Logical problem determination for SNA networks

by R. A. Weingarten
E. E. Iacobucci

Problem determination on a Systems Network Architecture network has dealt mostly with error detection on physical network components. Adequate logical error-detection mechanisms associated with the logical network (software-related) errors have been only recently provided with the announcement of a new on-line interactive package called the Network Logical Data Manager (NLDM). This paper discusses the physical and logical network environments, logical network problems, and functions provided by the two releases of NLDM for logical problem determination.

Systems Network Architecture (SNA)[1] can be viewed as consisting of two networks, a physical one and a logical one. The physical network is made up of a number of hardware nodes interconnected by links. The function of the physical network is to control the flow of user data to, from, and between the physical network nodes. The logical network consists of the management of protocols that support the exchange of user data. Although these two networks can be viewed separately, they interact closely since the logical network operates via the physical network.

In the past, problem determination has mostly dealt with the physical network. Error-reporting and error-recording mechanisms evolved from stand-alone diagnostic support on individual hardware nodes to on-line interactive diagnostic support of remote hardware nodes. Such IBM program products as the Network Problem Determination Application (NPDA)[2] have been developed that allow an SNA user either to receive alerts from the network nodes of pending or failed conditions or to query network nodes and intelligent modems[3] for status data.

SNA network logical problem determination capabilities have not evolved as rapidly. Until the announcement of the Network Logical Data Manager (NLDM),[4] the logical problem determination mechanisms were considered cumbersome and time-consuming.

NLDM was created to provide on-line support that enables the user, normally the network operator or trained diagnostician, to interactively display data about the logical SNA network for problem determination purposes. NLDM constantly collects and maintains status information about the SNA network. This information is presented in the context of logical conversations, called sessions in SNA, such that information is available leading up to the occurrence of a failure. The initial release of NLDM collects and displays session awareness and trace information about the session end points. The second release[5] expands this visibility to the network routes that are traversed by the sessions. Also included is support for the SNA Network Interconnection[6] environment.

This paper first describes in more detail the differences between the physical and logical networks and how they intersect. Next it discusses the difficulties associated with logical problem determination and how NLDM can be used to aid in identifying logical

Figure 1 Physical network

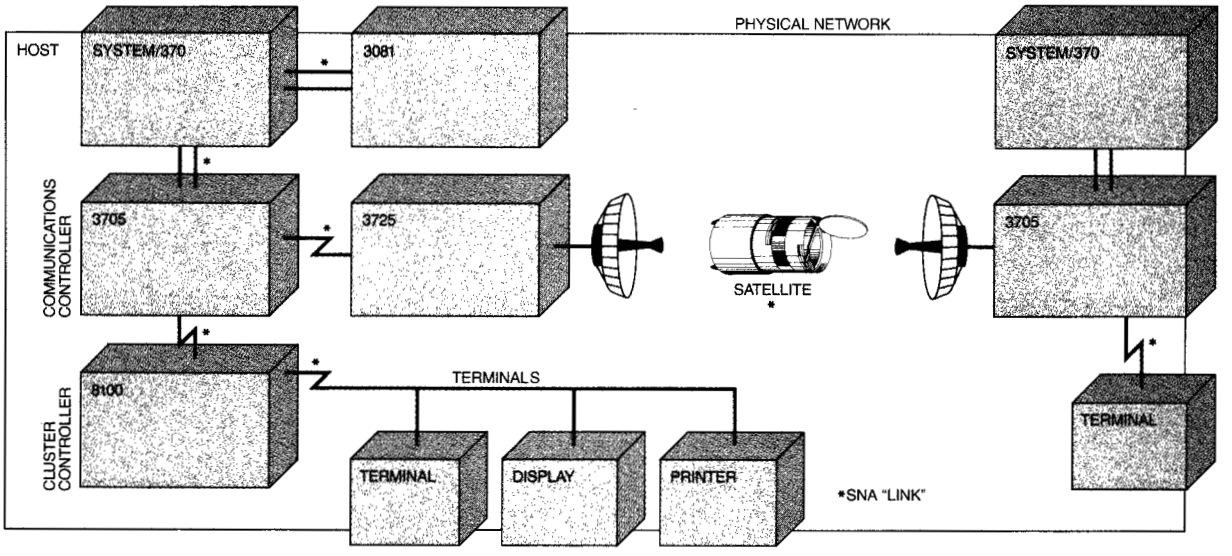

errors. Last, a discussion is given on the need to understand the mapping of the logical into the physical networks and how NLDM provides that support. Physical problem determination techniques and tools are not specifically discussed since they are the subject of numerous other papers.[3,7,8]

The physical and logical networks

The physical network. The SNA physical network consists of a number of hardware nodes interconnected by links. The hardware nodes are divided into classes according to their functions. Currently there are four types: host, communications controller, cluster controller, and terminal nodes. The terminal nodes encompass a considerable variety of devices such as video displays, keyboards, printers, cash dispensers, etc. Figure 1 shows how a physical network can be connected.

The functions and capabilities of each of these nodes are defined by SNA. Each of the physical nodes has associated with it an SNA function called the Physical Unit Services, which provides management functions for that node.

Understanding portions of the physical network is the responsibility of the System Services Control Point (SSCP). An SSCP has the responsibility for the network operator interface, communications network management interface, configuration control, network startup, network recovery, and participation in the creation of sessions.

