SNA Perspective

Volume 12, Number 11
November, 1991
ISSN 0270-7284

The single source, objective monthly newsletter covering IBM's Systems Network Architecture

Whither the 3174?: Part I

It is almost impossible to discuss SNA without mentioning the IBM 3174 establishment controller—the most ubiquitous of all SNA physical units, the heart of the 3270 Information Display System (IDS), and nursemaid to the three most abundant SNA logical units. The 3174 is a significant communications product; no longer merely a terminal controller.

This article, the first of a two part series, briefly reviews the origins of the establishment controller (formerly called the subsystem control unit and cluster controller) and examines the standards developed for the 3270 IDS. We then focus on new key 3174 feature options—token ring support, upstream and downstream support, multihost support, ESCON attach-ment, PC and asynchronous support, coaxial, twisted-pair, and fiber-optic links, and NetView interface. Finally, we examine three major direct competitors—IDEA Courier, McDATA, and Memorex Telex—and their positioning and value-added. The second article in this series will delve further into the role to be played by establishment controllers. It will also examine the impact on the 3174 of the IBM 3172 interconnect controller, LAN communication servers, and gateways.

(continued on page 2)

Impact of Frame Relay and SMDS on SNA

Frame relay and switched multimegabit data service (SMDS) are both targeted for high speed LAN interconnection within the emerging enterprise internetwork. Both offer a radical departure from traditional, wide area networking technologies and are generating a great deal of interest in the networking community. Frame relay is now available from several carriers, including WilTel, US Sprint, MCI, BT TYMNET, and US WEST. The first announcement of a trial SMDS service on the east coast was made at INTEROP 91—Bell Atlantic will introduce SMDS for a flat rate of $500 per month for T-1 (1.544 Mbps) access. RBOCs and other carriers will also be bringing their offerings to market in the next six months.

Frame relay and SMDS offer SNA users some new, valuable alternatives as well as some challenges. This article explores the possible impact on SNA and the opportunities offered by these new technologies.

(continued on page 12)

In This Issue:

(continued from page 1)

Humble Origins

Surprisingly, the designers of the original IBM 3270 Information Display System (3270 IDS) did not have SNA in mind. Rather, they wanted to improve on the 2848/2260 combination introduced, with System/360, in 1964.

2260 Terminal

By today's standards the 2260 was a primitive terminal. It had an immovable keyboard and a very limited display capacity (240, 480, or 960 characters). It was connected to the 2848 Display Control, which had very limited magnetic-core memory that enabled it to support a model-dependent maximum of eight 960-character, sixteen 480-character or twenty-four 240-character displays. Also, depending on the model, the 2848 could be attached to a System/360 channel or, via a modem and communication link, to a channel-attached 2701, 2703, or 2703 transmission control unit (all hard-wired).

2848 Controller

The limitations of the 2260 were such that buffer storage, keyboard handling, and character generation were all handled within the 2848. The 2265 display/controller was subsequently introduced to address the needs of the single-user remote site.

The interchange code used between the host and the 2848 was ASCII (despite IBM's resistance to its use on all other fronts, favoring EBCDIC). It was also transmitted high-order bit first (rather than the conventional—then as now—low-order bit first), as defined by the IBM Type III protocol.

The architecture of the 2848/2260 subsystem was constrained by the limitations of the time. Memory was very expensive and very bulky. Also, since the integrated circuit was not yet invented, the 2848's logical circuits were space and power hogs.

The placement of the 2848 between the 2260 and the host did have its advantages. The same 2260 display terminals could be used at the host site (with the channel-attached 2848) or at a remote site. Only the 2848 changed. This made the 2260 a fairly flexible asset and also allowed users at the host site to take advantage of a data transfer rate of 2600 characters per second (considered fast for the time). Users at remote sites were not so lucky, with a data transfer rate of 600 bits per second.

The First 3270 Information Display System

The first 3270 IDS borrowed much from the 2848/2260 subsystem. As before, the display terminals (3277s) were largely dependent on a controller. Again, the differences between local and remote connection lay in the controller (3271 for remote, 3272 for channel-attached). Interestingly, keystrokes were handled within the display terminal, unlike the later 3278, 3279, and other control unit terminals (CUTs). Character generation was wholly within the display. Connection to the controller was now by a coaxial cable at a data rate of 1 Mbps. Printers like the 3284 could also be attached.

Although the 3277 Model 1 preserved the 480-character display of the middle 2260 model, the Model 2 set what was to become an industry standard, with its 24-line, 80-column display format. This standard was adopted by the makers of the first ASCII display terminals (aka glass teletypes).

Two models of the 3271 cluster controller were eventually acknowledged by the designers of SNA. Models 11 and 12 supported the SDLC protocol and were subsequently classed as physical unit (PU) type 1. Models 1 and 2 used the binary synchronous communication (BSC) protocol.

The 2265 also had a successor in the standalone 3275, a BSC controller/display terminal with an optional 3284 printer.

Competitors

Unlike its predecessor, the 3270 IDS had its imitators. Memorex and Telex (competitors in those days but now merged) offered full plug-compatibility—they sold terminals such as the 1377 which could attach to IBM's cluster controllers. Subsequently, these terminals could attach to their own cluster controllers (e.g., 1371 and 1372) which were marketed as additions to or replacements of existing 3271 and 3272 controllers.

Courier (now IDEA Courier), Lee Data, and Braegen offered compatible subsystems, but not plug compatibility. Compatible terminal/controller subsystems together emulate the 3270 system but the individual terminals and controllers were not interchangeable with IBM models. This was a costly decision for Courier and Lee Data and a fatal one, eventually, for Braegen.

Protocol Converters

The fast-developing installed base of ASCII display terminals also created a market for 3270 protocol converters. These converters, which appeared to the remote host as 3271 controllers, made the attached ASCII terminals look, as far as their keyboards and display attributes would allow, like 3277s. Datastream and ICOT were early and quite successful entrants in this market. (Datastream was acquired by Lee Data. ICOT has since acquired of INS, one of many offering 3270 emulation products, both standalone and LAN-based.)

The early 3270 IDS left plenty of room for improvement—the 3277's use of LEDs as status indicators was very limiting; its internal keystroke processing required that unique models be produced for each national language; lowercase was available only as a request for price quotation (RPQ) feature; and the coaxial cable length was limited to 600 meters.

The SNA Generation

IBM's second-generation 3270 IDS introduced the Category A coaxial connection, relegating the old coaxial connection, in the process, to Category B status. The new 3274 cluster controller, available in SNA and non-SNA versions, both local and remote, supported adapters for both connection types (with a Category B maximum terminal count half that of Category A). The cable length was increased to 1500 meters and the data rate to 2.35 Mbps.

The Category A protocol made the display dependent on the controller for keystroke processing (as was the old 2260). Nothing was displayed except by controller action. This apparently retrograde step simplified national language support. The now-customizable controller (which included an eight-inch flexible disk drive) was provided with a language diskette. This gave the user a choice of language which determined, for example, whether the key to the left of the S would represent a Q (French or Belgian AZERTY keyboards) or an A (most other keyboards). For IBM, language-to-language display terminal differences were largely a keycap exercise, although some languages (e.g., Danish) required a revised character generator.

One of the reasons for the higher data rate was the need for the controller to process every keystroke, with no perceptible delay for the operator. Like the Category B protocol, the Category A protocol polled or addressed every attached terminal (display or printer) in turn (also called round-robin fashion). This meant that, between keystrokes at a single display terminal, up to thirty-one other displays or printers had to be polled or addressed, with data transfer occurring when required.

Category A display terminals, the first being the 3278, were freed from the 24x80 screen format limitation. Model-dependent formats were 12x80, 24x80, 32x80, 43x80, and 27x132 (which could be switched to 24x80). In all cases, there was an additional status line (the operator information area) which was able to convey more information than the previously-used LEDs.

Renewed Competition

The success of the new 3270 generation was reflected not only in renewed activity by the plug-compatible manufacturers, but also in the appearance of a number of companies offering microcomputer-based emulation products. The catalyst for this competition was, of course, IBM's belated entry into the PC business and its creation, thereby, of a de facto standard. Technical Analysis Corporation (TAC), which was quickly acquired by DCA, introduced the IRMA board with its E78 3278-emulation software; CXI (later acquired by Novell) introduced PCOX; Forte (eventually also acquired by DCA), Attachmate, and others followed suit. IBM's response was interesting, as discussed in the section, IBM's Emulation Entry, below.

