ▶ equal to the iSCSI (Initiator) Name
 - ▶ specific byte format for this is TBD
 - ▶ should go in SPC-3 (T10)
- Formats for long identifiers (e.g., in alias designation)
 - ▶ must contain iSCSI (Target) Name
 - ▶ optionally contain IPName/Address and/or TCPPorts
 - are these "hints" on where to find the iSCSI Node?
 - are these "required" addressing information?
 - ▶ specific byte formats TBD
 - ▶ should go in IPS document

Alternative Map 1

- SCSI Port = iSCSI Node
 - ▶ pros:
 - simple
 - static
 - ▶ cons
 - counter-intuitive
 - perhaps too restrictive
 - no more than one session per Node
 - moves Node name lower in the device (e.g., NOT an OS footprint)
 - adds management complexity - iSCSI Auth is on Name of each Port

Alternative Map 2

- SCSI Port = iSCSI Port
 - ▶ new construct
 - gets a static name derived from Node Name by extension
 - multiple sessions between Nodes pairs, but only one per Port pair (independent of ISID)
 - ▶ pros:
 - more static
 - maps nicely to HW iSCSI adapter and fits with SW iSCSI
 - allows multiple ports per iSCSI Node
 - allows iSCSI Auth based on iSCSI Name (of Node, not Port)
 - ▶ cons:
 - complicates the model, not clear there's a real gain

Requirements

- General
 - ▶ returning data to the same initiator that sent the command.
 - ▶ maintaining per-initiator mode pages
 - ▶ maintaining per-initiator task sets
 - ▶ identification is implicit; i.e., the initiator identifier is
 - ▶ not carried as a payload in any data structure
- Reservations
 - ▶ recognizing initiators as different to allow or reject commands in the presense of reservations
 - ▶ persistent reservations; identifying the same initiator through logout/login phase.
- Access controls
 - ▶ to identify initiators to allow or deny access to logical units based on either TransportID or AccessID enrollment state.
 - ▶ TransportID is not carried in any SCSI payload; it's part of the transport level and is exchanged between SCSI and transport in transport specific ways
 - ▶ AccessID sent through SCSI payload establishes specific enrollment state for the nexus on which it was carried
- Third party commands
 - ▶ to identify third party targets
 - ▶ third party identifier carried in SCSI payload (e.g., Extended Copy target descriptors, and some XOR block commands)
 - ▶ alias descriptors that enable use of long third party identifiers