The logical network. The logical network consists of entities called Network Addressable Units (NAUs) and the SNA sessions connecting the NAUs. An NAU represents a port through which end users may access the communications facilities. Although an NAU is a logical entity, it is assigned a unique network address which depicts its physical location in the network. This network address provides routing information so that user data can be directed to the correct physical node within the network. Three types of NAUs are defined in SNA: the system services control point itself, the physical unit, and the logical unit.

The logical unit (LU) is a set of function management services directly supporting an application program or terminal operator end user. The LU provides the services and protocols necessary for the end user to communicate with other end users. This communication, or exchange of user data, is produced by establishing a session between the two LUs on behalf of the end users. Each session has a set of agreed-to protocols which are supported in the logical unit services at each session end. These protocols establish such session characteristics as data syntax and meanings, data flow methods, and other session properties that will be used during the session.

Establishment of a session also implies the use of certain physical network entities. Among these entities are nodes and links in the path between the two session end-points and the node resources at each session end, including buffers, storage, and processing capabilities. Figure 2 illustrates how the logical network resides on the physical network.

A session is initiated by an end user, which may be either an application program or a terminal user. The initiator of a session uses logical names, called network names, to identify the session end users. The SSCP has the responsibility to resolve the logical network names into the physical network addresses assigned to the session end users.

Logical network problems are typically attributable to software errors and can be generally classified as either detectable or undetectable.

End users, including the network operator, do not use network addresses to refer to sessions. Even the physical problem determination packages, such as NPDA, use network names in communication with the network operators. Correlation of the logical network names to physical addresses is accomplished using the services of the SSCP. For instance, physical failure notification, which is communicated internally in the network with network addresses, is translated into the logical network names by the SSCP prior to being sent to the network operators or communications network management applications[9] for processing.

The logical problem

Logical network problems are typically attributable to software errors and can be generally classified as either detectable or undetectable. A detectable error manifests itself in several ways: (a) an error message, which can be logged in sequential files or displayed to the network operator, (b) session failure notification to the end user, or (c) a storage dump caused by the abnormal termination of a software product. Any combination of these symptoms may be experienced as a result of a detectable logical error. Typically, these can be well-documented and resolved since pertinent information is normally available at the time of failure.

Figure 2 Logical network usage of physical network

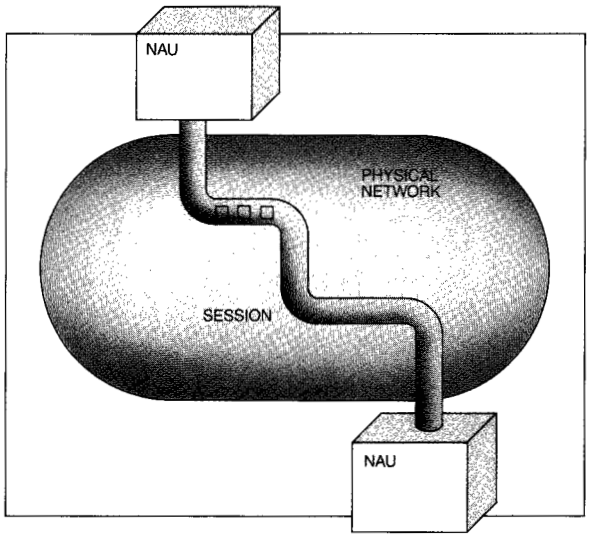

The undetectable errors appear to the session ends as if the network has "gone to sleep" and are therefore more difficult to resolve. These problems exhibit no external signals or notifications. The causes of this class of errors range from session protocol violations to network congestion. This paper examines two such problems, the protocol error and the loss of a message.

Protocol errors. As already mentioned, an SNA session follows specific agreed-to protocols. These protocols are established at session initiation. Two forms of undetectable software protocol errors can occur:

1. Mismatch or misunderstanding of protocols required for session initiation. This type of error can occur because of program error during session setup, or a mismatch can occur because of programming support levels of the software involved in SNA sessions.
2. Incorrect setting of protocol states by either the receiving or sending logical unit. Each protocol in SNA has a defined set of actions that it can

accept (or reject), depending on the current protocol state of the logical unit.

Loss of message. Another potential problem is the loss of a message or Path Information Unit (PIU) within the network without appropriate notification. The loss may result in an end user waiting for a message. It should be noted that NLDM addresses undetectable errors since SNA describes the protocols that are executed as a result of both detectable software and hardware errors or failures.

Past approaches to the problem. SNA provided facilities to support logical problem determination although these facilities were considered cumbersome and difficult. Methods that could be employed included the following:

1. Activating traces on a teleprocessing link, on a transmission group,[10] on an Advanced Communications Function/Network Control Program/Virtual Storage (ACF/NCP/VS) resource,[11] or within a host access method. This method requires a batch process to print the collected data. Since the printing is only available at host processors, there is usually a need to transmit large amounts of data through the network.

   Traces have several drawbacks associated with their usage. First, a trace will degrade network performance by using processing cycles of the node executing the trace, of nodes in the path used to transfer the traced data to a host processor, and of the host processing the received traced data for storage and printing. Naturally, the amount of degradation depends on the time the trace remains active and on the trace location.