Distributed Function Terminals

One of the most intriguing products designed for attachment to the 3274 cluster controller was the IBM 3290 Information Panel Display Station. Not only was it striking to look at, with its orange gas plasma flat-panel display and 122-key keyboard, but it also introduced a completely new mode of operation. With its capacity to handle all controller functions except the link or channel protocols and to support up to five concurrent host sessions, the 3290 was the first distributed function terminal (DFT). To distinguish the mode of operation of 3278s and 3279s, IBM redesignated them control unit terminals (CUTs). Not surprisingly, the 3274 required considerably revised software (configuration support D) to handle the 3290, including the ability to download software to the terminal at power-up time.

IBM's Emulation Entry

IBM's initial response to the emulator vendors was not a PC-based coaxial adapter with emulation software but the 3270 PC, whose software could be customized for either CUT- or DFT-mode operation. Its DFT operation was limited to four display sessions. Only days later, the company announced an availability date for the IBM Emulator Adapter board and 3278 emulation software for its other PCs. (The Emulator Adapter was the same board used in the 3270 PC, which had first priority on the board's initial production.)

Important Standards Established

Between the introduction of the 3274 and the 3174, several standards were established. Because these standards are reflected in large user investments in host software, they continue to be important.

The most important of these standards are:

- The 3270 data stream, supported by very specific code in host application subsystems.
- The extensions to the 3270 data stream (extended data stream), including IBM's approach to 7-color support, extended highlighting, and structured fields.
- SNA character string (SCS), consisting of control codes used to format a visual presentation medium—typically a printed page. Used most often for support of LU 1 printers.
- Programmed symbols, available as options on the 3278, the 3279 Models 2X and 3X, and the 3270 PC, and standard on the 3279 Model S3G and the 3290. Many host graphics display applications generate programmed symbols directly (i.e., not through generalized graphics subsystems such as IBM's GDDM).
- APL keyboard and display support (APL2 on DFT- and certain CUT-mode terminals).
- Vector graphics, using "drawing orders" contained in structured fields, as implemented on the IBM 3270 PC/G, 3270 PC/GX (and PC-AT variants of these), and 3179 Model G (and, subsequently, on the 3192 Model G and 3472 InfoWindow Graphics-5 (3472-G)). IBM vector graphics terminals also accept host orders to redefine function keys.
- Intelligent printer data stream (IPDS), used with advanced dot matrix and laser printers. Requires structured fields.

It is interesting to note that, except where CUTs are involved, all of these standards are transparent to the controller.

The 3274 set high standards for response time, even when it was not channel-attached. Given an adequate link data rate (up to 56 kbps), the avoidance of large printer data streams, and the relegation of graphics terminals to their own controller, 3270 display terminal users have become accustomed, in many cases, to subsecond responses.

You've Come A Long Way

The 3174 was first introduced with the designation subsystem control unit, later changed to establishment controller. This renaming ties in with IBM's current terminology for enterprise systems, with an establishment being a location or site within an enterprise.

The 3174's key feature options now include, but are not limited to:

- Support of token ring, both upstream (e.g., local 3745 attachment) and downstream (token ring gateway) supporting downstream PUs (DSPUs).
- Support of up three direct 3270 datastream host links (one channel attachment and up to two remote communication links).
- Support of ESCON channel attachment.
- Single-link multihost support.
- Asynchronous Emulation Adapter (up to three, each with eight ports), allowing connection of ASCII terminals and/or connection to ASCII hosts, with operation at up to 19.2 kbps.
- 3270 port expansion feature, increasing the total number of Category A devices supported to 64 (from the old maximum of 32).
- Support of coaxial or twisted-pair connections to Category A terminals.
- Support of 3299 Model 3 or Model 032 via coaxial, twisted-pair, or fiber-optic link.
- ISDN Basic Rate Interface Adapters (up to two, three, or four, depending on the 3174 model), each supporting up to eight remote PU 2.0 devices at 64 kbps.
- Node Type 2.1 (NT2.1) support, allowing the 3174 to operate as an APPN network node for devices on the token ring gateway or to support peer communication for Category A terminals.
- Response time monitor.
- Operation as a NetView entry point.

The evolution of the 3174 and the current 3270 IDS is depicted in Figure 1.

