iSCSI Architecture

Specification-related drawings

4.12 The SCSI model for distributed communications

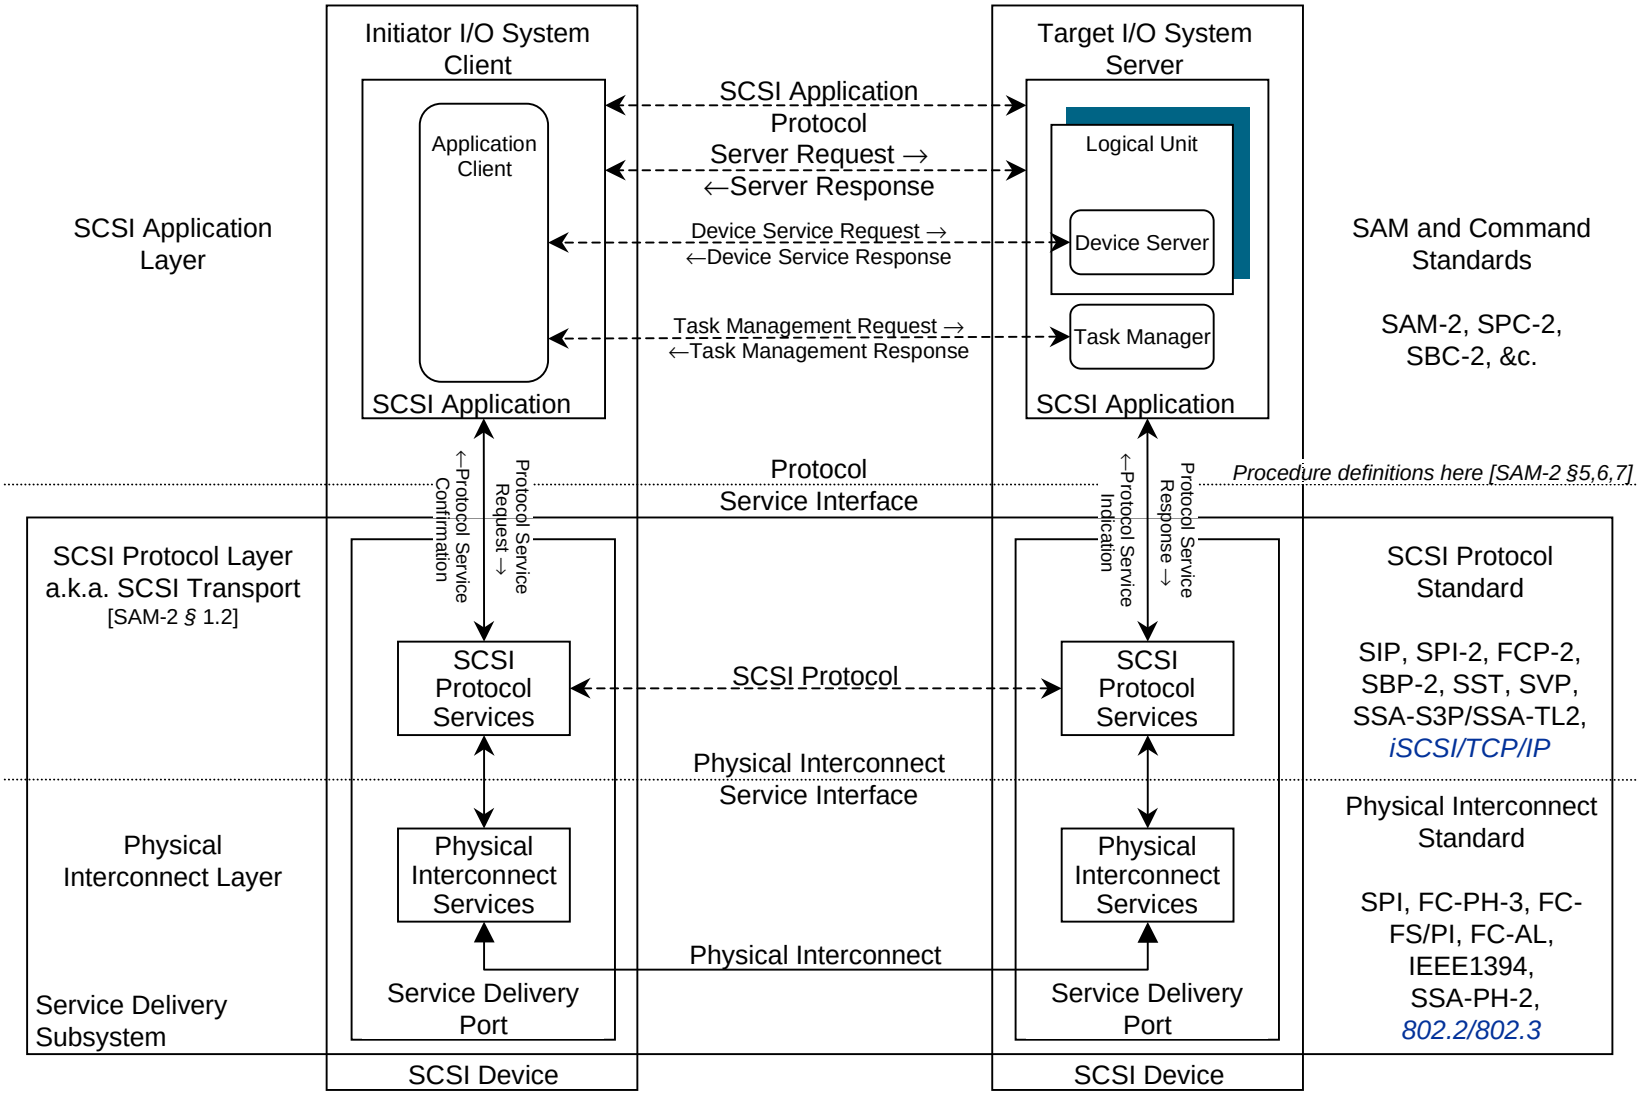

Composite of SAM-2 Fig.s 2, 5, 6, 7, 9, 26, 28

SAM-2 Service Delivery Port

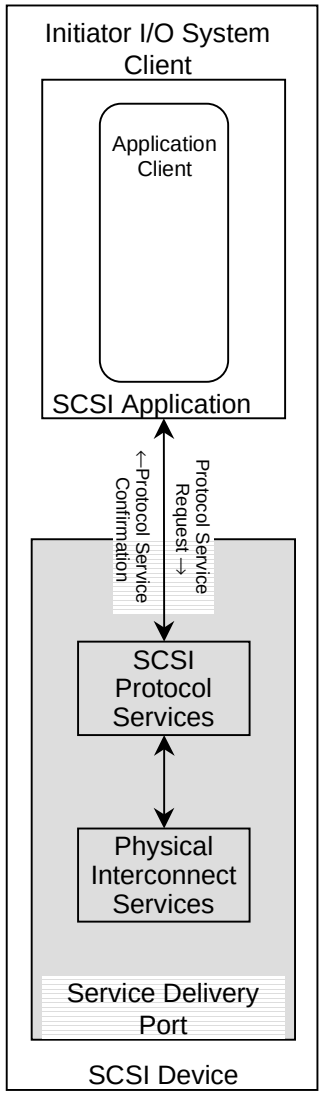

3.1.89 **service delivery port:** A device-resident interface used by the application client, device server or task manager to enter and retrieve requests and responses from the service delivery subsystem. Synonymous with “port” (3.1.61)

4.6 …the Service Delivery Port object represents the hardware and software that implements the protocols and interfaces between servers or clients in the SCSI Device and the Interconnect Subsystem.

3.1.81 **SCSI Multi-port unit:** A device that has multiple service delivery ports (see 3.1.89) or responds to multiple SCSI device identifiers (see 3.1.79)...