   Second, the amount of trace data to be stored, printed, and reviewed may be extremely large. The data that are generated typically identify the network traffic via the SNA resource addresses. Correlating this physical information to the logical network names used by the SNA network operator is a manual exercise involving the analysis of the system definition that was active at the time of failure. This aspect coupled with the large amount of data may make finding a specific PIU extremely difficult.

   Third, in order to find a previously undetected error, the error must reoccur while the trace is active. This can be difficult to predict both in time (session activity) as well as trace location. The length of time required to recreate a specific problem has a direct impact on the amount of data generated.

2. Taking storage dumps of all suspected software components that are contained within the physical path used by a session. This method has several drawbacks. First, dumps that can deactivate the physical component in the network must be taken. Second, the failure data may have been overlaid prior to the dump occurring. Third, several storage dumps from several software components may have to be correlated to understand the cause of the error associated with a specific session. Fourth, with the introduction of the multiple route release of SNA in 1979 and the SNA Network Interconnection announced in 1983, it is difficult to determine which software components are within the path used by the session to be dumped.

NLDM collects, stores, and monitors network logical problem determination data.

3. Placing traps within suspected software components to cause a planned abnormal termination when the logical error reoccurs, then taking a storage dump and handling this error as a detectable failure. Several drawbacks are also associated with this method. First, the software component that must contain the trap must be deactivated and reactivated with the trap included. Second, the error must reoccur while the trap is active. Third, the execution of the trap will cause the component that could affect network operations to fail. Fourth, the trap may be placed in the wrong part of the software component, and if the condition persists, a new trap must be tried.

An IBM study[12] was conducted in 1979 to review 36 logical problems reported by SNA customers. The

study showed that, on the average, resolution of these undetectable errors required an extended period of time and required several attempts to recreate the problem and detect them. When the problems were resolved, the logical errors spanned software from the host access methods, to the communications controller, to the cluster controllers.

The study group concluded that IBM required a product for logical problem determination similar to the products used for physical problem determination, such as the NPDA. The group recommended that IBM provide the SNA user with an on-line interactive application program able to monitor software activity on a session basis.

The objectives. From this study, objectives were formulated for such a program product. They were

1. To provide a facility to assist the user with logical network problem determination in an SNA network.
2. To record relevant session data and activity in process at and just prior to the first occurrence of a failure.
3. To record the protocols associated with a session at session initiation time.
4. To record the relevant physical node data of the session end points to assist in problem isolation to a specific physical node.
5. To allow the above four facilities to be accessed on line from a single centralized or multiple distributed network operator terminal(s).

All of the above objectives were met with the introduction of the first release of NLDM.

The Network Logical Data Manager

An approach to resolving undetectable logical problems. NLDM collects, stores, and monitors network logical problem determination data. NLDM utilizes the services of the Network Communications Control Facility (NCCF)[13] to obtain and display the session-related data at a centralized or distributed NCCF network operator terminal.[8]

NLDM collects two types of session-related data: session awareness and session trace data. Session awareness is notification by the SNA access methods (Advanced Communications Function/Telecommunications Access Method, or ACF/TCAM,[14] and Advanced Communications Function/Virtual Telecommunications Access Method, or ACF/VTAM)[15] to NLDM that a session has been successfully started. The session awareness data consist of a session start and end indication, session partner network names and network addresses, session type, and configuration information about the session end points.

The collection and display of the network address is a new feature not previously available from the network. Prior to this, the network address of the LU was shielded from the network operator and the end user by the SSCP. Even the physical problem determination package, NPDA, used network names for its operator interface.

The session trace data are supplied to NLDM by the SNA access methods as well as the ACF/NCP/VS[16] program. The session trace data obtained from the host access methods include the following parts of a message or PIU: the transmission header, the request/response header, and the first 11 bytes of the request unit or user data. Session trace data are only collected for sessions involving a resource for which a trace has been started. NLDM collects these data in storage for active sessions and places the data on a Virtual Storage Access Method (VSAM) file at session termination for limited historical viewing purposes.

NLDM also collects session trace data from the boundary function serving a session end point in an ACF/NCP/VS boundary node. These data consist of the last four PIU sequence numbers, which are contained in the transmission header, and the appropriate control block information related to the ACF/NCP/VS resource involved in the session. These ACF/NCP/VS data are automatically sent to NLDM at session termination or can be solicited during the session, as long as the trace for the session is active.

NLDM uses NCCF services to establish access to the network to collect the session awareness data, the session trace data, and control information needed to start and/or stop the session trace. Two paths are used: the communications network management interface and an LU-to-LU session.

The communications network management interface was incorporated into the access methods in 1978 to allow an application program (i.e., NCCF) to interact with the SSCP. The interface allows communications network management requests to be issued to a network node to obtain statistical main-