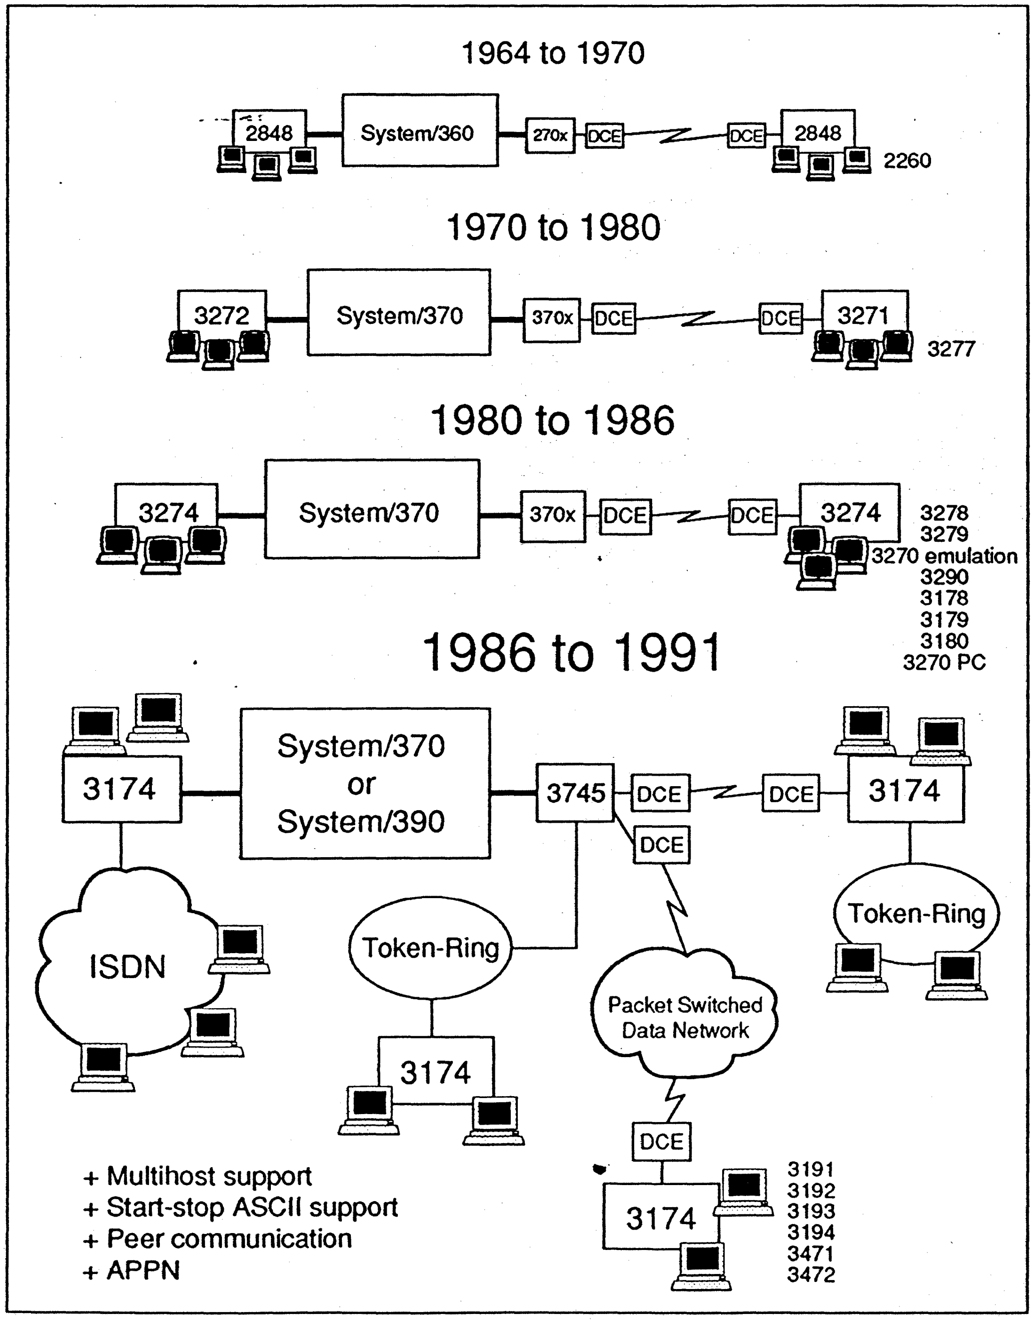

Figure 1

Feature Comparison of IBM 3174 and Competitors

Feature/Capability	IBM 3174	McDATA 7100	Memorex Telex 1174	IDEA Concert
Maximum Category A devices	64	128	96	64
Maximum start-stop ASCII devices (including hosts)	24	34	32[1]	32
Maximum host channel connections	1	2[2]	1	1
ESCON channel attachment				
Direct	Y	P	N	P
Via converter	Y	Y	NI	N
Maximum remote host connections	3[3]	4[3]	6[4]	4[5]
Maximum IBM hosts				
Via ESCON	8	TBA	NS	8 (P)
Via X.25 connection	4[6]	8	16	5
Via token ring	8	8	16	250
Via Ethernet	NS	8	NS	NS
Maximum overall	16	16	16	Unspec.
Maximum data rates				
Channel normal (Mbyte/s)	1.25	1.25	1.75	1.25
Channel data streaming (Mbyte/s)	2.5	3.0	NS	4.5 (P)
SDLC half-duplex (kbit/s)	64.0	128.0	64.0	64.0
SDLC full-duplex (kbit/s)	64.0[7]	64.0	64.0	64.0[7]
BSC (kbit/s)	19.2	19.2	19.2	NS
X.25 (kbit/s)	64.0	64.0	64.0	64.0
Start-stop to ASCII terminal or hosts (kbit/s)	19.2	19.2	19.2	19.2
Coexistence of remote connection link protocols				
SDLC and BSC	N	Y	Y	N
SDLC and X.25	Y	Y	Y	Y
BSC and X.25	N	Y	Y	N
SDLC, BSC, and X.25	N	Y	Y	N
Communication interfaces				
EIA-232-D (CCITT V.24/V.28)	Y	Y	Y	Y
EIA-449/422/423 (with DB-25 connector)	N	N	N	Y
CCITT V.35	Y	Y	Y	Y
CCITT X.21 switched	Y	N	N	N
CCITT X.21 nonswitched	Y	Y	Y	Y
Gateway support				
Token ring maximum DSPUs	250	240	240	250[8]
Group poll support	Y	Y	Y	Y
Ethernet maximum DSPUs	NS	240	NS	NS
Group poll support		Y		
DEC LAT connection (via Ethernet)	N	Y	N	Y
APPN support (via token ring)	Y	P	N	P
Peer communication support (Category A terminals)	Y	P	N	N
Local format storage	Y	Y	Y	N
APL2 support in conjunction with appropriate CUT[9]	Y	P	Y	N
Intelligent printer data stream	Y	Y	Y	Y

Table 1

Strong Competition

IBM cannot be all things to all people, so there are some obvious deficiencies in the 3174's list of features. Some are significant for those who have made large investments in Ethernet and/or Digital Equipment Corporation (Digital) systems; others are significant for the very large SNA-only network user.

IBM has three major direct competitors for the 3174. Two—McDATA and IDEA Courier—also address non-IBM connectivity and interoperability issues. The third, Memorex Telex, hews closely to the IBM line, while offering greater capacity at a competitive price. McDATA's 3270 line includes the controllers, while both IDEA Courier and Memorex Telex also offer terminals and terminal emulation products.

All three provide the 3174's essential features, including token ring host connection, token ring gateway with DSPUs, support of CUT mode, DFT mode, and start-stop ASCII terminals, terminal multiplexers, and NetView. All three claim improvements over IBM in several areas, including customization and the maximum number of LUs supported.

In addition to the features discussed below, Table 1 provides a comparison of some of the major features of the IBM 3174, McDATA 7100, Memorex Telex 1174, and IDEA Courier Concert.

IDEA Courier

IDEA, a nine-year-old company with headquarters in Billerica, Massachussetts, was an early entrant into the 5250 terminal emulation business (System/36 and AS/400 connection). In 1988, it acquired Alcatel Information Systems (formerly ITT's Courier Terminal Systems division) and thus entered the 3270 controller and terminal marketplace. The Courier division of IDEA is located in Tempe, Arizona.

Its current controller product line reflects IDEA's 5250 expertise. The IDEA Concert can be configured as a superset of the IBM 5394, when it is called the Concert 394 and supports connection to up to four AS/400s. Alternatively, it can be configured as a 3174 superset—most 3174 features plus Digital VAX interoperability via Ethernet, using Digital's Local Area Transport (LAT) protocol. Combined 3174/5394 operation is planned for next year, with internal 3270/5250 protocol conversion. This would eliminate the need for 3x74 Remote Attach on the AS/400 and also provide the functionality of a 5x94 Remote Attach on the mainframe.

The Ethernet/LAT support is fairly straightforward, with no apparent exploitation of Ethernet beyond the Digital connection. Figure 2 (page 8) illustrates connection to one IBM host and one Digital host.

The IDEA Concert does not support BSC and IDEA Courier has no stated plans to offer such support.

McDATA

This Colorado company, also nine years old, was formerly the OEM manufacturer of controllers for both Memorex and IDEA Courier. Overseas,

Feature Comparison of IBM 3174 and Competitors (continued)

NI = No information
NS = Not supported
P = Planned
TBA = To be announced

1 Each group of eight reduces the number of coaxial devices by eight.
2 Both Linkmaster 7100 channel connections can be SNA, or one can be non-SNA, or a combination of SNA and non-SNA. For a non-SNA-only connection, the 7100 can appear as four control units, allowing up to 128 sessions. With the 7100's MLT support, a user can hot-key between an SNA session and a non-SNA console session.
3 Minus the number of channel connections on local models.
4 Each X.21, X.25, or full-duplex SDLC link reduces the maximum by 1.
5 Reduced to 2 if there is also a host channel connection.
6 Eight hosts may be accessed if X.25 is the primary link.
7 Full-duplex link available in conjunction with token ring gateway only.
8 May be constrained by other configuration options.
9 In IBM's case, 3191 Models D, E, and L, and CUT version of 3192.

Table 1 (cont.)

McDATA still provides the controllers sold by such major European-based companies as Nokia and Siemens-Nixdorf.

McDATA competes with the IBM 3174 on two fronts. First, it offers features that are directly comparable to IBM's, but with extended capacity. Second, it offers features that are not offered by IBM in a manner transparent to the IBM host software.

Although IBM has made a number of Ethernet announcements in the last two years (ES/9370 direct attachment, 3172 and 3745 (September 92) attachment, etc.), it has yet to provide, for Ethernet users, the SNA support enjoyed by token ring users.

McDATA's Ethernet support, at its most basic, parallels IBM's token ring support. Just as the 3174 can communicate with one or more hosts via a token ring on which one or more channel-attached 3745s reside, the Linkmaster 7100 (McDATA's latest 3174 replacement) can communicate with one or more IBM hosts via an Ethernet LAN on which one or more channel-attached Linkmaster 6100 or 7100 controllers reside. The 7100 can provide access to ASCII hosts via the same Ethernet LAN, using the LAT protocol, for both CUT-mode terminals in ASCII emulation mode and ASCII terminals. (The Linkmaster 6100 is an IBM 3172 Interconnect Controller replacement. Just as the 7100 offers capabilities not offered on the IBM 3174, so the 6100 offers more features than the 3172. In this article, we only consider its role as a device for connecting up to six Ethernet LANs to one or two IBM host channels.)