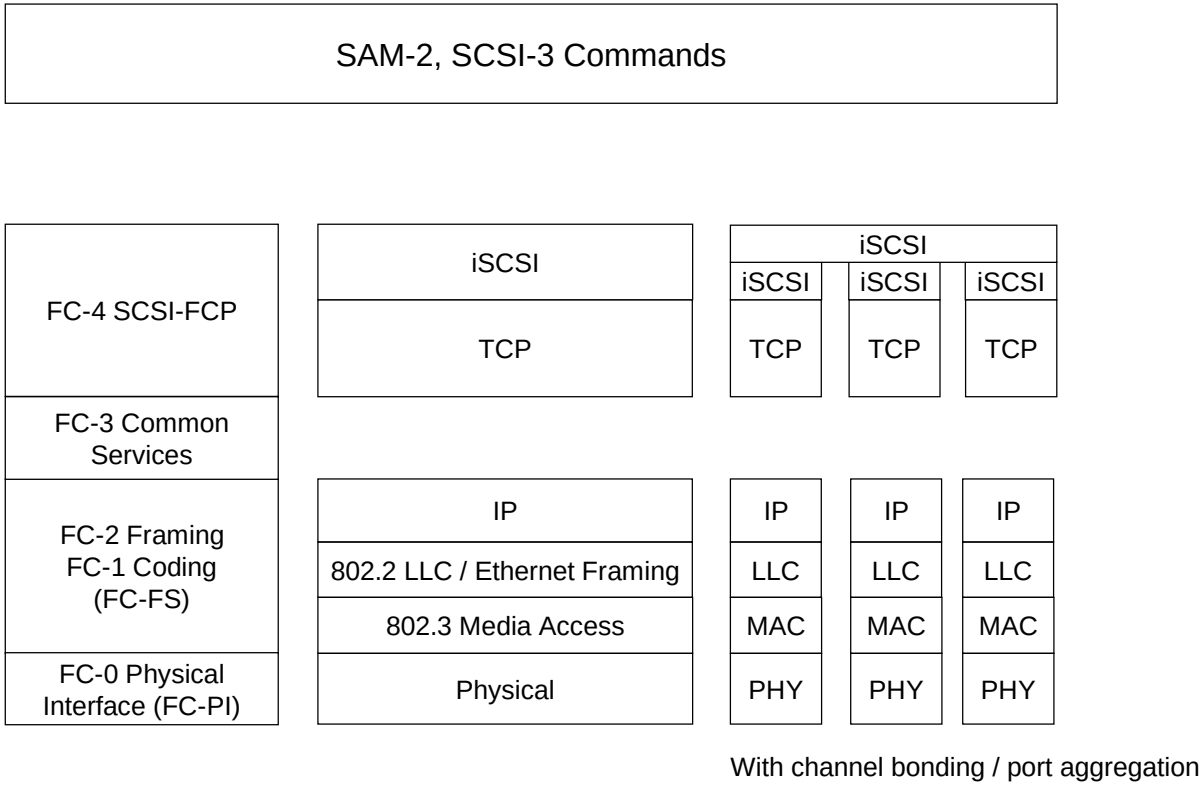

With channel bonding / port aggregation

Aggregation Alternatives

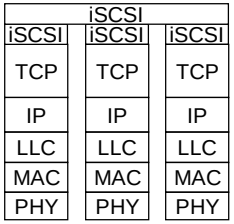

Proposed for iSCSI. Commands and status iSCSI messages are sequenced independently, in a central iSCSI module. Other iSCSI functions can be delegated to the individual protocol stacks.

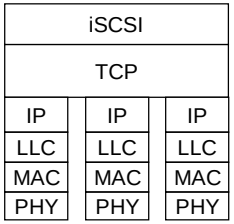

TCP is modified to aggregate over multiple IP addresses. That means that an end node can have multiple IP addresses, and the TCP implementation is able to load balance across them. Segments for the TCP connection arrive out of order at the several interfaces, but TCP is able to put them back in order using its sequence numbers. Problem: TCP connections are currently defined by the (IPaddr, Port, IPaddr, Port) 4-tuple. There is no TCP-layer connection ID to relate segments arriving on different IP addresses. Potential problem: One TCP engine must service all links.