Figure 3 NLDM's interfaces to access methods

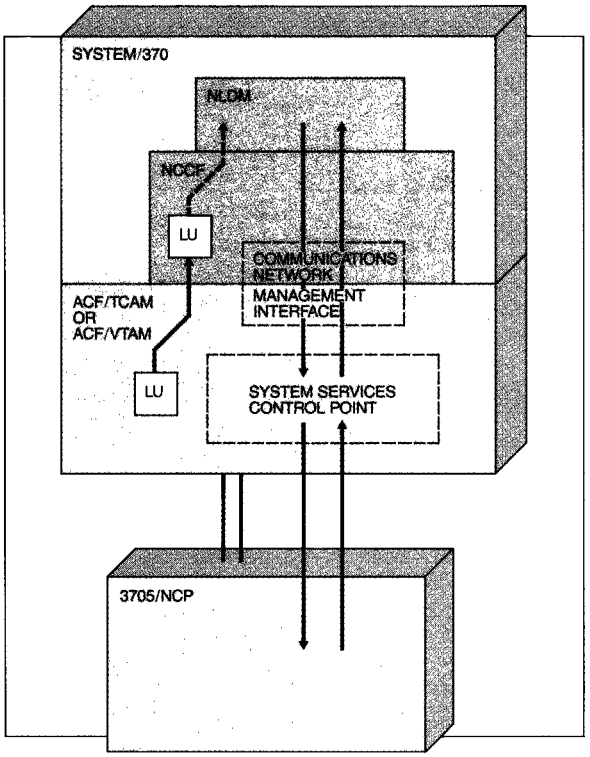

Figure 4 Network configuration

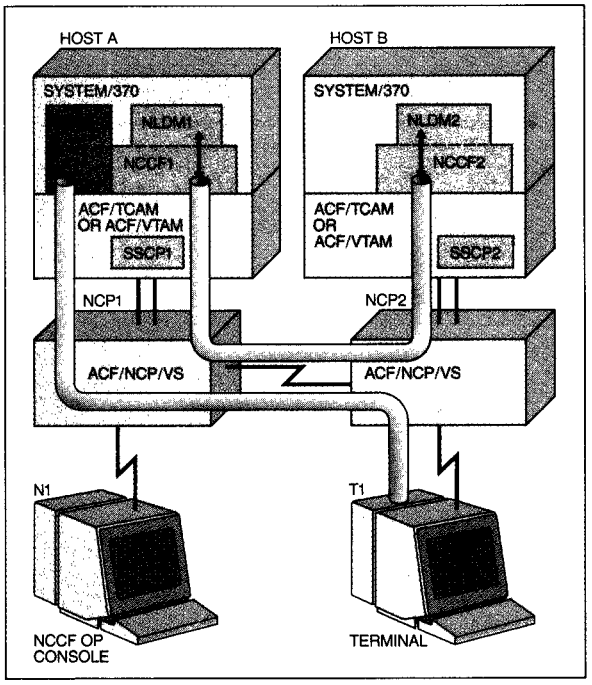

tenance data for use in problem determination. The first communications network management product to utilize this interface was NPDA. This interface, as illustrated in Figure 3, is also being used by NLDM for the following purposes:

1. To request that an SSCP send a session trace activation/deactivation request to network nodes under its control, including itself as a host node and any communications controller it owns.
2. To obtain the session awareness data generated by the SSCP participating in session initiation and termination.
3. To request that the ACF/NCP/VS provide it with the session trace data stored for a specific session.
4. To receive the session trace data from an ACF/NCP/VS at session termination for which the ACF/NCP/VS is the boundary node for the session end point.

NLDM does not use the communications network management interface to obtain the session trace data from the host access methods. It uses an LU-to-LU session also shown in Figure 3. The host access method captures and blocks the session trace data in a buffer whose size is determined by the user. Normally two buffers are used for this purpose, in an alternating fashion. As one becomes full, its contents are sent to NLDM, and the other buffer is placed in use. These buffers continue to alternate until all session trace activity has been halted in the access method.

NLDM uses the SSCP for the collection of session awareness data, for trace activation and termination, and for the collection of trace data obtained from the 3705/NCP. This fact makes it necessary to have an NLDM package at each host node containing the session end's controlling SSCP. As you may recall, the objectives specified that the session-related data were to be displayable at a single network operator terminal. As pointed out in the previous discussion, there is the possibility that two NLDM packages will be collecting data relating to the same session for session ends controlled by different SSCPs (called Cross Domain sessions). To see this concept illustrated, see Figure 4.

Collected session data are displayed by an NCCF network operator utilizing the supplied hierarchical display structure of NLDM. The network operator uses the network names to display session-related data. The operator can display such data for both

