

iSCSI Primer

An introduction to the iSCSI protocol






Outline

- iSCSI – what is it?
 - Concept
 - Strengths and weaknesses
- iSCSI – making it work
 - Targets and Initiators
 - Hardware considerations
 - iSCSI “HBAs”
 - Internal CPU load
 - iSNS (Internet Storage Name Service)
- iSCSI -- security
 - CHAP (Challenge Handshake Authentication Protocol)
 - IPSec (IP Security Protocol)
 - ACL (Access Control Lists)
 - VPN (Virtual Private Network)

iSCSI – what is it? -- Concept



- Instead of buying and using cables and interfaces that are useful for only handling local block-level data, why not use the in-place network to move all data?
- In addition to handling “network” traffic, the network could send IO (input/output) from hosts to storage devices such as disk drives.
- A layer is added to the network protocol to encapsulate the SCSI (Small Computer System Interface) protocol for transferring data between devices.



iSCSI strengths

- Inexpensive to operate
 - There's no need for new equipment such as cables and switches. The iSCSI protocol is designed to work with any existing IP (Internet Protocol) network architecture.
 - Software drivers are available for free from most OS (Operating System) vendors such as Microsoft.
- Leverages existing human expertise
 - Maintaining an iSCSI network is most similar to maintaining the IP network at-large.
 - System administrators don't need to manage an extra set of devices and cables.
- Distance is not a problem
 - Target storage devices can receive data from any machine accessible via the network. There's no need to directly connect them via a single cable.
 - New "fast" Ethernet speeds have the potential to solve latency issues.



iSCSI weaknesses

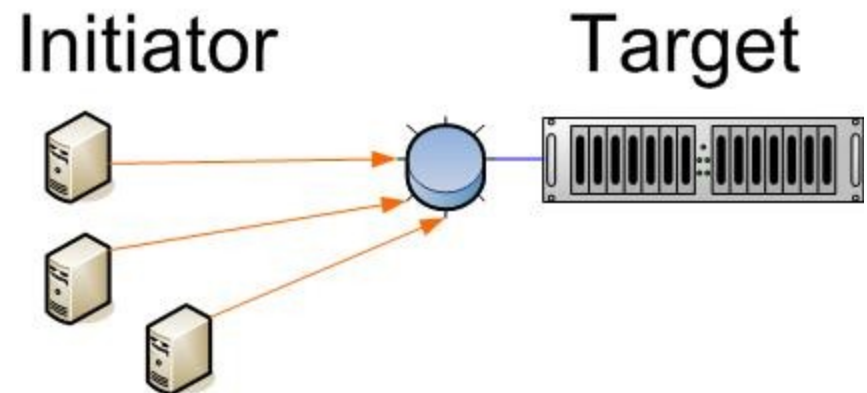
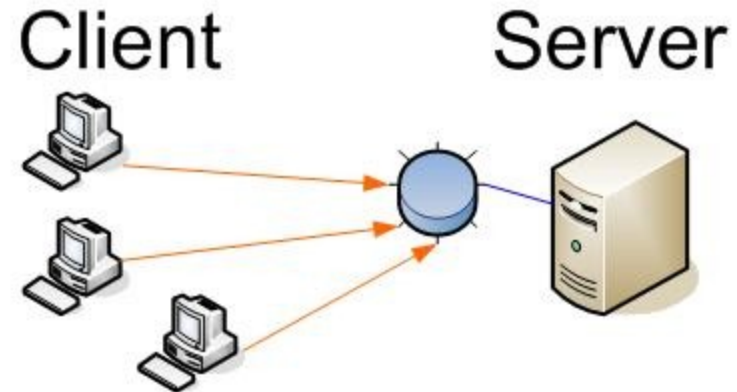
- Lack of dedicated IO path
 - Any network outage could lead to loss of data access.
 - Data must share Ethernet bandwidth with all other network traffic.
- Extra protocol layer can slow down system performance
 - CPU cycles are used to manage iSCSI protocol layers on top of basic TCP/IP commands, taking away cycles that could be used for other applications.
 - Memory space is needed for iSCSI applications such as software initiators.
- Performance is usually not equal to that of other storage devices
 - Current Ethernet is not capable of speeds equal to that of fibrechannel and SAS (Serial Attached SCSI) devices.
 - Sharing resources on the network and host machines makes it difficult to sustain high performance over any length of time.

