

Terminology

SCSI is a block-transfer protocol that enables data transfer between server and disc in a storage array

iSCSI SCSI transport protocol that operates over TCP and encapsulate the SCSI command and data in TCP/IP byte stream. Works with any Ethernet switch without distance limitation. Not suitable for application with high I/O requirements

iSCSI Initiator an iSCSI client which transport SCSI requests and responses, encapsulated into the iSCSI protocol, between the host and the iSCSI target

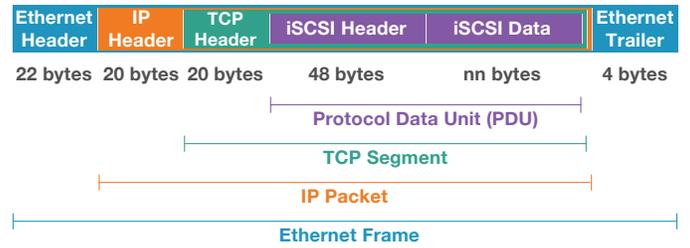
iSCSI Hardware Initiator PCI-E card, which offload iSCSI functionality for better performance or booting a server through iSCSI is possible using this card

iSCSI Software Initiator an iSCSI Initiator implemented by software and computing operations of the iSCSI initiator are performed by the server's CPU

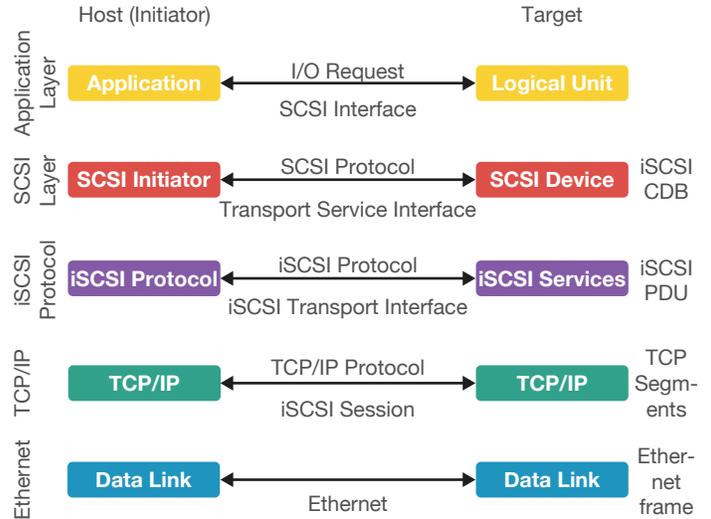
iSCSI Target A server with storage that can provide one or more logical units number (LUN) to client (iSCSI initiator)

iSCSI name is unique worldwide name (WWN) which the iSCSI node is known. The iSCSI name use on of the following formats:
iSCSI Qualified Name (IQN): can be up to 255 characters and the format takes the form of "iqn.yyyy-mm.naming-authority:unique-name" such as : iqn.2016-09.com.cisco.iscsi:array1
Enterprise Unique Identifier (EUI): takes the form of "eui.16 hex digits" for example "eui.0123456789ABCDEF"

iSCSI Encapsulation



iSCSI Protocol Layers



Traditional iSCSI vs Lossless iSCSI

	Compatibility	Prioritisation	Flow Control	Bandwidth
Traditional iSCSI	Work on all Ethernet Switches & Usefull for network with multiple hops between initiator and target	No prioritisation of iSCSI traffic	TCP layer takes care of flow control	Bandwidth shared with all the traffic on wire
Lossless iSCSI	Works only on DCB enabled switches	Can prioritize iSCSI traffic as "no_drop" using PFC	Additional Flow control using pause frames at the Ethernet layer	Bandwidth allocation to iSCSI traffic using ETS

Lossless iSCSI

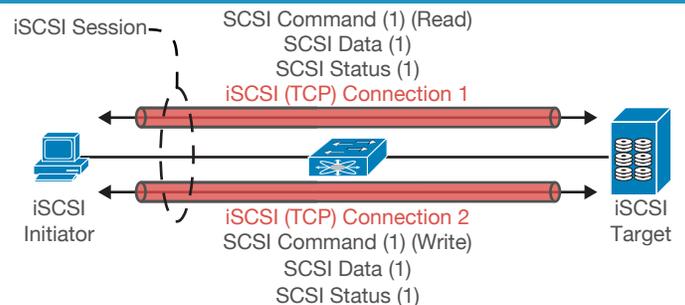
Ethernet networks are highly susceptible to broadcast storms leading to congested network. The Data Center Bridging (DCB) feature extends lossless capabilities and provide a network suitable for storage traffic. DCB feature help ensure bandwidth, throughput and performance for storage traffic. DCB provide the following features:

Priority Flow Control (PFC) - 802.1Qbb Enables lossless Ethernet using PAUSE frame, CoS assigned to "no-drop" will be PAUSED

Enhanced Transmission Selection (ETS) - 802.1Qaz Prevents a single traffic class of a "bursty" nature to starve other classes by allowing to create priority group and guarantee bandwidth

Data Center Bridging eXchange (DCBX) - 802.1Qaz Negotiate Ethernet capability's (PFC, ETS, CoS) using LLDP with other DCB capable device to simplifies management

iSCSI Session



iSCSI Session highest level of an iSCSI communication path that form between initiator and target. two types of session are defined in iSCSI: iSCSI discovery and login session used by the initiator to discover available targets. General iSCSI session to transfer SCSI data and commands after the login