

iSCSI – is it the future of Cloud Storage or doomed by NVMe-oF?



Widespread use and support mean iSCSI will still be widely deployed for at least the next few years, but its growth prospects beyond that are very unclear.

By John F. Kim, Chair, SNIA Networking Storage Forum.



What is iSCSI?

iSCSI is a block storage protocol for storage networking. It stands for “Internet Small Computer Systems Interface” and runs the very common SCSI storage protocol across a network connection which is usually Ethernet with TCP. You can read more about iSCSI at the “[What Is](#)” section of the SNIA website. (The SCSI storage protocol is also used to access block storage as part of the SAS, SRP, and FCP protocols, which run over various physical connections.)

Originally, the SCSI protocol was used only for local storage, meaning individual disk drives or direct-attached storage (DAS). Then around 1993, [Fibre Channel](#) came along and enabled SCSI to run the Fibre Channel Protocol (FCP) on top of a Fibre Channel Storage Area Network (FC-SAN). iSCSI was submitted as a standard in 2000 and grew in popularity as it was supported by more operating systems, first requiring dedicated iSCSI HBAs but later using a software iSCSI initiator that ran on top of any type of Ethernet NIC.

The dedicated iSCSI HBAs gave iSCSI faster performance that was closer to Fibre Channel performance at the time, while the software iSCSI initiator made it easy to use iSCSI from many servers without buying special HBAs for each server. Probably the biggest boost to iSCSI adoption was when Microsoft Windows Server 2008 included a software iSCSI initiator (starting in 2008, of course).

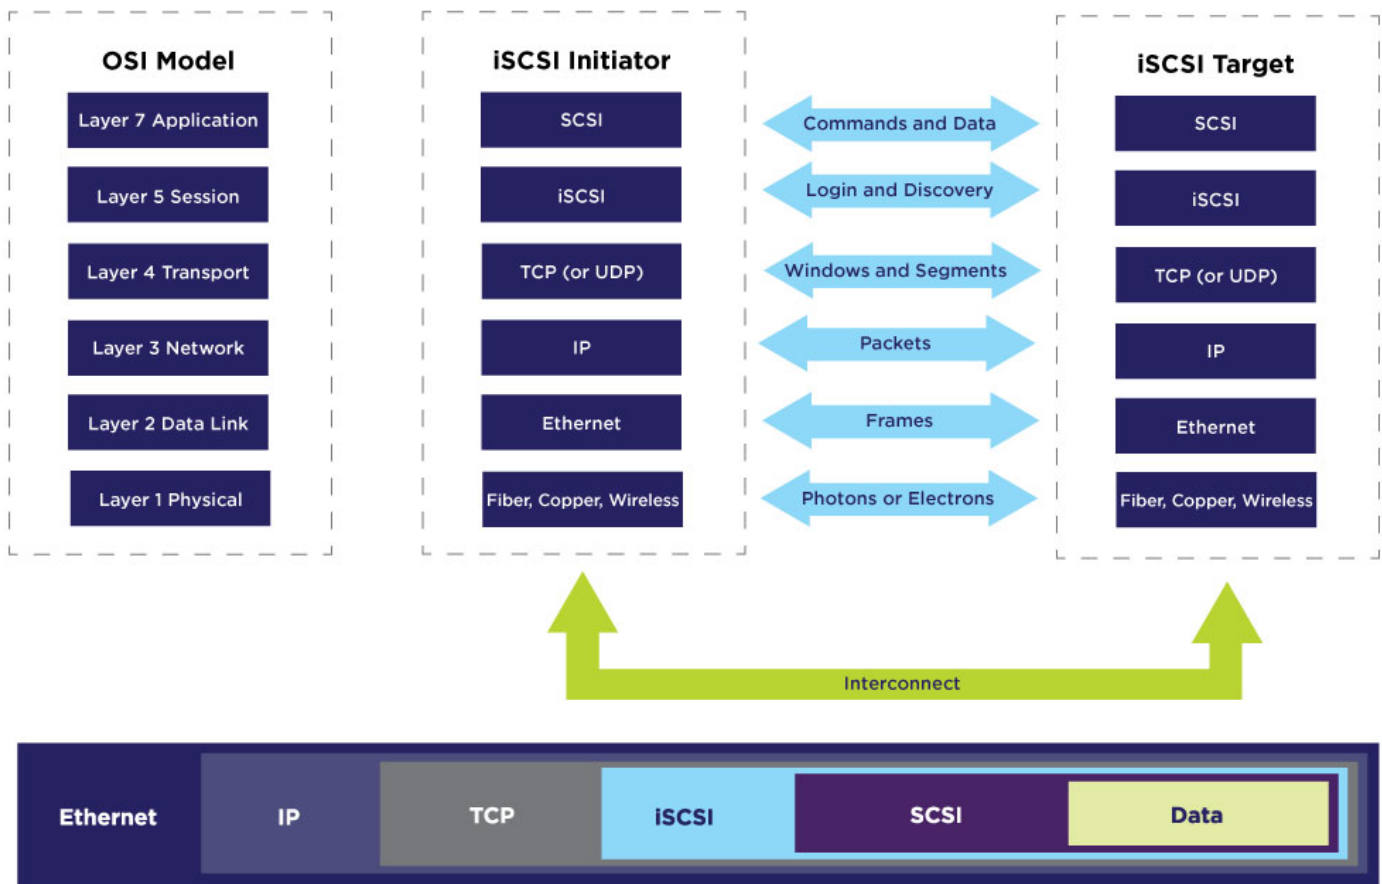
Date	iSCSI and Related Technology Milestones
1988-1994	Fibre Channel work; ANSI approval of the FC standard in 1994
1993	Arrival of first Fibre Channel products, carrying SCSI over FC
1997	First 1G FC products
1998	iSCSI technology developed
2000	iSCSI standard submitted for approval
2001	First 2G FC products
2002-2003	Solaris, Windows, NetWare, and HP-UX add iSCSI support
2002	First iSCSI HBA (1GbE)
2003	First 10GbE NIC (10GbE shipments didn't really take off until 2010)
2004-2005	First 4G and 8G FC products
2006	iSCSI Extensions for RDMA (iSER) standard; VMware adds iSCSI support
2008-2009	FreeBSD, MacOS, and OpenBSD add iSCSI
2008	Windows adds software iSCSI initiator
2010	10G Ethernet high-volume shipments begin
2011	NVMe 1.0 standard released; First 16G FC availability
2013, 2014	iSER added to Linux targets TGT (2008), LIO (2013) and SCST (2014)

