

STORAGE PROTOCOLS

Storage is a major consideration for cloud initiatives; what type of disk, which vendor, and as importantly which protocol? Experts will tout one over the other based on cost, performance, throughput, etc. Let's take a look at the major storage protocols at play in the data center:

Small Computer System Interface (SCSI):

SCSI is the dominant block level access method for disk in the data center. Blocks are typically the smallest unit that can be read or written to on a disk, they exist in various sizes depending on disk type and usage. Block level access means that the server can directly access the disk blocks without the need for a file system in place on top of them, this is opposite of file-based storage discussed later.

SCSI has been in use since the early 1980's and was originally used to move data within a single server. The operating system handles writing data using the SCSI protocol to a SCSI drive controller which managed one or more devices on a SCSI cable within a system chassis. The SCSI controller ensured that only one device would be active on the cable at any time which prevents contention on the SCSI bus. Because SCSI was managed by a single controller and contained within a system the chance for data loss, or contention were minimal, this meant that SCSI did not require control mechanisms to handle data loss or contention as with

networked protocols. SCSI itself is still widely used in its native format but it has also been encapsulated into other protocols for use within storage networks for consolidated storage.

Fibre Channel (FC):

Fibre Channel was designed to extend the functionality of SCSI into point-to-point, loop, and switched topologies. This allows for longer distances as well as storage consolidation. FC encapsulates SCSI data and Command Descriptor Blocks (CDB) into the payload of Fibre Channel frames. Fibre Channel networks provided the addressing, routing, and flow-control required to support SCSI data. Additionally Fibre Channel networks are designed to meet the needs of SCSI by providing 'lossless' in order delivery. This means that in a stable network FC frames will not be dropped, and are delivered in order ensuring that the Upper Layer Protocols (ULP) will not be forced to reorder or resend frames.

Fibre Channel networks are typically carried over fiber-optic links on dedicated infrastructures. These infrastructures are traditionally built-in pairs as exact mirrors of one another. This provides complete physical redundancy end-to-end. Additionally these networks provide high bandwidth and low-latency. FC networks come in 1/2/4/8 Gbps speeds with 16/32 Gbps in the works. Additionally 10Gbps FC links are typically available on a proprietary basis for links between switches.

internet/IP Small Computer System Interface (iSCSI):

iSCSI takes SCSI data and CDBs and places it in the payload of IP packets. This allows the SCSI protocol to be extended across existing IP infrastructures. While IP is routable within the data center and across the WAN iSCSI is not traditionally used/supported over routed boundaries (exceptions do exist.) The draw of iSCSI has been that storage data can be extended across the existing infrastructure with minimal additional cost.

iSCSI has not gained the market share many have predicted over the years due to flaws in the protocol and limitations of the traditional Ethernet based data center networks. until the standardization of 10 Gigabit Ethernet most data centers relied on 1GE links which were typically saturated already. This meant implementing iSCSI required new switching infrastructure. 10GE has changed the bandwidth limits but still not catapulted iSCSI into the mainstream. There are several reasons for this, one being that there is large existing investment in Fibre Channel, and two being the iSCSI protocol itself.

The problem with iSCSI from a protocol standpoint is that it takes the SCSI protocol which expects lossless, in-order delivery, and places it in TCP/IP packets which are designed to support heterogeneous WAN networks and experience packet loss and out-of-order delivery frequently. This is done without providing any additional tools to either SCSI or TCP/IP for handling the SCSI payloads in

the expected fashion. This in no way means iSCSI is unusable or should be written off it just means that additional considerations must be made when designing iSCSI, especially in the Enterprise or larger environment.

In order to provide proper performance for iSCSI on shared networks Quality of Service (QoS), physical architecture, and jumbo frame support must be taken into account. Because of these considerations many iSCSI networks have traditionally been placed on separate network hardware from the data center LAN (isolated iSCSI networks.) This has minimized some of the benefits of consolidating on a single protocol. With 10 Gigabit Ethernet and the standardization of Data Center Bridging (DCB) iSCSI looks more promising for a greater audience. For more information on DCB see my previous post (<http://www.definethecloud.net/?p=31>.)

Fibre Channel over Ethernet (FCoE):

FCoE was ratified in 2009 and provides the functionality for moving native Fibre Channel across consolidated Ethernet networks. FCoE relies on the DCB standards referenced above. FCoE encapsulates full Fibre Channel frames inside Ethernet Jumbo Frame payloads. Utilizing jumbo frames ensure that the FC frame is not fragmented or changed in any way. The FCoE and DCB standards provide a robust tool set for consolidating existing Fibre Channel workloads on shared 10GE networks while providing the lossless, in-order delivery SCSI expects. FCoE does not modify the existing Fibre Channel protocol suite and allows for the same

management model including zoning, LUN masking, etc. FCoE has started gaining ground over the last two years pushed by several large hardware vendors in the storage, network, and server markets. For more information on FCoE see my post (<http://www.definethecloud.net/?p=80>.)

Common Internet File System (CIFS):

CIFS is a file based storage system based on Small Message block (SMB.) This is a shared storage protocol typically used in Microsoft environments for file sharing. Windows-based file shares rely on CIFS as the transfer protocol of the file level data. File based storage relies on an underlying files system such as FAT32, XFS, NTFS or otherwise which differs from block based storage which does not. File level storage is an excellent medium for some applications but is not traditionally effective in others. When an application needs direct block access to disk file based storage is not appropriate. Deployments that fall into this category include some databases and most Operating Systems.

Network File System (NFS):

NFS is another file based storage protocol. NFS is traditionally used in Linux and Unix environments. NFS is also a widely used protocol for VMware environments and can offer several benefits for virtual machine storage. As a file based storage protocol NFS experiences many of the same limitations as stated for CIFS above.

Hyper Text Transfer Protocol (HTTP) and others:

When the cloud discussion leaves the data center (private/internal cloud) and moves up to the service provider level such as Google, Amazon, or the TelCos the protocols listed above may not have the necessary scalability. When you begin talking about supporting thousands of customers with multiple Terabytes each, traditional storage protocols may not suffice. It has to do with both the scalability of the systems and the administration of the disk. iSCSI and FC both require a fair amount of management for the RAID, volumes, and LUNs, whereas CIFS and NFS require a fair amount for the security and volumes. Protocols such as HTTP based storage are being used to simplify storage configuration and increase its scalability.

Which is the right protocol to use when moving to the cloud? Obviously there is only one answer! As always in IT 'it depends.' Each protocol has its uses, benefits and drawbacks. The most important thing to remember is that most environments can benefit from more than one or all of these protocols. Every application is different and any given protocol may have advantages for a particular app. The only universal truth in cloud storage is that protocol flexibility will be key.

Source: <http://www.definethecloud.net/storage-protocols/>