SCSI Standards Architecture

a report by
John Lohmeyer

Chair, InterNational Committee for Information Technology Standards (INCITS) T10 Technical Committee

John Lohmeyer is Chair of the InterNational Committee for Information Technology Standards (INCITS) T10 Technical Committee. He is also currently a principal engineer for LSI Logic and has previously worked for NCR and Symbios Logic as a design engineer and storage architect. He has participated actively in SCSI standards development for 22 years and served as Technical Team Leader for the Serial Attached SCSI (SAS) Working Group as it developed the proposal that served as the basis for the SAS standard project. Mr Lohmeyer obtained his BS in Electrical Engineering from Purdue University.

'Divide and conquer' is a strategy that military planners have employed since the beginning of history. It is also the name of an algorithmic technique that conquers a problem by dividing it into two or more smaller problems. Small computer system interface (SCSI) standards architecture also divides and conquers.

The first two SCSI standards (SCSI and SCSI-2) were all-in-one standards. They included everything from cables and connectors to protocols to command sets in one standard. SCSI took three years to develop and is about 200 pages. SCSI-2 took nearly eight years to develop and is over 600 pages.

When the InterNational Committee for Information Technology Standards (INCITS) T10 Technical Committee was contemplating writing an SCSI-3 standard, it was apparent that another all-in-one standard was disadvantageous. As the size of the standard grew, it took longer to develop and approve. The entire standard would be delayed by any aspect that was not quite finished. Finally, people wanted to map the SCSI command sets onto newer serial interfaces; it did not make sense to include several physical interfaces in one standard.

So, T10 adopted a layered standards approach much like the International Organization for Standardization (ISO) Reference Model for networking standards. This approach divides SCSI into multiple layers of standards. It also separates the pieces so they can proceed at their own pace. The lowest layer deals with physical interfaces (also called transports). The next layer up deals with transport protocols usually directly associated with one physical transport standard. The top layer consists of command sets associated with specific devices such as disk drives or tape drives. Instead of one big standard, there are many smaller standards.

The downside of layered standards is the confusion brought about as a result of there being over 30 SCSI standards (see *Figure 1*).

SCSI Standards Building-blocks

Each box in *Figure 1* represents one or more standards. A number after the acronym for a standard means it is the nth generation in that family of standards. So, SBC is the first SCSI Block Commands standard, and SBC-2 is the second-generation SCSI Block Commands standard.

The green standards are the seven different transports; each SCSI product uses only one column of these transport standards. All SCSI products conform to the SCSI Architecture Model (SAM) (yellow) and all SCSI products conform to the SCSI Primary Commands (SPC) (orange). Finally, each SCSI product incorporates one (or more) of the command set standards (blue).

Fortunately, only a few of these standards apply to a given product. Building an SCSI product involves picking the right set of building-blocks – one standard from each colour group. Currently, shipping parallel SCSI disk drives conform to SCSI Parallel Interface-4 (SPI-4), SAM-2, SPC-2 and SBC-2. Serial Attached SCSI (SAS) disk drives will conform to SAS, SAM-3, SPC-3 and SBC-2. A tape drive for SAS would replace SBC-2 with SCSI Stream Commands-2 (SSC-2).

Leveraged Software

Layered SCSI architecture is mirrored in the software architecture as well. Operating systems use layered driver structures such as SCSI Common Access Method (CAM) or the Microsoft Windows miniport drivers. This architecture has one or more class drivers that share one or more host bus adapter drivers. This allows both disk drives and tape drives to coexist on the same parallel SCSI bus. It also permits a disk class driver to control both parallel SCSI and Fibre Channel disk drives. This protects software investment as people migrate to newer physical interfaces.

Is it SCSI, SCSI-3 or SCSI-4?

It was natural to expect the generation of SCSI standards after SCSI-2 to be called SCSI-3. The titles of these standards all started with 'SCSI-3'. However, when T10 began to revise these standards, naming conventions became an issue.