2015	Availability of 25G and 100G Ethernet products
2016	NVMe-oF 1.0 standard released; first 32G FC availability
2017	VMware ESXi previews iSER (GA in 2018); Linux kernel adds NVMe-oF
2018	NVMe-oF able to run on TCP (in addition to RDMA and Fibre Channel)
2019	First shipment of 200G Ethernet products (and 400G Ethernet switches)

iSCSI Use in the Enterprise

In the enterprise, iSCSI has been used mostly for so-called “secondary” block storage, meaning storage for applications that are important but not mission-critical, and storage that must deliver good—but not great—performance. Generally, the most critical applications needing the fastest storage performance used FC-SAN, which ran on a physically separate storage network. FC speeds stayed ahead of iSCSI speeds until 2011, when 10GbE reached high volumes in servers and storage arrays, equaling 8GFC performance. Starting in 2016, Ethernet (and iSCSI) speeds pulled ahead as 25G and 100G Ethernet adoption far outpaced 32GFC adoption.

The fact that iSCSI runs on Ethernet and can be deployed without specialized hardware has made it very popular in clouds and cloud storage, so its usage has blossomed with the growth of cloud. Today, iSCSI is the most popular way to run the SCSI protocol over Ethernet networks. The rapid growth of faster Ethernet speeds such as 25G, 50G and 100G (replacing 1G, 10G and 40G Ethernet), along with increasing support for congestion management and traffic QoS on Ethernet switches, have greatly improved the performance, reliability, and predictability of iSCSI as a storage protocol.



Other Storage Protocols Threaten iSCSI

However, the emergence of NVMe over Fabrics™ (NVMe-oF) now threatens to displace iSCSI for high-performance block storage access to flash storage. Simultaneously, the growing use of file and object storage poses a threat to both iSCSI (and to FC-SAN).