Customer Exercise



Making it Work

- Targets and Initiators
 - Work much like the client-server network model
 - Multiple initiators can contact a single target
 - Initiators can connect to multiple targets
- iSCSI connections are simply TCP network connections
 - Standard port for starting iSCSI connections is 3260
 - Groups of iSCSI connections are called “sessions”
 - An initiator must “log in” to a target to establish a session. Multiple TCP connections may then be added to the session, but all will be managing IO from the one initiator to the one target
- iSCSI is placed in the 5th layer of the OSI (Open System Interconnection) network model






Making it Work (continued)

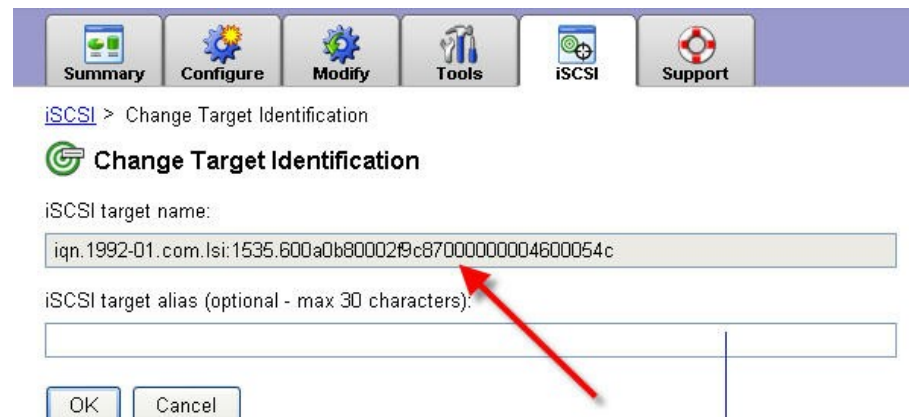
- iSCSI names
 - A unique string that identifies a particular iSCSI device
 - Regardless of the number of Ethernet ports or TCP connections involved, an initiator has only one iSCSI name. The same rule applies for targets.

iSCSI name example



- Just like the URL address of a website is independent of the IP address, an iSCSI name is independent of its IP address
 - Allows use of DHCP
 - Can be managed via an iSNS (Internet Storage Name Service) to coordinate devices


- iSCSI name is based partially on the local domain
- It can be manually edited, but it is not advisable, as the name must conform to a certain standard



Discovery Sessions

- A discovery session links a specific IP address and TCP port to a specific iSCSI device name
 - Established via IP address
 - If you know the IP address of the iSCSI device, most initiators give an option to enter this information
 - Useful if the IP addresses are static
 - Established via iSNS
 - An iSNS program runs on a server with a well-known IP address
 - Targets and initiators register with the iSNS application
 - iSNS presents targets to initiators via iSCSI names
 - Useful when using DHCP
 - iSNS can also act in manner similar to a FC fabric switch, by setting “domains” where given iSCSI targets and initiators are grouped together, but other devices are not allowed to see them

iSNS Discovery and Connection





iSCSI Portals

- A “portal” is simply the TCP component connection (Ethernet port) to the network for a given iSCSI device
- Portals can be grouped together for a single initiator/target, or they can each represent an individual initiator/target
- Portals typically have unique IP addresses, but it is possible to use a single “virtual” IP to represent multiple portals and balance network traffic across multiple portals
- Portals are important to path management drivers that must direct the IO to a specific hardware location

iSCSI Security

- CHAP (Challenge Handshake Authentication Protocol)
 - Is basically a password, stored on a given iSCSI device, that must be known by any other iSCSI device trying to establish a session with the protected device.
 - Is typically used by targets to give only “trusted” initiators access to storage, but can be configured to go both ways (mutual authentication)



Other network security

- iSCSI works with most every other network security protocol, such as IPsec (IP security) data encryption, use of VPNs (Virtual Private Network), and ACLs (Access Control List)
- Since iSCSI operates at layer 5 of the OSI model and these other methods at layer 3, iSCSI is virtually oblivious to these security methods
- CHAP provides security at the host/user level. IPsec, VPNs, and ACLs provide security at the network level