Figure 1: SCSI Standards Architecture

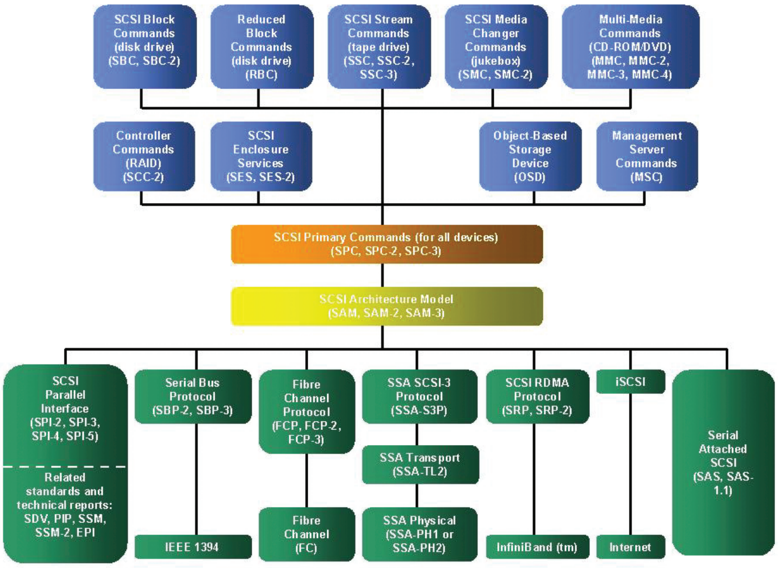

Were they still SCSI-3 standards or SCSI-4 standards? T10 decided that the '-3' added no value. So, SCSI-3 Parallel Interface (SPI) was followed by SCSI Parallel Interface-2 (SPI-2), and the SCSI-3 Block Commands (SBC) was followed by SCSI Block Commands-2 (SBC-2), etc.

There is not really an SCSI-3 standard. If anything, SCSI-3 is the whole family of standards shown in *Figure 1*. Also, there is no SCSI-4 and probably never will be.

SCSI Architecture Model

The SAM standards, shown in yellow in *Figure 1*, deserve special attention. They define what it means to be SCSI. Some people interpret the SAM standards as 'standards for writing SCSI standards' because they define services that must be provided in the other SCSI standards. However, the SAM standards are not only relevant for the SCSI standards writers. People who implement SCSI products will find that the SAM standards contain requirements not repeated elsewhere.

The standards shown in *Figure 1* conform to the SAM. There are some other standards and specifications that resemble SCSI but do not comply with the SAM. They leverage parts of the SCSI command sets, but, for a variety of reasons, do not comply completely with the SAM. Examples include ATA Packet Interface (ATAPI) devices and Universal Serial Bus (USB) storage devices.

SCSI Primary Commands

The SPC standards (orange) include commands that are mandatory or optional for all SCSI devices. Like the SAM standards, all SCSI devices must conform to an SPC standard.

Parallel SCSI products should use SPC-2 and SAM-2, the last generation of these two standards to support the unique requirements of the parallel SCSI bus. Serial SCSI products may use SPC-2 and SAM-2, but will probably find that SPC-3 and SAM-3 are more suitable for serial interfaces.

Transport Standards for Interfaces

The first SCSI-3 transport standards were SCSI-3 Parallel Interface (SPI) and SCSI-3 Interlocked Protocol (SIP). (Both these standards have since been withdrawn.) While these standards conformed nicely to the ideals of layered standards, it quickly became obvious that there was not going to be another transport protocol standard to use with SPI in place of SIP. Since SPI and SIP were always used together, it made sense to merge them. Starting with SPI-2, the parallel SCSI physical interface and the parallel SCSI transport protocol were merged. This also made SPI-2 much easier to read than the separate SPI and SIP standards. This is a case of the divide and conquer strategy going too far.

SAS has followed the SPI-n lead and also used a single standard for both the physical interface and transport protocol.

In the cases of Institute of Electrical and Electronics Engineers, Inc. (IEEE) Standard 1394, Fibre Channel, InfiniBand™ and iSCSI, it makes sense to divide these standards because the physical interface standards were developed by separate groups of people from those who developed the transport protocol standards for SCSI.