Figure 5 NLDM hierarchy for obtaining logical problem determination data

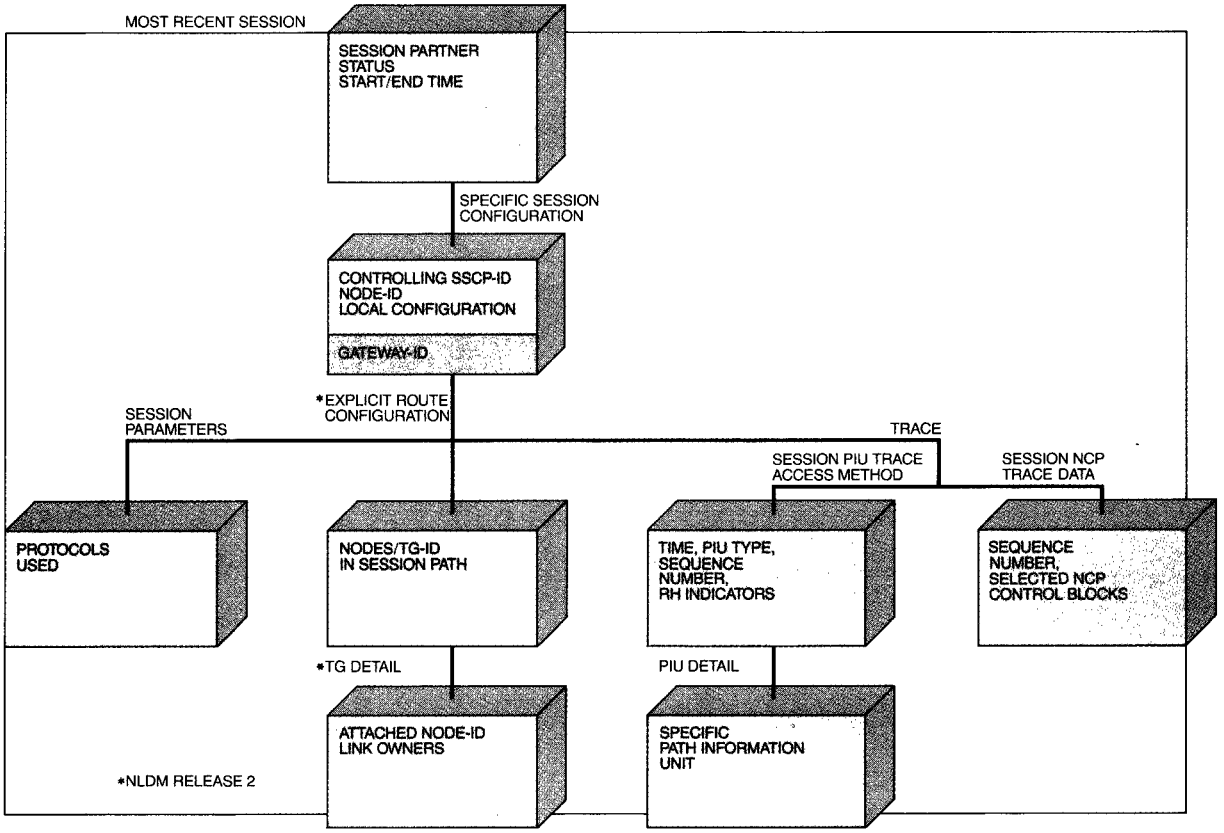

session ends at the same display station, although two NLDMs may be used to collect the session data for a cross-domain session.

To provide the network operator with a single network operational view in an SNA cross-domain environment, NLDM uses the services of NCCF to establish communications between itself and the other NLDMs. From the session awareness data obtained at session activation, an NLDM can determine which other NLDM to access to obtain data relating to that session.

Figure 4 illustrates this configuration. NLDM1 in HOST A is provided with the configuration information data by SSCP1 for an application whose network name is APPL because SSCP1 is the owner of APPL and participates in the session set up between APPL and T1. NLDM1 cannot directly obtain configuration information or trace data relating to terminal T1 since T1 is not under the control of SSCP1. NLDM1 must locate the appropriate resource owner and request the session awareness and trace data from the NLDM associated with that resource owner.

In this example, if the operator requires session trace data from T1, NLDM1 will request these data automatically from NLDM2. The data will be returned to NLDM1 for display to the requesting NCCF network operator. Note that if the two resources (APPL and T1) had been under the control of the same SSCP, only one NLDM would have been involved in the data collection and display.

NLDM provides the user with a hierarchy of displays that help in obtaining session-related data from both session ends. Figure 5 illustrates the NLDM hierarchy used to obtain logical problem determination data for a specific session.

The following example illustrates how the NLDM displays can be used in solving a logical or software-related problem. This problem appears to the terminal end user as if the "system has gone to sleep."

Figure 6 Session history for selected NAU

The example illustrates a bracket protocol error attributable to a software programming error. But before the NLDM methodology used to identify the undetected error is described, a short description of the bracket protocol is presented.

Bracket/data flow protocol. The bracket protocol[17] defines a specific set of rules for controlling a related "conversation" between two logical units connected in a session. The specific set of bracket rules to be followed for a specific session are agreed to at session initiation. Bracket/data flow rules identify who can start or end a new related conversation, when it can be started, and when each LU can send data within the related conversation. The use of the protocol ensures that data from each session end are not sent at the same time.

The bracket protocol defines several states that must be obeyed in order to have an orderly conversation. Bracket indicators placed in the Request/Response Header (RH) are used to control the bracket states. The intent of this section is not to describe the bracket protocol in total, but only that which is necessary to describe the logical error.

To start a related conversation, the authorized session end places a "Begin Bracket," or BB, indicator in the RH. When all the user data are sent and a reply is required, a data flow control indicator, "Change Direction," or CD, will be placed into the RH. The Change Direction indicator signals the receiving session end that it is its turn to send data. This flip-flop of the conversation controlled by the Change Direction indicator continues until the conversation ends with an "End Bracket," or EB, indicator.

The setting of the Change Direction indicator is extremely important. If the receiving LU is not given this indication, it cannot send data back to the requestor.

Let us take an example to indicate how an undetected error can be diagnosed as a bracket protocol error. Consider a new application program called APPL that uses bracket protocol between itself and various terminals. Let us assume that a specific session exists between APPL and T1 as in Figure 4. In this example, the application program starts a related conversation with T1 to ask a specific question and receive an answer. Then, based on the answer, the application starts one of several other related conversations to obtain more information.