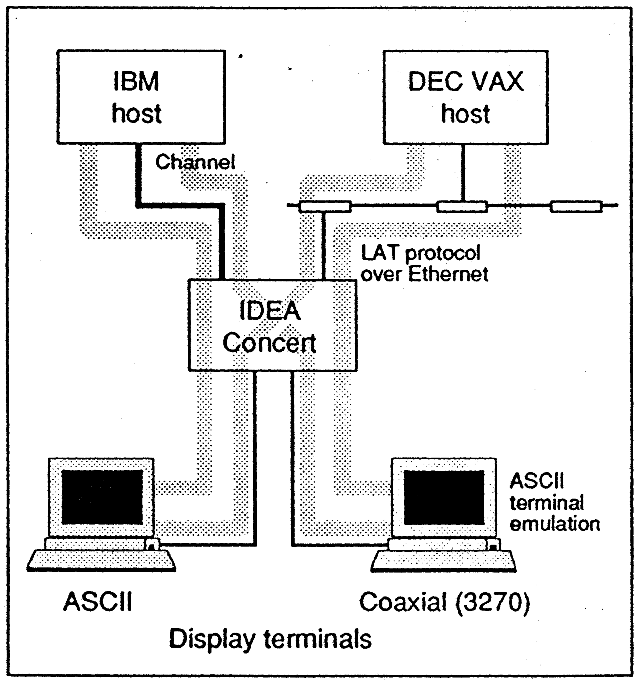

Figure 2

McDATA's latest release includes some very impressive capabilities, which are too numerous to discuss in this article. Figure 3 illustrates the use of Ethernet for both DSPU and gateway support, in this case providing terminal users on all three 7100s with access to both IBM hosts, plus any Digital hosts which may also reside on the Ethernet. If Digital hosts are not a consideration, the same IBM multihost support can be accomplished with the 7100's token ring support.

McDATA's real non-SNA strength is at the channel attachment level, allowing a non-SNA connection to one host and SNA to the other. Better yet, the 7100 allows mixed SNA and non-SNA operation on a single channel connection, using two subchannel addresses. This allows a console operator using multiple logical terminal (MLT) support, for example, to switch between the non-SNA console session and an SNA production session. The 7100 also allows eight controller addresses per remote BSC link, allowing the addressing of 256 logical units.

Memorex Telex

The Memorex Telex 1174 provides the greatest number of host connections and the greatest number of controller appearances (sixteen) per remote link. It is also the only match for IBM's APL2 CUT-mode support, the dual language feature, and explicit partitioning. Memorex addresses interoperability issues with its LANSYS product, which is not discussed in this article.

Memorex Telex's enhanced non-SNA support is quite impressive, allowing a single 1174 to appear, on each BSC link, as 16 controllers. This would appear to allow host addressing of up to 512 LUs (based on the limitation of 32 per non-SNA controller). But in fact, up to 576 LUs can be supported, presumably via 18 perceived controllers on two or more links. Moreover, the 576 units may be placed in a contention pool (dynamic host connection), accessible to CUT-mode and ASCII terminals.

Is Interoperability Enough?

Using a single display terminal, whether 3270 or ASCII, to communicate with both IBM and non-IBM (e.g., Digital) host applications can be a significant challenge.

For many users, especially programmers, switching terminal personalities when switching hosts is not an issue. For many others, it pays to hide all host systems completely. For such users, interoperability requires more than a dumb terminal.

IBM's Systems Application Architecture (SAA), with its common user access (CUA), is intended to make applications appear the same to the user, regardless of the type of system on which they reside. With CUA, at least at the enhanced graphics workstation level, terminal personality is not an issue.

The Open Systems Interconnect (OSI) model addresses interoperability but does not address the user interface. Large users are demanding (if not buying) OSI conformance, a fact not overlooked by IBM. IBM is supporting, or will support, an OSI-compliant interface below SAA's common programming interface for communications (CPI-C) for each of SAA's platforms.

The Future of the 3174

To return to the 3174, this product line has evolved significantly from its early days of dumb block-mode terminal support. But does the ability to interoperate 3270 and ASCII terminals with IBM and Digital hosts across Ethernet and token ring LANs still represent a short-term solution? If so, where does the 3174 go next?

IBM has added NT2.1 capability and the competitors will follow suit. Does this make the 3174 merely a transitional device, supporting existing terminals as the world of the advanced workstation takes over? After all, NT2.1 can run on many platforms. Will the CUT, the only remaining link with the antiquated 2260, disappear in five years? In ten years? The answer will depend on how heavily committed users are to existing applications and how aggressively they pursue open systems.

In a future article, we shall delve further into the role to be played by establishment controllers, the IBM 3172 interconnect controller and compatible products, and LAN communication servers in the world of LANs, SNA, LAT, TCP/IP, and open systems. ■

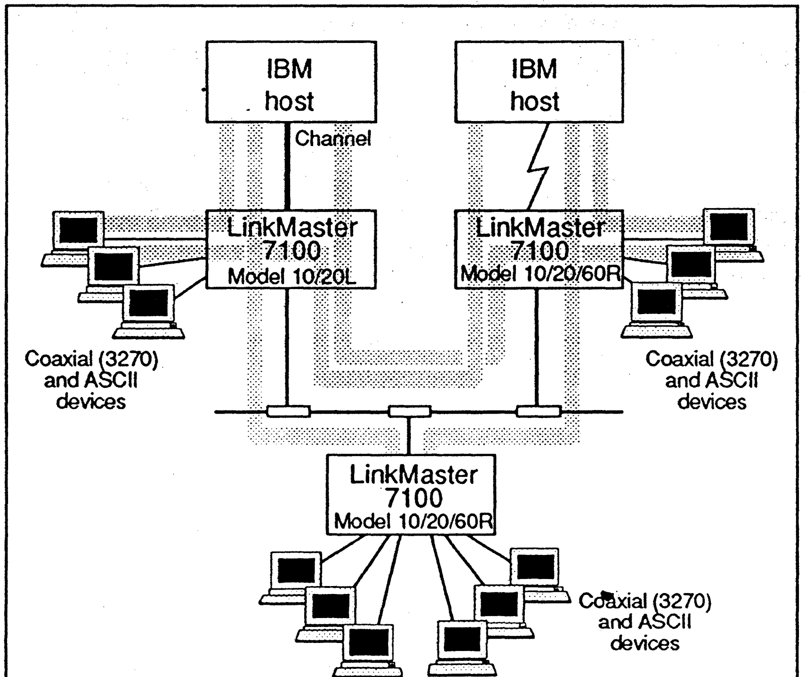

Figure 3

Readers may contact these companies at the following addresses:

IDEA Marketing Department
P.O. Box 29039
Phoenix, AZ 85038-9039
Phone: (602) 894-7000

McDATA Corporation
310 Interlocken Parkway
Broomfield, CO 80021-3464
Phone: (303) 460-9200

Memorex Telex
6422 East 41st Street
Tulsa, OK 74135
Phone: (918) 627-1111

3270-Related Definitions

Asynchronous Emulation Adapter (AEA)
A feature of the 3174 establishment controller, providing full-duplex, start-stop mode communication at up to 19.2 kbps to up to eight ASCII display terminals, printers, and/or host systems. Depending on the model, the 3174 can accommodate up to three AEAs. The following logical connections are supported:

- ASCII display terminal or ASCII printer to ASCII host
- ASCII display terminal, emulating 3270 display terminal (with 24x80 or 32x80 screen format), to IBM host
- ASCII printer, emulating IBM 3287, to IBM host
- 3270 display terminal, concurrently emulating up to five ASCII display terminals, to up to five ASCII hosts (mixed IBM and ASCII host sessions may be supported on the same 3270 display terminal, to a maximum of five)
- 3270 printer (3287 Model 1 or 2, 3262 Model 3 or 13, 4224 Model 201 or 202, or 5204 in 3287 emulation mode), emulating an ASCII printer, to ASCII host

Control Unit Terminal (CUT)
A terminal whose keystrokes are processed and whose presentation space (device buffer) is managed by the cluster controller to which it is attached. A CUT is capable of supporting only one host communication session at a time. (However, the IBM 3174 establishment controller will allow the operator of a CUT to switch among up to five host sessions, all of which are maintained by the controller.) Examples of CUTs are the IBM 3178, 3179, 3180, 3278, and 3279. IBM 3270 series printers (e.g. IBM 3287) are also CUTs.