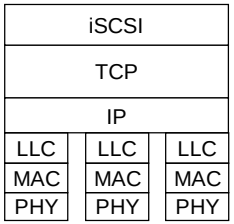

IP does the aggregation, balancing traffic over multiple links. Current routers would have difficulty preserving parallel flows, as they would tend to discover (through ARP) only one destination MAC address for a given IP address.

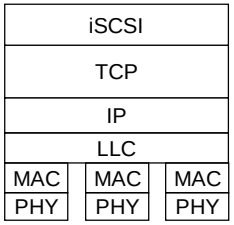

Effectively the same as above, with the additional problem that it adds a link dependency.

Aggregation Alternatives (2)

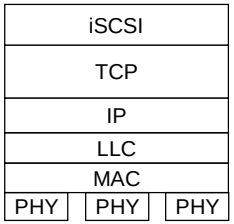

As specified by 802.3ad. Problem: frames for the same TCP connection will take the same link in a link bundle (so that they will arrive in order, which is not what's desired here).

iSCSI Session Concept

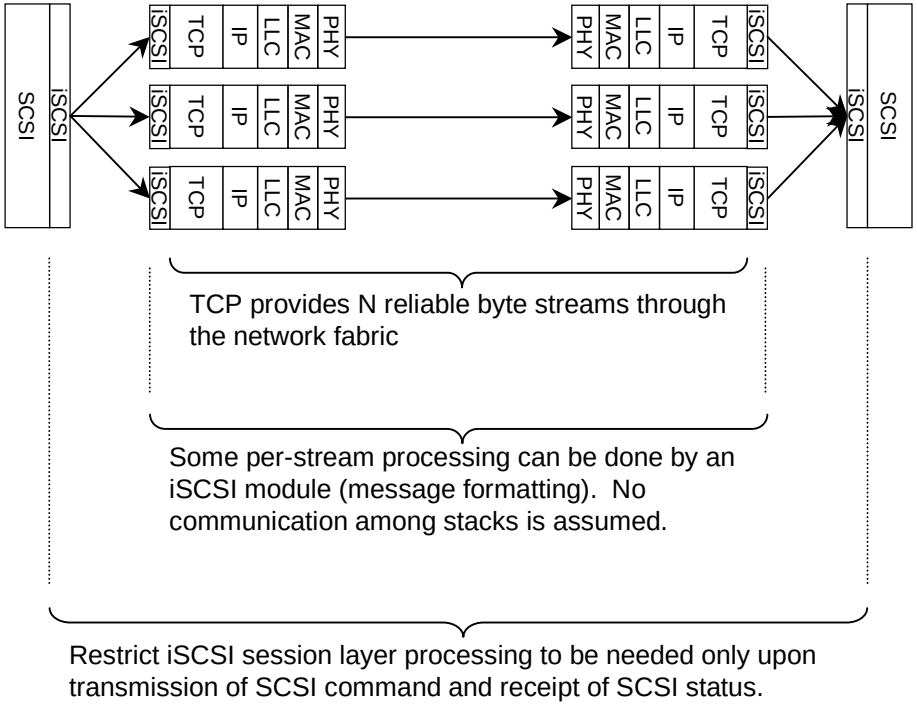

SCSI Multiport Target Unit

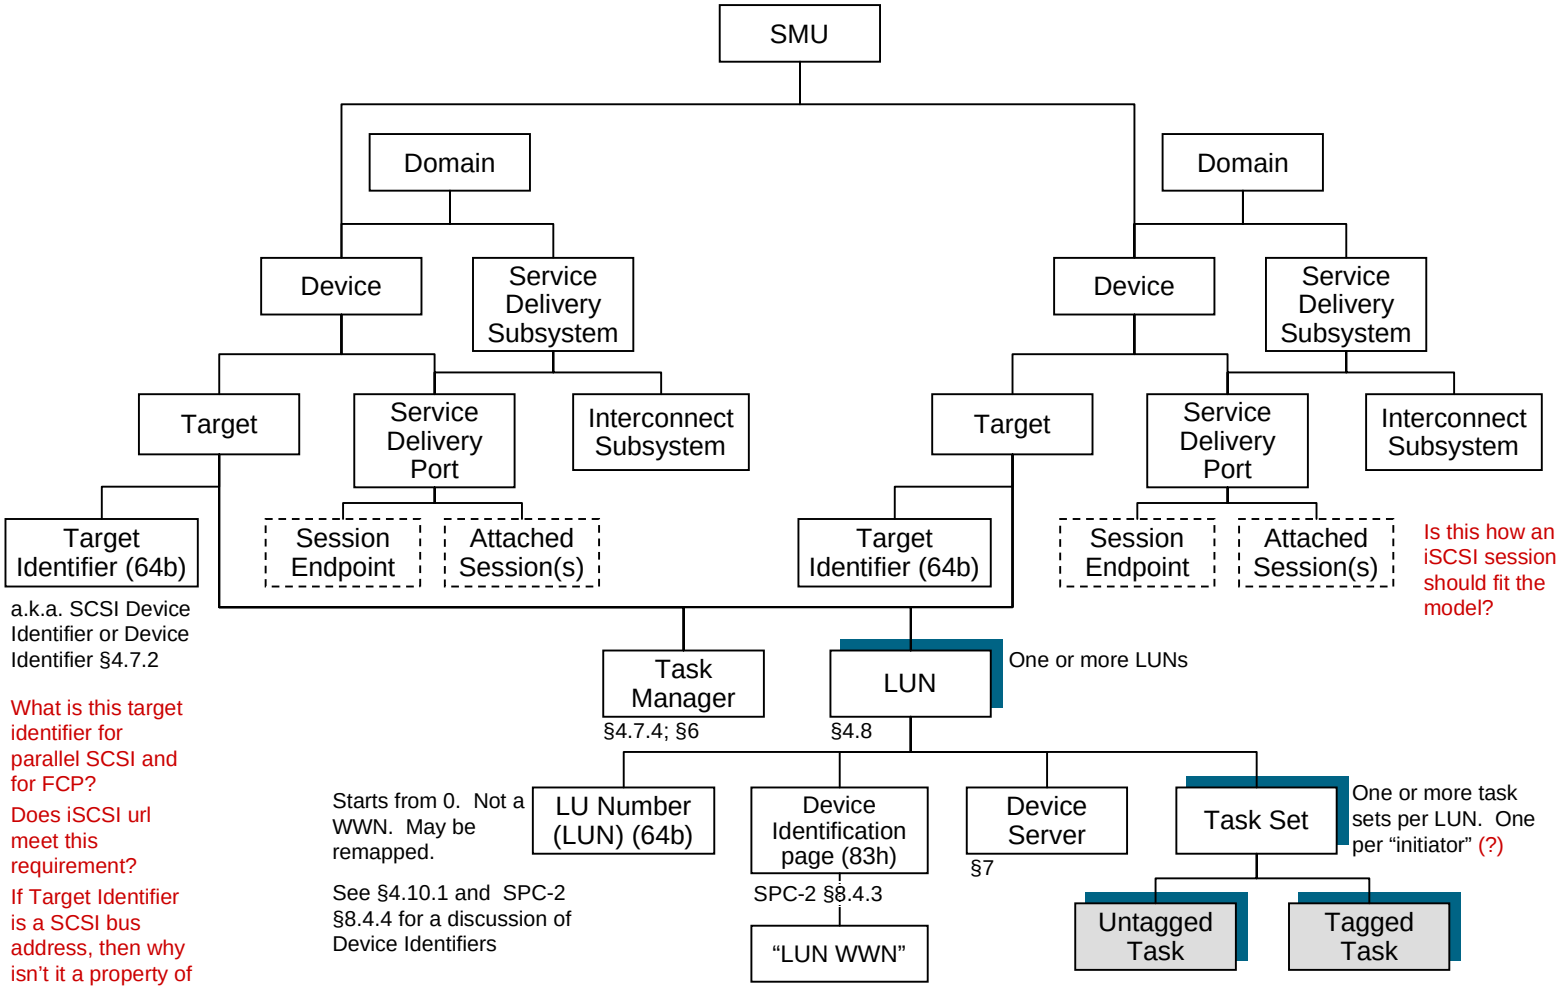

SCSI SMU Target Identification of Tasks (SAM-2)