Let us assume that the terminal operator has no indication that an error occurred. To the terminal operator, the logical error appears as a "long wait." This wait continues until the terminal operator becomes frustrated and requests assistance from the NCCF network operator. At that point, the network operator can use NLDM to display the session activity associated with the terminal T1 as illustrated in Figure 6. From this screen, the operator can see that an active session is still in progress between APPL and T1 and that further problem determination is needed to isolate the problem.

The NCCF operator has several options. One is to display the configuration data screen associated with terminal T1 and then utilize the physical problem determination techniques to see whether a physical node problem caused the error. Another is to continue to use NLDM in an attempt to find the source of the problem. To continue with the example, the operator can use NLDM to display the session activity associated with each session end. Let us assume that the NLDM operator displays the session-related data for APPL as in Figure 7. As seen there, the first four entries under the SEL (SELection) column are equal to the following SNA messages:

- 1 is APPL initiating a related conversation with a Begin Bracket (BB) indicator.
- 2 is APPL asking the T1 operator a question, then allowing T1 to respond by sending a Change Direction (CD) indicator.
- 3 is T1 answering the question, then allowing APPL to respond by sending a Change Direction (CD) indicator.
- 4 is APPL ending this related conversation with an End Bracket (EB) indicator.

Figure 7 Specific session trace data—APPL

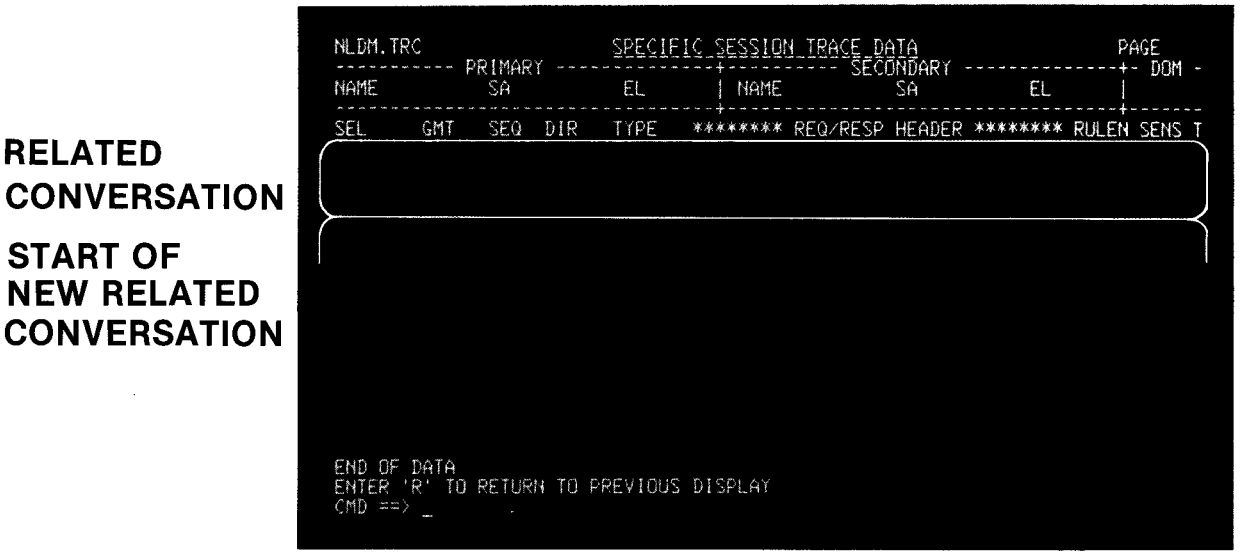

These four messages should appear normal to the NLDM operator since they adhere to the application scenario and the bracket protocol. As the scenario continues, the application now analyzes the received data and starts a different related conversation to gather other information. The application will use a bracket sequence similar to that used on SEL lines 1 through 4.

The data on the NLDM screen indicate that the next related conversation was started on SEL line 5 with a Begin Bracket (BB) indicator. This message was followed by another message sent by APPL on line 6, but note that no Change Direction (CD) indicator was placed on this message. This is an indication to a trained NCCF network operator or NLDM diagnostician that an error has occurred since the terminal cannot respond without receiving a Change Direction indicator.

Since this application program was new, the NLDM operator can assume that the application program failed to request that the access method place the Change Direction indicator in the message for this specific related conversation.

In this case, further diagnostic tests will not be necessary. The application program must be fixed to request that the Change Direction indicator be placed in the appropriate message field for this specific conversation.

Protocol error diagnosis requires that the NCCF operator have SNA knowledge or diagnostic procedures identified by the installation. If neither is available, the NLDM operator can request that NLDM record these data on its VSAM file for later use by a trained diagnostician.

A lost PIU is easier to detect since the NLDM operator need only detect a mismatch between the last four sequence numbers. What is not easy to determine is which software component in which network node "lost" the PIU. The only configuration data available with NLDM Release 1 are physical unit network names identifying the boundary nodes supporting the session end points. By knowing the complete physical route assigned to the session, the network operator could utilize other problem determination methods to attempt to isolate the problem further.