The T10.1 task group elected to have three sub-layers (protocol, transport and physical) in Serial Storage Architecture (SSA) standards. While this partitioning made sense at the time, it is unlikely to be repeated.

SCSI Command Sets Overview

The first SCSI command set was for disk drives at a time when disk drives contained less than 30 megabytes and transfer rates were about one megabyte per second. A lot has changed since then, but software written for those old disk drives would still work with current SCSI disk drives. This remarkable feat comes from the abstract view of storage that was embodied in the SCSI block device command set. The SCSI block device command set dealt with logical blocks rather than cylinders, heads and sectors. It presumed that the disk drive had enough intelligence to manage defective media transparently to the system. This meant that as disk drives evolved in performance, capacity and capabilities, the system software did not have to change.

A lot has changed in the last 20+ years and system software has also made some changes. One change was increasing the size of the logical block address field (twice) to support the enormous capacity increases. (This is one good reason for not using that 20-year-old software.) Nonetheless, the stability of the SCSI command sets is a major part of the value of SCSI. The SCSI command sets are supported by all modern operating systems and have been mapped to almost all input/output (I/O) interfaces.

SCSI-2 included command sets for disks, tapes, optical disks, CD-ROMs, printers, scanners, medium changers and communications devices. Some of these command sets were not migrated to the SCSI-3 generation of standards and several more command sets were added.

SCSI Command Sets

The command sets that survived to the third generation of SCSI are described briefly here. Since SCSI-2 is still an approved standard, there are products that continue to use the SCSI-2 command sets. Of course, using the newest applicable standard is always recommended.

Disk Drives

SBC and SBC-2 define the full-featured command sets for disk drives. Reduced Block Command (RBC) is a reduced command set used mostly with SBP-2 and IEEE Standard 1394.

Tape Drives

SSC and SSC-2 define the command sets for magnetic tape drives. SSC-3 has just started development.

Media Changers

SMC and SMC-2 support medium changers (i.e. juke-box devices).

CD-ROMs and DVDs

MultiMedia Command (MMC) support optical storage devices through MMC-4.

Redundant Array of Independent Disks Controllers

SCC-2 is a command set for storage controllers such as redundant array of independent disks (RAID) devices. (SCC was withdrawn after SCC-2 was published.) Unfortunately, most RAID products still use proprietary command sets.

Enclosure Services

SCSI Enclosure Services (SES) and SES-2 support enclosure service devices. These devices manage power supplies and cooling fans, etc. SES was based on an industry specification called SCSI Accessed Fault Tolerant Enclosures (SAF-TE), but SES supports more capabilities and does not need a separate device address.

Object-based Storage Devices

Object-based Storage Devices (OSD) is a command set for a still more intelligent view of storage devices. One of its goals is to create operating system independent file systems.

Protocol Bridge Controllers

The Management Server Commands (MSC) is a management command set for SCSI protocol bridge controllers that convert between various SCSI protocols like SPI-n, Fibre Channel Protocol (FCP-n), SCSI Remote Direct Memory Access Protocol (SRP-n) and iSCSI.

SCSI Standards Organisations

The INCITS T10 Technical Committee is responsible for most of the SCSI standards. IEEE is responsible for IEEE 1394. The INCITS T11 Technical Committee is responsible for Fibre Channel. The InfiniBand Trade Association is responsible for InfiniBand™. The Internet Engineering Task Force (IETF) is responsible for iSCSI and the Internet.

Conclusions

The SCSI Standards Architecture is a flexible and powerful tool. Command sets are migrated easily to new I/O interfaces, even the Internet. The divide and conquer strategy permits the SCSI architecture to evolve sensibly without impacting parts that are stable. ■

References

InterNational Committee for Information Technology Standards (INCITS): http://www.incits.org

T10 Technical Committee: http://www.t10.org

T11 Technical Committee: http://www.t11.org

IEEE: http://www.ieee.org

InfiniBand Trade Association: http://www.infinibandta.org

Internet Engineering Task Force (IETF): http://www.ietf.org