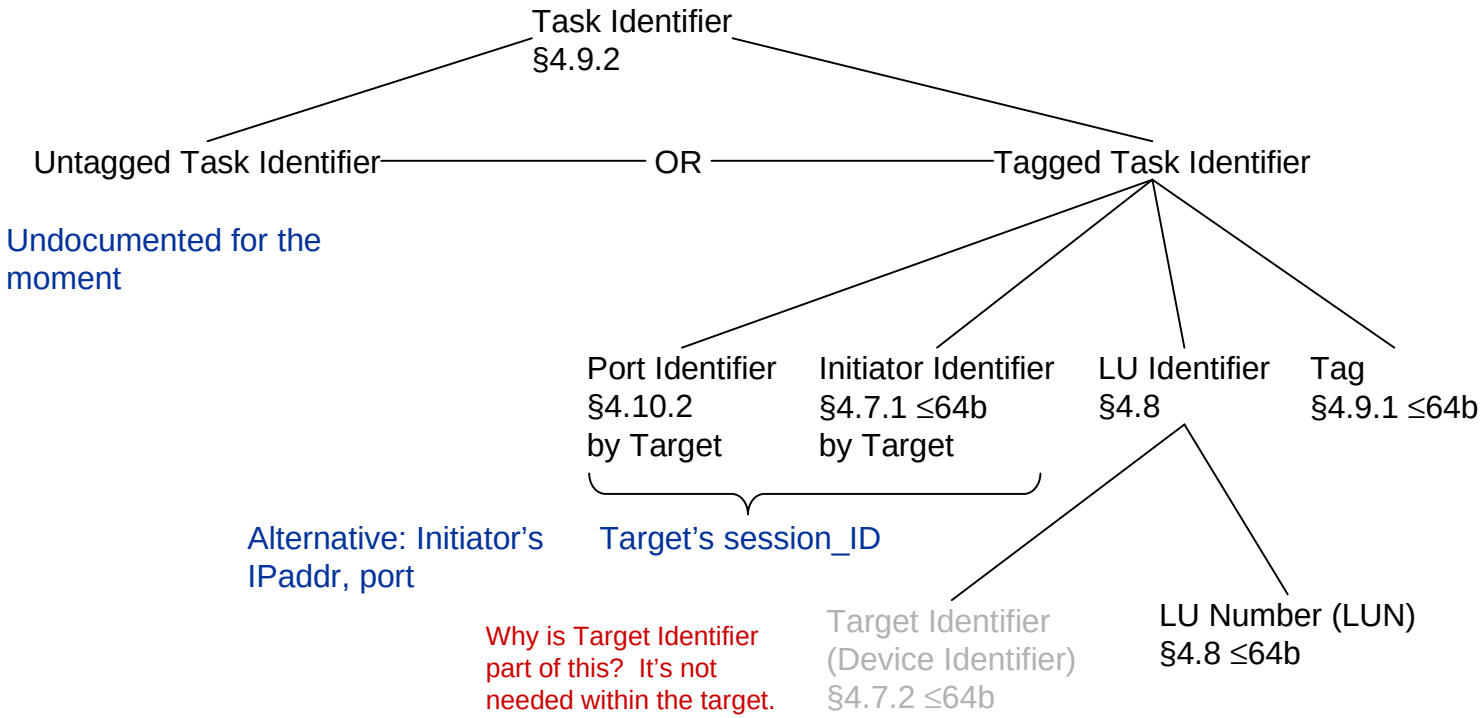

SCSI SMU Initiator Identification of Tasks (SAM-2)