NETWORK COMPUTING

TCP/IP vs. OSI

Or How I Learned to Stop Translating and Love Standards

By KEVIN L. MILLS

Predominant opinions are generally the opinions of the generation that is vanishing.

— **Disraeli**

For five years in the late 1970s and early '80s, I worked on numerous government contracts as an employee of System Development Corp. I made a good living designing systems to connect one type of computer to another over a wire.

I helped Digital Equipment Corp. PDP-11s emulate Honeywell Inc. VIP terminals; I faked numerous IBM Corp. 370 systems into believing that a PDP-11 was an IBM cluster controller; and I cajoled Honeywell Level 6 minicomputers into masquerading as IBM 3704 front-end processors.

Of course, although they worked, these solutions were never perfect. But who minded an anomaly here or a disconnect there? Two incompatible devices could now communicate — most of the time.

Or at least until something changed in some component: a new operating system release, a new protocol version, a more efficient but plug-compatible printed circuit card.

The important thing was, I had a job: a never-ending job, an intricate job (somewhat like solving a crossword puzzle), a job that made money for my company, a job that cost the taxpayers countless millions of dollars, over and over again. In fact, you might say I was enrolled in a welfare program for data communications protocol translators.

In 1982 I signed on with the National Bureau of Standards, now the National Institute of Standards and Technology, which was just beginning a program to establish a suite of computer networking standards so that computers of any manufacture, built independently, could exchange data without translators.

The new data communications protocols would be built into the vendor's commercial products, maintained by the vendor and sold to all customers. At some point, every computer installed in the government would be able to exchange information, as needed, without commissioning translators. The government intended to end the welfare program for translators.

In the 10 years since, two contradictory trends have occurred: The government has nearly achieved its objective of establishing a standard protocol suite, and the original problem has worsened considerably.

Postponing the government success story until later, first consider how the original problem of protocol cacophony has multiplied.

During the 1970s state-of-the-art computer networks consisted of hundreds of terminals connected over coaxial cable, or by modems, to front-end processors and then to mainframes. How did these terminals communicate with the computer? Using master-slave protocols: 3270 Bisynch, VIP, Uniscope, Poll-Select. Later SDLC and HDLC, and then X.25, along with X.3, X.28 and X.29, were invented.

These protocols were installed in thousands of systems worldwide. Of course, in universities and research institutes, a new generation of peer-to-peer protocols was being developed to enable computer-to-computer communications among equals. These protocols remained largely a curiosity because the number of computers deployed was small compared with the many thousands of terminals.

Two singular events, the result of two decades of innovation in electronics, occurred early in the 1980s: the Sun Microsystems Inc. workstation was introduced, and the IBM PC was announced. Each of these introductions established a de facto standard: the desktop scientific and engineering workstation and the professional, personal computer.

While this raging torrent of low-cost computing spread across the work place, healthy developments also were occurring in the supercomputer market. A number of Cray Research Inc. machines were released in many configurations. Several Japanese companies also introduced supercomputers.

And attention began shifting to parallel supercomputers comprising hundreds and sometimes thousands of microprocessors working in concert. More and more, large research laboratories and universities bought and installed several supercomputers.

By the late 1980s the proliferation of PCs created a market demand for local-area networks and the associated protocols, so-called LAN operating systems, to enable the exchange of data among PCs within an office and a building. Numerous solutions reached the marketplace: Novell Inc.'s Netware, Banyan Systems Inc.'s VINES, Microsoft Corp.'s LANmanager and more.

Thus the market is simply repeating

the 1970s, but this time, instead of terminal-to-host protocols, the cacophony consists of LAN operating systems and other proprietary, peer-to-peer protocols.

How did this latest cycle of diverse protocols occur? First, the computer user began to get more (much, much more) for less (much, much less): i386 computers for less than $2,000 and i486 machines for less than $3,000.

This led to the availability of vastly improved desktop software: Microsoft Windows, IBM's OS/2, Word, WordPerfect, dBase, Lotus 1-2-3, Excel and so on. The user has become king: Processing at the desk is facilitated with these marvelous tools.

Of course, the user needs to access data from mainframes and to share data with other users on other PCs. This leads to demand for networked PCs. Several fast, agile companies supply solutions. Now divisions, departments and sometimes individual users buy from among the available offerings.