The knowledge to understand the physical route being used by the session has become more difficult with each release of SNA. The latest release, SNA Network Interconnection, has made this task exceedingly difficult without the use of NLDM. The support needed to identify the physical path

Figure 8 Session to virtual route to explicit route

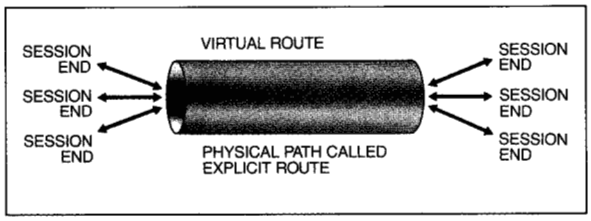

assigned to a session was implemented in NLDM Release 2 and was announced with the SNA network interconnection release.

There are several reasons why the path display data were placed into NLDM Release 2 at this time and not in the previous release. The following historical perspective of the growth of SNA session usage of the physical network explains what was behind these reasons.

Logical to physical mapping

From 1974 through 1978, sessions within an SNA network always traversed the same physical nodes and links if they originated and terminated in the same two nodes. These paths were predefined by the network administrator. If the end user could identify which node the terminal and application resided in, the network administrator could determine, via either a wall map or a user-provided application program, which physical nodes the session traversed.

With the introduction of multiple routes between host nodes and communications controllers in 1979, determining this route became more difficult. Multiple routes allowed up to eight distinct and unique routes to be created between any two network subarea nodes (i.e., host or communications controller nodes). Sessions between these nodes could be assigned to any of the eight routes, depending on the level of network service they needed. The sessions were not directly mapped to a physical path, called an explicit route,[18] but to a logical path called a virtual route. Many sessions could be mapped to the same virtual route, where many virtual routes could use the same explicit route.

For the network operator to understand the physical path assigned to a session in the multiple route environment, the operator would have to first know to which virtual route the session was assigned. Then the operator would have to map the virtual route to the explicit route to obtain the physical path data as shown in Figure 8. The virtual and explicit route data assigned to a specific session could be obtained on a session basis with a combination of session status[19] data requested from the access method via operator command and then either viewing a "wall map" or a user-supplied application program to determine which physical nodes were involved in an explicit route. Although possible, this method presented difficulties.

With the development of SNA network interconnection, it became apparent that a new method was required to obtain the physical configuration traversed by a session.

With the development of the SNA network interconnection, it became apparent that a new method was required to obtain the physical configuration traversed by a session. The fundamental structure of the SNA network interconnection dictates this conclusion. When a session spans multiple networks, each session segment is assigned to a different logical and physical path in each interconnecting network. The path identification (virtual route and explicit route) changes in each interconnected network. Total path data relating to all the routes spanning multiple networks could no longer be reasonably obtained for an internetwork session. Additionally, the old "wall-map" approach would not have been reasonable since each network could be independently owned, and therefore, the total information may not have been available.

A mechanism was created to allow the NLDM network operator to obtain path configuration data on the entire physical path between two session ends whether or not the session spanned an SNA interconnected network. The additional configuration data, which can be displayed, include the network names and/or subarea portion of the network address of

physical nodes traversed by the session, the transmission group numbers, the owners of any link within a transmission group, the names of any gateway nodes traversed by the internetwork session, and the virtual and explicit route number in use by the session. This support was placed into NLDM Release 2.[20]

Two additional capabilities were required by NLDM for the interconnected environment. The first deals with modifications of NLDM Release 1 to acquire session trace data from the real session ends contained in nodes in the separate networks. NLDM Release 1 has the capability of displaying session trace data about both session ends at the same NCCF operator console. It uses an NLDM-to-NLDM session to obtain the data if either of the session ends are in another SSCP control domain. To obtain data about both session ends, NLDM needed a unique identifier to obtain the data from another NLDM. The identifier selected was the SNA address pair.

This unique address pair is not available in an SNA interconnected network environment because the address pairs are translated as they traverse an ACF/NCP/VS gateway. Therefore, NLDM was modified to utilize the Procedure Correlator (PCID) as the identifier to retrieve the session trace data across an interconnected network.

The other capability that has been provided by NLDM Release 2 is an enhancement to a facility developed and released in 1979 and called Route Test.[18] Route Test is essentially an "echo test" which traverses a specific set of physical network nodes to test for physical connectivity. If the test fails, notification of the failing location is sent to the test requestor as well as the owner of the failing element. The network operator can request the test to be executed on either the logical or physical route between any two host and/or communications controller nodes.

The Route Test utilizes the virtual or explicit route notation. This mechanism cannot be used within an interconnected network to test the total end-to-end route utilized by a session. Just as the session address pairs are translated as the session traverses a gateway node, the routes utilized change. To allow the Route Test to continue in an interconnected network, NLDM has been upgraded to execute a series of Route Tests in all the networks traversed by a session. Prior to this enhancement, the Route Test could only be used to test connectivity for either a virtual or explicit route using the appropriate route notation. The network operator had to determine to which route a session was assigned. Determination was done by utilizing network operator commands appropriate to the access method to display routes utilized by a specific session. Then the operator would have to request a Route Test using the appropriate data. The new enhancement to NLDM makes this very easy. The network operator need only ask for a connectivity test to be executed for a specific named session.