Definition Matching Exercise



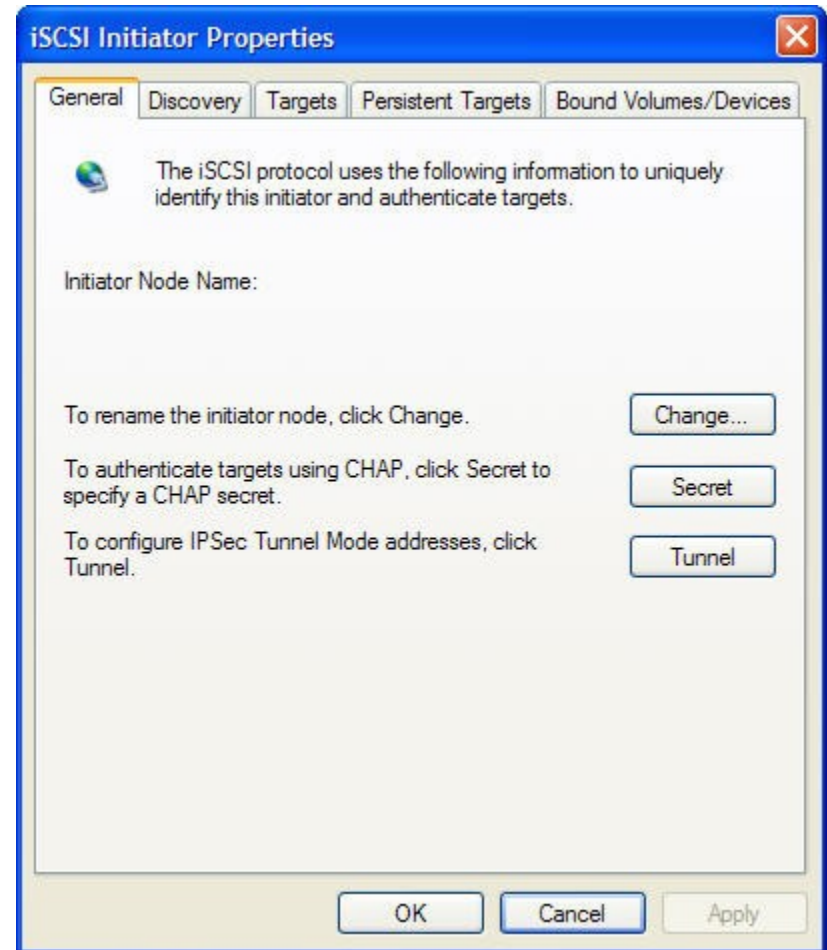
Hardware-based iSCSI implementations



- Since iSCSI does create additional work for the transport of IO, hardware solutions are available to alleviate most of the burden
- iSCSI “HBAs” (Host Bus Adapter) cards create a connection to your network specifically for iSCSI, and relieve the CPU of most iSCSI duty. Separate management tools configure the HBAs, and iSCSI devices are presented to the host as simple disk devices, like any other SATA, SAS, or FC storage device

Software-based iSCSI implementations

- Most OS vendors now provide a free software-based initiator program. These programs use existing NICs (Network Interface Card) or Ethernet ports already known to the system and append iSCSI functionality on them in the OS layer
- Software drivers consume more resources than hardware-based solutions in the form of using CPU cycles and memory



Hybrid models

- Some vendors offer NICs that provide TOE (TCP/IP Offload Engine), but not integrated iSCSI functionality (CHAP, pre-defined iSCSI names, etc.)
- These boards relieve some burden by handling most network layer functionality on a chip residing on the NIC, but the OS (and hence the CPU) must still handle iSCSI-specific functions via a software initiator or similar driver



A liger. It's pretty much my favorite animal

iSCSI host solution table

	NIC	TOE	iSCSI HBA
Definition	Network Interface card – provides Ethernet connectivity for a host	TCP/IP Offload Engine – a specialized NIC that has additional functionality	HBA card that manages an iSCSI interface independent of the host CPU
What “work” do they offload?	Physical and data link communication	TCP/IP protocol, physical and data link communication	iSCSI protocol, TCP/IP protocol, physical and data link communication
Work that must be done by the CPU	iSCSI protocol management TCP/IP protocol management	iSCSI protocol management	None
Requires SW initiator?	Yes	Yes	No – driver is provided with the HBA
Physical location	Installed in the server, usually available on the motherboard	Installed in the server, usually as an expansion card	Installed in the server using an expansion card slot
iSCSI Performance	Poor at best	Moderate iSCSI performance	Offers the best iSCSI performance by handling IO and freeing the CPU
Additional notes	Easiest to use – virtually every server has one	Can help with other networking applications besides just iSCSI	Requires more setup time and configuration



Final Exercise

- Plan an iSCSI solution