And all this happens quickly, without the knowledge or oversight of a corporate information resources manager, because the price to buy the necessary equipment is too small to require corporate consent. Before anyone realizes, the organization has made a major, uncoordinated investment in an information infrastructure that is sprawling, diverse, complex and incompatible.

Will This Madness Ever End?

The answer is yes and no. Yes because there already exist two solutions today that enable multivendor data communications across the full range of computer equipment: PCs, workstations, minicomputers, mainframes and supercomputers. These solutions are the so-called Transmission Control Protocol/Internet Protocol suite of protocols and the Open Systems Interconnection.

At least one of these solutions, OSI, is on a long-term evolutionary path to enable strategic, corporate investments that will continue to provide multivendor data communications even as hardware, operating systems and applications evolve.

The answer is also no because these standard solutions, TCP/IP and OSI, will never embody every feature of the latest, hottest new network wrinkle that users might employ — at least not until the data network is viewed as a utility akin to the electric power grid or the voice telephone network.

So organizations need to invest in a strategic solution, as the government is, that will allow multivendor data exchange today and in the future.

At this point, I have identified two strategic solutions that are available today, but I have not recommended a specific solution. Before I do so, I need to describe what TCP/IP and OSI are, to account for the growing popularity of TCP/IP and to explain how OSI is positioned for long-term success.

TCP/IP and OSI: Sibling Rivals

I think TCP/IP and OSI can best be viewed as sibling rivals, both belonging to the family of protocols aiming to provide data exchange among networks of heterogeneous computers.

Of course, one is older, and one is prettier, and one is stronger, and one is smarter, and one is better educated, and one is richer. But still, all in all, each accomplishes the same thing.

You can attach your computers to a network, be it a WAN or a MAN or a LAN or a DAN; you can route information in packets or datagrams between networks, and your computers and their routers can find each other, even if they are moved occasionally; you can send data, as a reliable stream or as best-effort messages, between software processes on various computers. You can transfer files or log in remotely to a computer, and you can send mail electronically.

Despite the striking similarity in function between OSI and TCP/IP, some significant differences exist. TCP/IP has made more substantial progress in implementing network management features, while OSI is working toward more robust, and more complex, management specifications.

TCP/IP benefits from its long life because a whole host of proprietary products have been put to work using TCP/IP networking to distribute services across a network; for example, Structured Query Language access to relational databases, network file services allowing remote mounting of file systems and remote windowing for bit-mapped graphics displays.

For most of these added features, OSI provides equivalent services that can be used by the same third-party software vendors to provide the same added features over OSI networking. In addition, for some of the features, such as SQL access and windowing, standard

SEE TCP/IP, PAGE 34

TCP/IP,

FROM PAGE 32

specifications are being developed to integrate the services into the OSI architecture in a standard manner.

On the upside for OSI, existing applications are being enhanced, and new applications are being developed to provide additional user services. Message handling system (also known as X.400) applications soon will provide standard security and directory services, along with the ability to interchange electronic data using an electronic data interchange user agent.

The largest advances for OSI over TCP/IP will come from new applications services. Directory service (X.500) applications will enable retrieval of information from locally maintained directory servers distributed throughout a network. Remote database access (RDA) products will extend SQL access across a network of heterogeneous databases.

The distributed transaction processing (DTP) service will provide synchronized transactions distributed across a set of network nodes.

OSI and TCP/IP solve the same basic problems by using different technical specifications. Like a pair of bickering siblings, advocates for each approach stress one's strengths and the other's weaknesses, but neither talks very much about their common heritage, nor about a few other facts.

TCP/IP and OSI will coexist, will interoperate and will complement each other in sometimes surprising ways. Today on Internet, a global network of largely TCP/IP-based computers and routers, OSI is used routinely.

The OSI protocol for routing packets (CLNP) is deployed in a significant and growing segment of the Internet infrastructure. A number of gateways exist for interoperation of TCP/IP mail (simple mail transfer, or SMTP) with X.400. TCP/IP file transfers (FTP) occur routinely on the Internet, as do OSI file transfers. Finally, some of the earliest pilots for X.500 are being conducted on Internet.