Distributed-Function Terminal (DFT)
A Category A coaxially-connected device that does not require cluster controller interaction to respond to keystrokes. A DFT terminal may contain one or more SNA network addressable units, typically LU 1, 2, or 3. Communication between a DFT and a cluster controller is in the form of EBCDIC message blocks.

DFT Extended (DFT-E)
A DFT extension that allows the display terminal to be used for controller customization (normally only accessible by a CUT), provides access to an ASCII host session, and supports X.21/X.25 and controller diagnostics.

IBM 3172 Interconnect Controller
A microchannel-based intelligent controller, used to attach token ring, MAP 3.0, or Ethernet local area networks (LANs) to the channel of a System/370 or System/390 host processor. May also be used for remote channel-to-channel connection over a T1 link (1.544 Mbps). The 3172 Model 2, which uses an Intel 80486 processor, also supports the 100-Mbps Fiber Distributed Data Interface (FDDI). Host channel connection may be made by System/370-type parallel channels or ESCON fiber-optic channels.

IBM 3174 Establishment Controller
An IBM 3270 IDS controller, capable of supporting up to 32 (64, with the expansion feature, on some models) Category A coaxially-connected devices (in CUT mode or DFT mode, single and multiplexed), token ring devices, and up to 24 ASCII terminals. Depending on the model, the 3174 may communi-cate with the host system via a parallel or ESCON channel attachment (at 1.25 or 2.5 Mbyte/s), a remote communications link (up to 64 kbps) or via token ring (4 or 16 Mbps). Remote attachment options include CCITT V.24/V.28 (EIA-232-D), CCITT X.21, and CCITT V.35 interfaces, and SNA/ SDLC, CCITT X.25, and BSC operation. May also communicate in start-stop mode with an ASCII host.

IBM 3270 Information Display System
A terminal system consisting of cluster or establishment controllers, keyboard/display terminals, printers, etc.

IBM 3270 Workstation Program
An IBM software product that provides much of the capability of the IBM 3270 PC on a variety of IBM PCs, including especially IBM PS/2s.

3270-Related Definitions *(continued)*

IBM 3274 Cluster Controller
A 3270 series cluster controller capable of supporting Category A coaxially-connected devices (in CUT mode or DFT mode, single and multiplexed). Depending on the model, the 3274 may communicate with the host system via a channel attachment (at 640 kbyte/s), or a remote SDLC or BSC communications link. Early models of the IBM 3274 can also support Category B coaxially-connected devices. It has been replaced by the 3174 establishment controller.

IBM 3278
A monochrome CUT. Although it is no longer manufactured, the 3278 is used as a reference point in describing the functionality of many 3270 terminal emulation programs.

IBM 3287
Coaxially-attached dot-matrix printer. Used as the reference point for host-addressable non-IPDS printer support in 3270 emulation products, even where the printer actually used is other than a dot-matrix.

IBM 3299 Terminal Multiplexer
A multiplexer for up to eight (Model 2 or 3) or up to 32 (Model 032) Category A devices. May be connected via coaxial or twisted pair cable to the terminals and to the terminal adapter of an IBM 3174, to a multiplexed port on an IBM 3274 (Models 2 and 3 only), or to the multiplexed port of another 3299. The 3299 Model 032 also supports a fiber-optic connection to the 3174.

IBM 3708 Network Conversion Unit
A device providing concurrent line concentration, protocol conversion, protocol enveloping, and ASCII pass-through for start-stop ASCII devices. ASCII terminals can communicate with a host system in either native mode or as full-screen 3270 displays or printers. ASCII terminal, ASCII host, and SNA/SDLC host connections operate at up to 19.2 kbps. For ASCII to 3270 protocol conversion, its appearance to an SNA/SDLC host is the same as that of a 3274 Model 51C or 61C cluster controller.

IBM 3710 Network Controller
A device providing the following concurrent functions:

- ASCII (start-stop) to SNA 3270 protocol conversion
- BSC 3270 to SNA 3270 protocol conversion
- Asynchronous and BSC (RJE) protocol enveloping
- ASCII pass-through
- Concentration of ASCII, BSC, and SNA/SDLC traffic

Host communication is over one or more user-defined SNA/SDLC or X.25 links.

IBM 3745 Communication Controller
A successor to the IBM 3705 with low, medium and high capacity models. Depending on model and expansion options, supports from 16 to 896 communication lines and from 1 to 16 host system attachments. Some models are designed for remote operation as line concentrators.

Intelligent Printer Data Stream (IPDS)
IBM's approach to host-initiated printer data streams using advanced printers with multifont and graphics capabilities. It uses structured fields to send data and commands to a printer independently of the attachment protocol and, therefore, of the system to which the printer is attached.

Multiple Logical Terminals (MLT)
Support for multiple logical units via a single terminal port on an IBM 3174 establishment controller, 3274 cluster controller or equivalent. Depending on the level chosen when the controller is customized, can support up to five display sessions on a CUT (3174 only), or up to five display and/or printer sessions, in any mix, on a DFT.

Operator Information Area (OIA)
A non-data area occupying the last row on a display terminal. Used to provide error and status information. Often referred to as the status line.

(continued on page 20)

(continued from page 1)

The Emerging Technologies

Some points will be covered here to set the stage for examining the relationship of frame relay and SMDS to SNA later in this article.

Both of these technologies evolved to provide better LAN interconnection services for enterprise internetworks. Their important characteristics include:

- Higher access speeds
- Bandwidth on demand
- Multicast operation
- Savings
- Best-effort delivery

Higher Access Speeds

Frame relay carriers offer speeds up to T-1 with T-3 (45 Mbps) expected within the next year. SMDS will also offer T-1 and T-3 with SONET (155 Mbps and above) planned for the mid-1990s.

Bandwidth On Demand

A particular network session can utilize varying amounts of bandwidth as needed without prior signalling with the network. For example, fractional T-1 allows a session to use varying amounts of capacity; however, the host must signal the change to the network and wait while channels are reassigned. SMDS uses an access approach which is designed to ensure fairness among all equipment sharing the access link. Users can have all the bandwidth they can get within these fairness rules. Frame relay uses the concept of committed information rates to monitor usage. If a particular session uses more than the committed bandwidth, the excess frames are marked for discard if congestion occurs.

Multicast Operation

Frame relay and SMDS support multicast operation whereby any attached system can send one frame to all other members of the multicast group. Multicast operation extends the LAN across a wide area, allowing remote clients and servers to interact as if they were both on the same LAN. Systems are able to join and exit from different groups as needed. Without this feature, application designers must account for local or wide area communication and take different approaches for each.

Savings

Frame relay and SMDS are switched services. This reduces the number of lines needed for full connectivity. In addition, a switched topology scales well since new sites can be added at the same incremental cost. New tariffing also offers a potential for substantial savings. Most vendors will offer a flat rate that will be substantially less than current services.

Best-Effort Delivery

Both frame relay and SMDS offer a best effort at reliable delivery. A frame may occasionally be discarded because of congestion or transmission errors. The networks do not inform the attached systems; instead, they continue normal operation. This absence of error correction in the network simplifies network design. It also reduces processing overhead which increases throughput. However, the impact is that this overhead does not disappear; is merely moved back to the end points. The attached equipment is responsible for detecting lost data and recovering it through retransmission.

How Does SNA Fit?

Both frame relay and SMDS offer the SNA user some interesting new alternatives. One of the most exciting is the option of creating a single backbone "highway" for SNA and other traffic (see Figures 4 and 5 on page 13). Many organizations already have substantial non-SNA internets, usually based on TCP/IP or IPX (Novell), or which may be based on OSI in the future. Usually two separate transports are used, one for SNA and the other for the remaining traffic. Building a single backbone consolidates expensive transmission links and leverages economies of scale. However, the single backbone must be able to provide adequate support for all of the protocols using it.

Economies of Scale

For example, T-3 costs are approximately four to nine times those of T-1 links for twenty to thirty times the speed. Consolidating several T-1 links into a T-3 trunk will increase capacity and may not add much cost. The advantages of consolidating management functions and increasing the interconnectivity of the entire organization are also attractive

Simplification

The traditional internet backbone operates with an autonomous set of routers that make their routing decisions through exchanges with neighboring routers. This strategy has simplified network initialization and configuration since routers learn from each other and adapt to changes in the topology.

