

ISCSI delivers storage over Ethernet

By Bryce Mackin

Network World, 09/09/02

ISCSI is a block-level storage protocol that lets users create a separate storage network using Ethernet. ISCSI uses Ethernet as a transport for data from servers to storage devices or storage-area networks. Because iSCSI uses Ethernet, it doesn't suffer from some of the complexity and distance limitations that encumber other storage protocols.

The iSCSI protocol puts standard SCSI commands into TCP and sends those SCSI commands over standard Ethernet. An iSCSI SAN consists of servers - with an iSCSI host bus adapter (HBA) or network interface card (NIC) - disk arrays and tape libraries. Unlike other SAN technologies, iSCSI uses standard Ethernet switches, routers and cables, and the same Ethernet protocol deployed for communications traffic on LANs (TCP/IP). It can take advantage of the same type of switching, routing and cabling technology used for a LAN.

Because iSCSI uses SCSI commands, relying on Ethernet only to transport the SCSI commands, operating systems see iSCSI-connected devices as SCSI devices and are largely unaware that the SCSI device resides across the room or across town.

Most components inside these iSCSI devices are very familiar to network professionals, including RAID controllers and SCSI or Fibre Channel drives. The only added feature is the iSCSI protocol, which can be run on standard NICs in software or on specialized iSCSI silicon or HBAs that off-load the TCP/IP and iSCSI protocol.

ISCSI is built using two of the most widely adopted protocols for storage (SCSI) and networking (TCP). Both technologies have undergone years of research, development and integration. IP networks also provide the utmost in manageability, interoperability and cost effectiveness.

When SCSI is mapped to TCP/IP through the iSCSI layer, SCSI is freed from its parallel bus structure, letting users extend network storage using Ethernet.

SCSI traffic at the application layer moves to the physical network interface (either an NIC using software or an HBA) then out over the Ethernet-based SAN to a storage device, where the SCSI command is decoded and presented to the storage operating system.

There are other components that can be added to an iSCSI SAN to facilitate inclusion of other storage protocols. These devices bridge the iSCSI protocol to Fibre Channel and SCSI devices. Bridging products let users deploy an iSCSI SAN without requiring retrofits to the entire storage infrastructure. This means that storage can be transitioned to become iSCSI-compatible, and the existing storage can stay in place.

This bridging functionality lets users preserve their existing investment while migrating to an IP-based SAN. In addition to the bridging function, some of the protocol-conversion devices also can provide switching capability. This means that these devices can bridge iSCSI to another protocol and simultaneously route the command to the appropriate device.

An Internet Engineering Task Force working group created the iSCSI standard. This group has moved quickly to offer a universal standard. However, the organization's job is not done yet. Ratification of the iSCSI standard by the entire organization is still in process but quickly moving to completion. Most are expecting ratification of the standard this year.

Devices such as iSCSI HBAs, iSCSI disk arrays, iSCSI tape and iSCSI to Fibre Channel/SCSI bridging products are available today, and more will be available throughout this year. Most of the products available today are prestandard, but all have undergone rigorous interoperability testing to ensure that they work in accordance with the standard.

ISCSI provides IT managers with a new option for solving increasing storage demands. By delivering the advantages of SANs using a familiar and ubiquitous Ethernet infrastructure, iSCSI unifies network and storage deployment, making SANs available to the masses.