Both TCP/IP and OSI have a lot to learn about upper-layers architecture and services. TCP/IP uses an antiquated upper-layers protocol encoding technique that is inferior to the OSI solution, ASN.1. TCP/IP uses well-known addresses for connecting to network services, whereas OSI relies on a directory of names to find the address for a needed service.

OSI forces an arbitrary three-layer structure (session, presentation and application) onto the upper layers, creating built-in inefficiencies and making certain operations, such as encryption, more difficult than necessary. Neither architecture provides the desired flexibility to construct new application services by combining existing, refined or newly defined components into a bundle of cooperating objects.

Both TCP/IP and OSI have deficiencies in system-level issues such as security, multicasting and multimedia. TCP/IP-related specifications (called requests for comments) are being developed for privacy-enhanced mail, and RFCs also are under consideration to provide security services for network management and routing. In addition, Kerberos, a secret key authentication, integrity and confidentiality system, has been developed at the Massachusetts Institute of Technology under Project Athena and is being deployed in portions of Internet.

For OSI, standards are under development for authentication, confidentiality and integrity at the network, transport, link and application layers; deployment of solutions is several years away.

OSI has a rich set of multimedia capabilities embedded in the electronic-mail standard, while TCP/IP is just developing such extensions for SMTP. Neither TCP/IP nor OSI standards have given much consideration to real-time multimedia services.

Both TCP/IP and OSI application services face increasing competition from LAN operating systems, from other proprietary protocols and from developing consortia solutions such as those endorsed by the Open Software Foundation, Unix International and X/Open.

Why Is TCP/IP so Popular?

Why, then, is TCP/IP, with its numerous deficiencies, so popular? The overriding reason is that TCP/IP answers an immediate, almost desperate, need — data communications in heterogeneous networks — very well.

Also, TCP/IP is bundled into the Berkeley Software Distribution (BSD) Unix, which provided the fundamental operating software for the Sun workstation — the de facto standard in the fastest growing computer market segment. Thus, TCP/IP became the minimum networking capability for any vendor entering the market for scientific and engineering graphics workstations.

Do not overlook Internet, a linkage of perhaps 5,000 networks and 1 million computers. TCP/IP is the *lingua franca* of Internet; if you are a researcher in a corporation, university or government lab, then an Internet connection is almost a necessity of modern business.

Understand, also, that BSD Unix, and thus TCP/IP, has been used by students in universities, especially throughout the United States, for about a decade; therefore, TCP/IP is understood by many users, systems integrators and developers.

The connection between TCP/IP and BSD Unix is a result of U.S. government funding to the University of California, Berkeley, and other universities. In addition to subsidizing directly the development of TCP/IP, the federal government continues to subsidize TCP/IP indirectly through programs such as the National Science Foundation Network (NSFnet), the NASA Science Internet (NSI) and the Energy Sciences Network (ESnet).

The results of such subsidies are freely available to commercial suppliers, who then can implement TCP/IP networking without a major investment in protocol software development. With an immediate revenue stream, suppliers can concentrate resources on improving the usability of their TCP/IP networking products.

Why Is OSI Inevitable?

If TCP/IP answers the need for data communications among heterogeneous computers, then why is OSI an inevitable replacement for TCP/IP?

OSI is the accepted international standard for data communications. As such, OSI is specified for use by a growing number of governments around the world: the European Community legislates OSI, the U.S. government mandates OSI (and the states are following) and the Commonwealth of Australia has adopted OSI, as have Japan, Taiwan and the Nordic countries.

OSI also is accepted by other groups with international scope: the World Federation of MAP/TOP User Groups. One reason for the acceptance of OSI is that the standards are created and evolve in an open process, visible to users and suppliers throughout the world.

The program of work is organized and scheduled so that plans can be drawn for developing and deploying solutions that use the resulting standards. In addition, OSI standards are augmented by a rigorous testing process that improves the quality of OSI products and aids in managing the evolution of change.

Apart from the process and political reasons for OSI's inevitability, there are technical reasons as well. The OSI application services provide increased functionality over those provided by the TCP/IP applications: SMTP, FTP and remote log-in.