Systems operations are also simplified—since the discovery of a local router is all that is necessary, all internet traffic is sent to a known router which takes responsibility for delivery. This flexibility provides a high level of resilience and reliability for the backbone.

Drawbacks

On the other hand, there are some drawbacks:

- Convergence
- Variable throughput
- Error recovery

Complex internets can take a long time to converge to a new routing environment when a topology change occurs. Variable throughput can also adversely affect applications and users who depend on a stable response time for their activities. This is particularly true of synchronous architectures such as SNA. Best-effort delivery also means that some data may be discarded and will require end-to-end recovery across the internet, which again affects SNA networks.

SNA over Multiprotocol Routers

Most non-SNA internetworks use bridges or routers to connect LANs to the backbone. The same approach can be used for SNA sites—internetwork

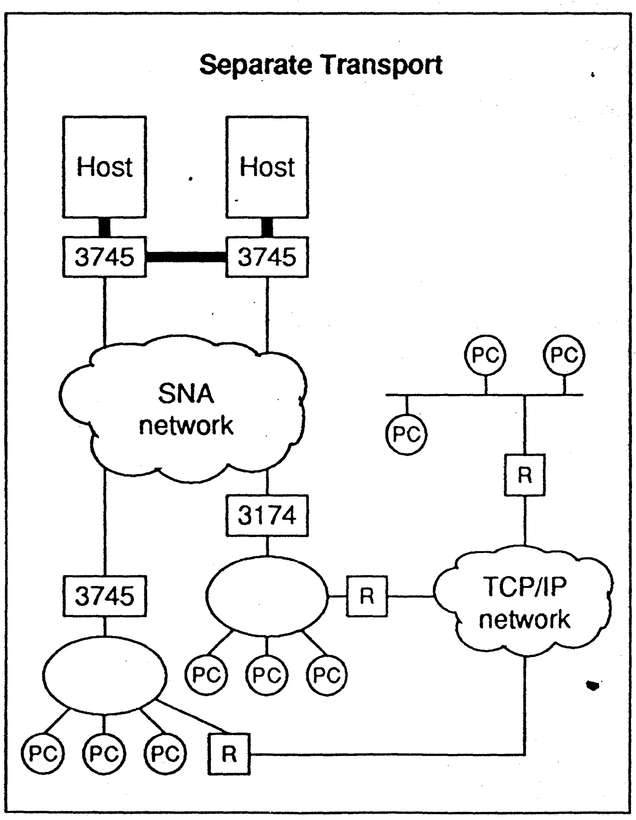

Figure 4

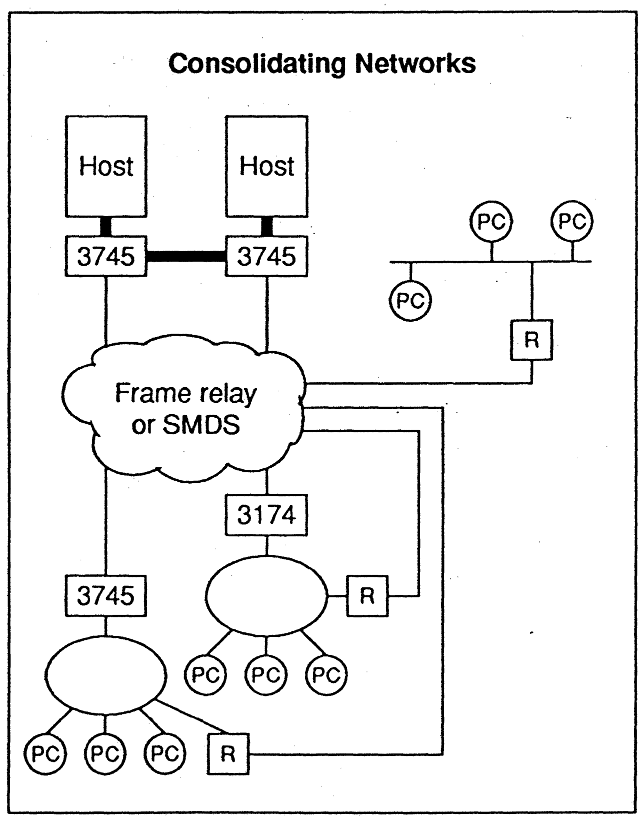

Figure 5

devices can be interposed between the internetwork backbone and SNA devices such as establishment controllers, communication controllers, and LAN interconnectors. This approach uses tunnelling to move SNA message units across the backbone. It requires no changes to the current SNA devices (as shown in Figure 6). Most bridge and router vendors have announced frame relay and SMDS support for their products as well as SNA/SDLC connections.

Concerns

There are several issues with sending SNA traffic over multiprotocol routers:

- Polling
- Timeouts from variable delays
- Network management
- Lack of SNA experience

One issue here is the polling traffic that SNA requires. Most multiprotocol router vendors will likely provide local "spoofing" at each end to eliminate polling traffic on the backbone and to avoid the timeouts that a variable delay backbone can cause (see "Serial Tunnelling: Spoofing Your SNA Network," *SNA Perspective*, October 1991). These vendors are likely to enhance the level of SNA support in their routers over time, incorporating selected SNA routing capabilities directly into the bridge or router. Most vendors have eschewed PU 4 to take a safer path, using the more open and stable node Type 2.1 and PU 2 interfaces.

Using these products may simplify or increase the complexity of enterprise management, depending on the user's environment. These products take advantage of the years of experience in the multiprotocol routing technology deployed today, but the concern for users is the level of SNA expertise that these vendors can deliver. SNA is complex and fundamentally different from TCP/IP, IPX, and other communication architectures. Most internetworking vendors are seeking strategic partnerships with firms who have the necessary SNA experience.

SDLC Tunnelling and Session Routing

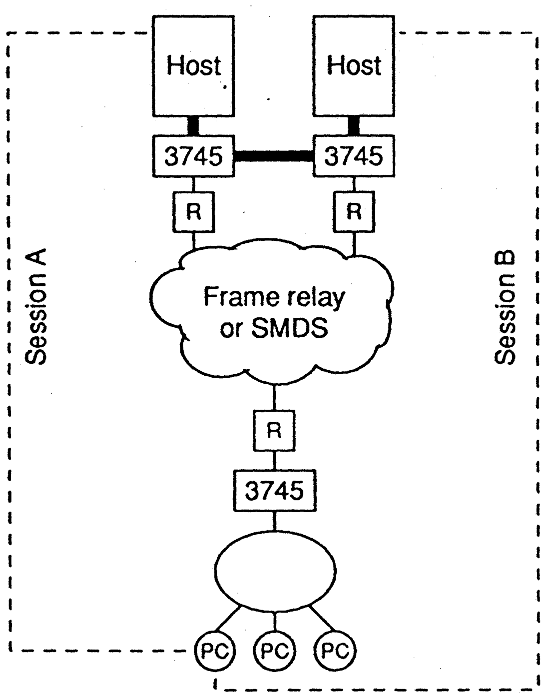

Figure 6

Session Level Routing

Internetworks are intended to allow direct communication between partners without any intermediaries. A desktop client should be able to access any server within the internetwork by sending messages directly to it. Of course, access control requirements such as passwords are used to limit access to authorized clients.

SNA, on the other hand, currently requires cross-domain services since a communication controller controls all session establishment procedures for those devices it controls. All session traffic flows to the device's owning communication controller and is then relayed to the desired target. Cross-domain services extend the reach of SNA devices at the expense of time and resources. Each communication controller in the path consumes resources for the relay function. Although this approach was designed to support network scalability, managability, and reliability, one impact is increased network traffic.

The directness of the new internetworking technologies can be exploited to reduce this overhead (see Figure 6). For example, a product from Netlink, the SNA_Hub, allows a PU 2 device to specify its intended target host before establishing a direct session. The device may change target hosts between sessions and may even carry on simultaneous direct sessions with different hosts across the internetwork.

Frame relay and SMDS will be good foundations for direct session routing solutions which will become more important in the future.