NVMe-oF is more efficient than iSCSI (and NVMe is more efficient than SCSI). It was designed as a leaner protocol for solid state (flash or other non-volatile memory) so it eliminates the SCSI layer from the protocol stack and delivers lower latency than iSCSI. Note that NVMe-oF can run over Ethernet TCP, Ethernet RDMA, Fibre Channel, or InfiniBand fabrics, with the RDMA options delivering the lowest latency, but all versions of NVMe-oF (including on FC or TCP) deliver faster performance than iSCSI on the same-speed connection. So now the fastest flash arrays and fastest applications (on Linux) are moving to NVMe-oF.

iSCSI Advantage: Broad Support

Probably the biggest advantage iSCSI still holds today is that it's widely supported by all major operating systems and hypervisors. iSCSI is still supported by almost all enterprise storage arrays. NVMe-oF is currently only fully supported on Linux, and maybe only 1/3rd of enterprise storage arrays today support NVMe-oF on the "front end" meaning from server to storage (and a few of those support NVMe-oF only on Fibre Channel, not yet on Ethernet). However, VMware has announced plans to support an NVMe-oF initiator and another vendor has independently developed a Windows Server NVMe-oF initiator.

In addition, some specialized SmartNICs are able to take NVMe-oF storage and make it look like a local NVMe SSD, meaning it can be used by nearly any OS or hypervisor. (While only Linux fully supports an NVMe-oF initiator today, nearly every modern OS and hypervisor does support local NVMe SSDs.)

iSCSI Advantage: Hardware Acceleration Options

iSCSI runs most commonly on top of Ethernet using the TCP protocol, but can also run on top of InfiniBand. Also, it can run on standard network cards or specialized Host Bus Adapters (HBAs) to take advantage of either RDMA (using iSER) or iSCSI offloads and/or a TCP offload engine (TOE). iSCSI can be accelerated by using network adapters with an iSCSI hardware offload and/or TCP Offload Engine (TOE). In the former, the adapter (host bus adapter or HBA) offloads the iSCSI initiator function from the server CPU. In the latter case, the adapter offloads the TCP processing from the server kernel and CPU. Use of a TOE has fallen out of favor in some circles due to limitations that arise from handling all TCP tasks in hardware, but other forms of stateless TCP offload are still popular, can be used to improve iSCSI performance, and are supported by most enterprise Ethernet adapters.

But... NVMe-oF can also be accelerated in the network adapters, can also use RDMA, and can run over a wider variety of networks than iSCSI.

iSCSI Limitation – Block Access Only

iSCSI can only support block storage access, but file and object storage capacity are growing more rapidly than block storage capacity because so much of the new content today—audio, video, photos, log files, documents, AI/ML data, etc.—are more easily stored and used as files or objects instead of blocks. File and object storage enable easier sharing of data across multiple users and applications than block storage.

If use of file and object storage continues to grow faster than use of block storage, it could limit the growth rates of all block storage, including iSCSI, Fibre Channel, and NVMe-oF.

iSCSI – An Uncertain Future

On the one hand, iSCSI use is being driven by the growth of cloud deployments that need block storage on Ethernet. On the other hand, it's being displaced by NVMe-oF in areas that need the fastest performance, and also challenged by file and object for the storage of multi-media content, big data, and AI/ML projects. Widespread use and support—and the current OS limitations of NVMe-oF—mean iSCSI will still be widely deployed for at least the next few years, but its growth prospects beyond that are very unclear.



About the Author

John F. Kim serves as Chair on the SNIA Networking Storage Forum. He has been Director of Storage Marketing at Mellanox Technologies for six years, where he helps storage customers and vendors benefit from high performance networking and RDMA (Remote Direct Memory Access). Follow him on Twitter: @Tier1Storage

About the SNIA Networking Storage Forum

The mission of the SNIA Networking Storage Forum (NSF) is to drive the broad adoption and awareness of storage networking solutions. NSF members are dedicated to providing the global IT community with vendor-neutral information, education and exposure to networked storage solutions. To learn more about the NSF's activities and how you can join, visit snia.org/nsf.