X.400 provides an extensible framework for carrying information of all kinds, not simply personal mail messages. The OSI virtual terminal service supports more than simple character or line terminals: forms, page and scroll modes also are supported.

X.500 is far more capable than the equivalent TCP/IP centralized directory service. And future OSI standard application services have no direct TCP/IP equivalent; recall DTP, RDA and the manufacturing messaging specification.

Closer to the network, OSI also provides enhanced capabilities over TCP/IP. For example, the TCP/IP address space encompasses 32 bits and is rapidly approaching exhaustion, while an OSI network address comprises 160 bits, a size that will provide global addressing into the foreseeable future.

In addition, the routing protocols used with TCP/IP are constrained by the flat, 32-bit address so that the routing tables maintained in Internet switching nodes are growing quite large and, thus, becoming unwieldy.

OSI routing protocols support a form of hierarchical routing so that address information can be represented more efficiently in summary form, reducing the amount of routing information that flows in the network and that must be stored in the switching nodes.

Internet will be the first large beneficiary of the advantages of OSI routing. Today OSI CLNP services and routing protocols are implemented in the NSFnet backbone, and CLNP with static routing is available in several of the regional networks attached to NSFnet.

Government networks, NSI and ESnet, also are prepared to forward OSI datagrams. OSI routing services will become implemented more fully over time in Internet, and, with luck or proper planning, OSI switching services will provide a natural transition path on Internet as the TCP/IP address space limits are reached.

The federal government, then, had a hand in two successful solutions to the problem of incompatible, heterogeneous computer systems. The government role in TCP/IP was one of direct, federal subsidy for a solution, while the government role in OSI was one of collaboration with industry to develop voluntary standards and then of endorsing the results as FIPS 146, the Government OSI Profile. And, of course, the government uses both solutions.

Practical Plans for Practical People

So what solution should you use: TCP/IP or OSI? In general, if you are installing a new network or acquiring new data communications services, then you should specify and implement OSI as your standard protocol for multivendor information exchange.

Where you have specific requirements that go beyond the capabilities available in OSI products today, you should augment OSI with other network protocols as needed to meet such additional requirements; usually this means accepting proprietary solutions.

In solicitations, you should make clear your intent to reduce proprietary enhancements over time, and you should require the offerer to specify plans for including additional OSI services into products as OSI specifications continue to mature.

There may be instances where procuring TCP/IP products is sensible. For example, suppose you already have deployed a large TCP/IP network, and you are procuring new systems to add to that network; then certainly the systems you buy should be capable of running TCP/IP protocols.

But, if the procurement is of significant size, then the systems should be purchased with a dual-stack capability so that they are bilingual: TCP/IP and OSI. At the same time that you procure your bilingual computers, you should upgrade the routers in your deployed network so that they are capable of routing both TCP/IP and OSI data.

Also, you should ensure that some of your newly acquired, bilingual computers (often called dual-suited hosts) include software to relay between TCP/IP and OSI applications.

Sometimes, even when you are procuring a new network, installing TCP/IP together with OSI might make sense.

For example, suppose you are acquiring a large network of routers, servers and workstations, but you need to integrate some older existing computers into the network.

Often TCP/IP implementations exist for older computers for which no OSI implementation exists and for which no OSI implementation is planned. In such circumstances, TCP/IP might provide the means of integrating the older computers into your new network.

Here, the migration path is straightforward: Procure routers (often called dual-suited routers) capable of switching both OSI and TCP/IP data and add some number of dual-suited hosts with application gateways. Your new network then will support information exchange between your old existing computers and your newly procured, OSI-capable equipment.

I recommend that you procure OSI products when you are installing a new network or making a significant upgrade to an existing network. I recommend that you buy TCP/IP, in addition to OSI, only when the network you are upgrading is already a TCP/IP network or when TCP/IP provides the only means of integrating older, existing computers into your new network.

Acquiring TCP/IP alone makes sense only when you are buying a single computer or a few workstations to connect to an existing, large, TCP/IP network such as Internet. Even here, because Internet is adding support for OSI coexistence and interoperation with TCP/IP, procuring OSI in addition to TCP/IP probably makes good sense. ◀

Mills is chief of the Systems and Network Architecture Division at NIST's Computer Systems Laboratory.