APPN

APPN looms large in some users' strategies. A multiprotocol router could contain the appropriate software to act as an APPN network node. In this capacity, it could participate in the routing operations of other APPN network nodes for SNA traffic. This same router could move other protocol traffic across the same internet backbone. This approach is considered safer than PU 4. First, APPN specifications are more open than hierarchical SNA routing with PU 4, though IBM has yet to announce the result of its "reconsideration of our ability to publish APPN network node" discussed in June 1991 by Ellen Hancock, IBM Vice President of Network Systems.

However, this solution will not help those organizations using subarea networking, since IBM has also not yet presented a complete strategy for integrating APPN and subarea networking. APPN will be most attractive to those users who have large numbers of token rings or IBM midrange systems.

Extending the Hop Count

Token ring interconnection has been limited by the seven-hop-count limit in IBM's source routing bridges. Such a hop count limits the scale and scope of larger logical token ring networks being designed.

New internetworking products can help since many bridges from other vendors, using transparent bridging instead of source route bridging, can support more than two token ring interconnections directly, reducing the number of hops between rings. Frame relay and SMDS can also be used to provide a wide area backbone between token rings without consuming the precious hop count. Since there is only a single hop across such a backbone, it becomes easier to create national or international scale logical token ring networks.

SNA and Frame Relay

Users will deploy private frame relay networks in several ways:

- New frame relay backbone
- Integrating frame relay into existing backbone
- Public frame relay access

Some may start from scratch, buying frame relay switches and building their own backbone. Others will add frame relay to their current enterprise network. Both packet- and circuit-switching vendors have added frame relay interfaces and software to their products, allowing users to introduce frame relay into their current environment. Others will use a public service which entails attaching a router to the frame relay network port. The router will be responsible for collecting LAN traffic and forwarding it to the frame relay network as well as distributing frame relay traffic to the appropriate target.

Frame Relay for the 3745

IBM has responded by announcing a frame relay data terminal equipment (DTE) interface for the 3745. This allows two 3745s to use the frame relay network as an SNA tunnel. SNA data units are encapsulated in frames and sent across the network where they are decapsulated. Encapsulation means that the 3745 cannot communicate with "native" frame relay systems sharing the same network. That is, IBM is not providing a full data circuit-terminating equipment (DCE) interface, which would allow the 3745 to be a frame relay router.

Shipment of the 3745 frame relay interface is scheduled for September 1992 when NCP version 6 is released. IBM has stated that this NCP version will accommodate frame relay characteristics.

Handling Error Recovery

Most of the NCP changes will be focused on the path control network. Traditional SNA networks operate as highly reliable, stable delivery services.

Data units are checked at each relay point and are retransmitted if they are corrupted. Thus, the NCP has been designed to expect such characteristics. Frame relay networks, by contrast, perform only simple error checking: if a frame is corrupted, it is simply discarded. The new version of NCP must have additional capabilities for recovering from unreported discards. These modifications will, no doubt, be functionally similar to layer 4 transport protocols such as TCP or OSI Transport class 4.

Variable Pacing Windows?

Other NCP changes for frame relay can be expected in the future. For instance, bandwidth on demand may not be completely available with this first-generation product. Frame relay allows any session to take as much bandwidth as it can obtain, subject to the committed information rate constraints. The equivalent in SNA would be the selection of very large pacing windows at session establishment. This would also imply knowledge of whether a frame relay network is being used. When a frame relay network is used for tunnelling, its performance will be reduced to the pacing constraints of the SNA network.

Congestion Notification

Frame relay has congestion management notification bits which inform the receiver and sender that the network is experiencing congestion. NCP will need to include mechanisms for recognizing these conditions and responding to them. The sending NCP must slow down the rate of offered frames until the congestion notification is cleared. The receiving NCP can also assist by reducing the session pacing window which will also serve to reduce the frame volume over time.

Modifying Applications

Applications must be modified to take advantage of frame relay capabilities. Further, VTAM will need extensions that will allow applications to demand variable bandwidth as they carry out their tasks. Taking full advantage of frame relay would be possible if additional session establishment parameters were used to specify the committed information rate and multicast groups. Additional services would be needed so that hosts could join and drop different multicast groups.

Frame Relay Management

Many traditional SNA environments will question the introduction of frame relay, especially if it will support a multiprotocol backbone. Managing bridges, routers, and a variable delay network may present new challenges. If users are building a private frame relay network, then switch and routing control issues must also be addressed. Those using a public frame relay service must also consider the local management interface (LMI) which is used to exchange operational information.

Frame relay will be managed by the simple network management protocol (SNMP) of the TCP/IP world. Additional capabilities will be needed for NetView to manage all elements from an enterprise platform. LAN managers will be able to obtain operational status and statistics locally, but control will reside in NetView. This structure facilitates the introduction of SMDS, which will also use SNMP.

Communication Controller—A Likely Target?

The communication controller is the target of many bridge/router vendors. Its relatively high cost, old technology, and low performance make it an attractive target for newer bridges and routers. Internetworking products have a lower purchase price, use more current hardware and software, and offer better performance, at least for small and moderate-sized networks. Router reliability in large networks can impact overall performance.

However, the communication controller would not be easy to displace. Bridge/router vendors have not yet demonstrated their ability to integrate native SNA subarea routing (PU 4 node) into their products. This would prove to be a more difficult task if IBM continues with its annual succession of NCP changes, which would keep the internetworking vendors off balance.

A more attractive option may be to focus on the 3172 LAN interconnect controller as an alternative target. A channel-attached device that can be enhanced may allow alternative products to do an end-run around the communication controller. Companies with channel experience that are expanding into the multiprotocol arena may find an opportunity here. For example, McDATA of Broomfield, Colorado, has recently started promoting its gateway/server as a 3172-compatible product. For companies not wanting to directly tackle the channel, interfacing into the 3172 rather than the 3745 may be advantageous. Its high-speed LAN interconnection plus the openness of its interface make this a more attractive target for internetworking vendors to gain exposure and leverage at the mainframe level.

Low IBM Profile on SMDS

IBM has not yet indicated its plans for SMDS. It has kept a low profile lately where attention has been focused on SMDS. IBM participated in the SMDS Solutions Showcase at INTEROP 91, showing an RS/6000 with SMDS interfaces. This, of course, gives no indication of IBM's plans for its SAA mainframe, midrange, and desktop systems. IBM is certainly following developments and learning about the potential. Later, IBM will decide what it will do.

SMDS is attracting more attention as its advantages are recognized. High-speed access, eventual national and international coverage, cell relay, and isochronous support are important technical features that users are demanding. Security, closed private networks, and access control are also built into the SMDS environment. Advantageous tariffs and the relief of turning backbone management over to a carrier that can achieve economy of scale are also attractive.

Why IBM Downplays SMDS

IBM has downplayed SMDS in the past. Perhaps one reason is that SMDS is a public service provided by the regional Bell operating companies (RBOCs) and other carriers. If SMDS succeeds, IBM and other vendors would no longer have control of the enterprise backbone. The only opportunities would be in selling interfaces into the SMDS network. Many network management opportunities would also be lost to the carriers.

Conclusions

The arrival of new internetworking backbones is a fact. Frame relay services will be widely deployed by the end of 1992. Many carriers will be offering national and international frame relay service at T-1 or lower speeds. Higher T-3 speeds will be available in 1993, or sooner if the market drives them faster.

Concurrently, SMDS has moved into the market more quickly than expected. However, SMDS growth will be limited until the carriers define and implement standards for interconnecting SMDS networks into a worldwide service.

Both these technologies offer the SNA user new options and opportunities. Some may choose a frame relay network as a higher speed, lower cost backbone for a large number of sites. Others will want to develop a single, high-speed, multiprotocol backbone. In the future, pressure will increase to evolve SNA to accommodate bandwidth-on-demand and multicast operation.

IBM has been slow to respond to these new trends. It has been late announcing and delivering its first frame relay product. The much-discussed multiprotocol router based on the RS/6000 has not yet been revealed either, though IBM has stated it will begin limited shipments by year-end 1991. The multiprotocol router could be a key element in exploiting these new technologies. (This router may also be a safe IBM-sanctioned path for SNA users to consider for replacing certain remote communication controllers with a more modern communication engine.)

Many vendors are exploring opportunities in today's different SNA environment. They offer products that do not require a 3745 and extend basic SNA functionality such as direct session routing. They may be able to offer solutions that address other areas before IBM brings more products to market. These opportunists must also find the shelter of published specifications and niches that IBM is not expected to address soon.

These developments will accelerate the evolution of SNA to an internetworked, distributed foundation for future enterprise networks. Users will benefit from following frame relay and SMDS developments closely and taking advantage of new alternatives. ■

Architect's Corner

The Roaring 90's

by Dr. John R. Pickens

In a recent conversation about the future of networking, internet pioneer Dr. Vint Cerf likened the next ten years to the historical final decade of the 19th century—the "Roaring 90's." Considering the technological changes underway, especially in communications, this decade a century later will be as significant and tumultuous.

The coming decade is a decade of transition. For example, Vint postulates that considering the installed base of 600 million phones today, a forward-looking estimate of one billion networks (one network per phone—wireless perhaps) seems quite achievable. But do we have the technology (and routing algorithms) to support it?

After exploring this topic with Vint, I continued to contemplate and began to expand the analogy into the IBM environment. I placed myself in the year 2000, or in the year 2010, or 2020. What will the IBM internet look like? Will SNA still dominate? Will TCP be winning? Will OSI replace TCP? Will OSI replace SNA? Which protocol suite will win the performance race? How fast will LANs go? How many protocol layers will there be? Will applications "see" networks? Will networking technology in the future look like today's technology, only be faster and more pervasive?

I've thought a lot about these questions, and have recently come to some new conclusions. (Note: I expect controversy. . . .)

SNA vs. OSI vs. TCP

Let's start with the SNA versus OSI versus TCP question. Which will advance? Which will decline? Which will ultimately win?

None will.

They all will.

I now see no overriding reason for any of the three to become the single protocol of the future, even in IBM-dominated environments. TCP has too much de facto multivendor momentum. SNA has too much dominant-vendor momentum. OSI will rise, but does not have the power to displace either of the other two. (Other tributary protocols, such as Novell's IPX and AppleTalk will undoubtedly also continue, but drop beneath the noise level.)

I do believe that cross-fertilization will occur. Common security algorithms will be sprinkled across all three. TCP and SNA will tie into OSI/CCITT X.500 directory services. SNA will migrate to OSI-based and TCP-based (a concession) network management services. Common integrated routing mechanisms will be designed and implemented for all three (triple IS-IS, a speculative topic for a future Architect's Corner). The future promises greater integration of SNA/OSI/TCP. But in my opinion the three will continue to coexist independently in approximately equal proportions, both within the network and on the desktop.

TCP deserves special comment. I would note that TCP devotees no longer describe TCP as an "evolution" toward OSI, the political expediency of such a move having long since vanished. Rather TCP is described as a permanent, enduring protocol. To understand this, we only need to look at recent moves within the European ISO-related community toward embracing TCP (but, note, not at the exclusion of OSI). Or, if that's not convincing, consider the recent statement from Ellen Hancock of IBM, "It (TCP/IP) has become a standard in its own right and is no longer simply an interim step to OSI." (Keynote speech at INTEROP 91.) Nevertheless, while TCP's star has certainly risen, it too will be unable to completely dominate.

Actually, the question of which protocol will dominate is irrelevant—there isn't enough time for any of the three to achieve dominance. A significant paradigm shift is occurring. The signals are visible today.

The Ultimate (Internetworking) Paradigm Shift

Gigabit global LANs are on the horizon. Before the end of the decade it will be feasible to construct LANs with the physical diameters of the circumference of the earth's equator. Lobe lengths on the order of north/south polar longitudinal lines. Transmission rates on the order of gigabits. Desktop distribution transmission speeds at .1 to 1 gigabit rates.

At these proportions and dimensions, the traditional underlying physics of protocol engineering is tossed out the window. To explain:

OSI, SNA, and TCP all share a common design point—robust communications mechanisms that are adaptable to a wide variety of physical transmission environments. In other words, design for the lowest common denominator—slow link speed, high latency, high error rate. As a result, as the denominator migrates to high speed, low latency, and low error rate, all three protocols feature properties and characteristics that will not map forward, especially layers.

A favorite anecdote from recent protocol engineering milieux concerns the following question. How does one crank TCP (or SNA or OSI) up to subgigabit speeds (600 Mbps)? The answer? Simple. Two Crays. Fifty percent CPU utilization. Oh, and keep the distance between Crays very short.

To support the new generation of gigabit-based physical connectivity environments, new protocols are going to be required, e.g., XTP which is becoming an IEEE standard (under IEEE 802.4). New algorithms are going to be required, e.g., rate-based versus window-based flow control. Layers are definitely not going to be required. In fact, an extreme point of view would argue that the future design point for protocol design is workstation-based memory-to-memory transfers.

Intermediate Systems

What of the intermediate systems required to support these new gigabit multipoint switching fabrics? Hardware. The switches are going to be very hardware intensive. Early products supporting ATM, for example, use hardware end-to-end for the switching fabric. Software within the network will be required only for system administration and network management.

I'm not about to predict the final outcome for gigabit networking environments—ATM, frame relay, Y-networks, et. al.

The Final Resting Place for SNA/OSI/TCP

So, what becomes of SNA, OSI, and TCP in the next twenty to thirty years? They will continue to exist, but in a "third-world" kind of role. Tributaries, as opposed to the mainstream. SNA, OSI, and TCP products will become the "protocol convertor" targets of the next century. The networking 8086. The distributed equivalent to DOS. Many product opportunities will continue to exist—gateways between subgigabit SNA/OSI/TCP environments and the hardware-intensive gigabit switching fabric. Servers and SNA/OSI/TCP WAN-farms (like today's modem pools for digital PABXs). Even as end-systems migrate to this new switching fabric, there will continue to be a place for today's (lowest-common-denominator) protocol suites.

An Ongoing Role for the Experts

Now back to the present. Should the user worry today about the ultimate paradigm shift—the demise of OSI/SNA/TCP?

Not to worry—SNA, OSI, and TCP are very much alive. Much design and evolutionary work remains. And the transition period will be extremely elongated. Besides, in a world of a billion LANs, even 10 percent of ES/IS systems running OSI/SNA/TCP is a respectable installed base.

But, in this century's Roaring 90's, life will certainly be interesting and profitable. ■

3270-Related Definitions *(continued from page 11)*

SNA Character String (SCS)
Control codes used to format a visual presentation medium, such as a printed page. LU 1 printers use SCS controls whereas LU 3 printers print from a device buffer similar to that of a 3270 display terminal.

Structured field
For both displays and printers, a special field within the data stream neither destined directly for nor originating from the presentation space. May be used to download programmed symbols and vector graphics, and for the transfer of nonpresentation-space data (e.g., for file transfer or the control of non-IBM peripheral devices such as plotters), etc. IBM's IPDS makes use of structured fields.

Vector graphics
Graphics support characterized by the use of "drawing orders" which are used by the display terminal or graphics printer to create and display the desired graphics images. Drawing orders are not dependent on the geometry or other physical characteristics of the terminal device. With vector graphics, the display terminal operator can interact directly with the host graphics application to modify individual structures. This is in contrast to programmed-symbol graphics in which the image can only be displayed. ■

Copyright © 1991 CSI - Communication Solutions, Incorporated, all rights reserved. Reproduction is prohibited. • Subscription rates: U.S. - one year $350, two years $520. International - one year $385, two years $590 • *SNA Perspective* is published monthly by CSI, 2071 Hamilton Avenue, San Jose, CA 95125 • Telephone (408) 371-5790 • Fax (408) 371-5779 • Managing Editor: Louise Herndon Wells • Associate Editors: Vincent Busam, Alan Kucharski, Basil Treppa • Circulation Coordinator: Cheryl Roberts • Contributors: Bill Potts, John McConnell, Dr. John R. Pickens • Typesetting and illustration: Aaron Lyon at dSIGN • The information and opinions within are based on the best information available, but completeness and accuracy cannot be guaranteed.