NLDM Release 2

The first modification required to provide these functions was made to the session awareness data. New information was needed to identify the physical and logical route to which the session was assigned. These data were needed to initiate SNA commands to obtain the configuration data, which will be described later. A new session awareness has been created for internetwork sessions. For an internetwork session, the gateway NCP contains two session ends and the data needed to translate the network address session pair from one network to another. Those data include the network address pairs, identification of the next adjacent SSCP (and therefore NLDM) in the direction of the primary and secondary end points and route information in adjacent networks. The data are obtained via SNA flows from the gateway NCP to the gateway SSCP. These data are then forwarded to NLDM via a new parallel Logical Unit session used for session awareness data. The data are then used by NLDM to correlate requests from other NLDMs relating to a specific internetwork session.

Configuration data. As an internetwork session is created, each SSCP involved in the session setup provides its local NLDM with the session awareness data as shown in Figure 9. Each NLDM on the internetwork session setup path checks its individual configuration data base to ascertain whether the physical configuration explicit route data have already been collected. If configuration data are already available, no action takes place. If data are not available, NLDM issues a request via the communications network management interface to its SSCP to issue an SNA command to obtain the configuration data. SNA utilizes a modified Route Test for this purpose. The Route Test will traverse the physical network on the same route used by the session. As the test request returns to the requesting

Figure 9 Session awareness data

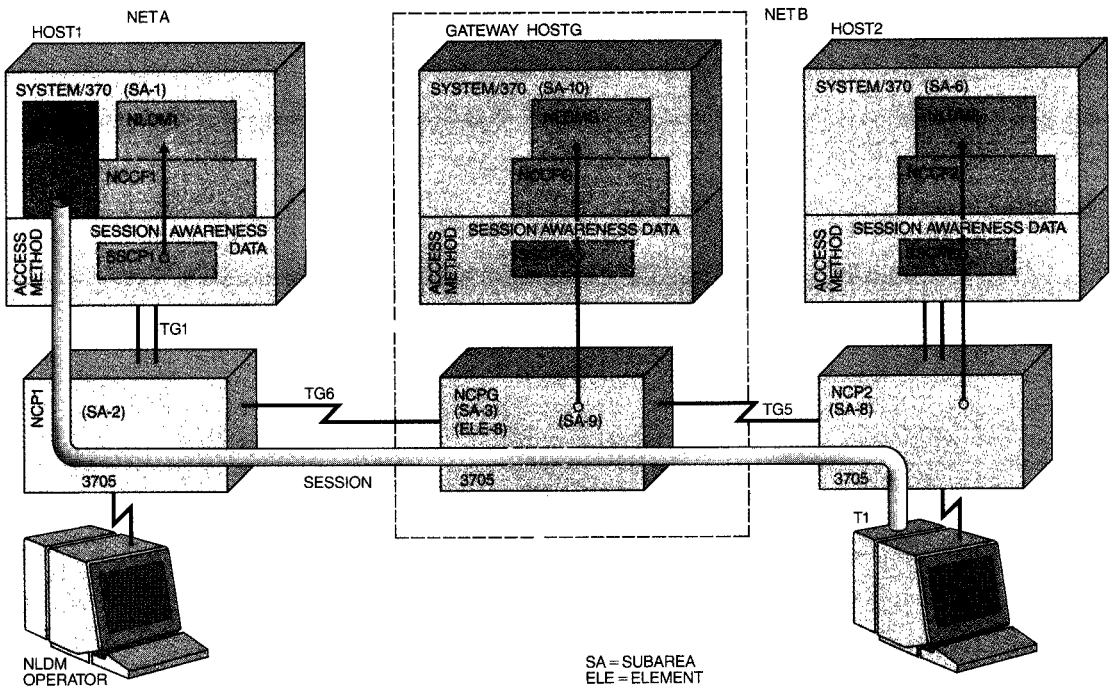

NLDM, the configuration data are appended by the nodes that comprise the route.

This mechanism is illustrated with Figure 9. After the session has been established between APPL and terminal T1, session awareness data are sent to SSCP1, SSCPG, and SSCP2. Let us consider what occurs in NLDM1. NLDM1 will use the session awareness data to determine what physical path has been assigned to the APPL session between the host and the gateway NCP (the other perceived session end within the network). If configuration data were not available, NLDM would send a request via SSCP1 into the network to the gateway NCPG to obtain the configuration data. The results would be returned via SSCP1. The NLDM at the gateway SSCPG host would also request the configuration data, but in this case, from both interconnected networks. This would occur since the session awareness it receives would contain the session data related to each session connected to the gateway NCPG. To obtain the configuration data, NLDM2 would use two separate configuration requests for the data. The requests would be issued to the gateway NCPG via the gateway SSCPG. One request would be for configuration data from the gateway NCPG to the APPL host, whereas the other would be from the gateway NCPG to node NCP2 supporting terminal node T1.

The configuration data is recorded at each NLDM that has recorded the session awareness data to be used for operator display.

NLDM can display the configuration data for a specific named session from end to end even if the session traverses a gateway. For a gateway environment, NLDM will only display this configuration data on a network-by-network basis. Each NLDM within the session initiation path of the session will display configuration data upon request. This aspect is illustrated by the network defined in Figure 9. Using the NLDM hierarchy, the network operator at the NLDM display console can display the most recent session screen that shows that APPL is in an active internetwork session with T1.

Using the appropriate display prompting methodology, the NLDM operator can display the specific configuration data for APPL as depicted in Figure 10. NLDM Release 2 has modified this NLDM

Release 1 display by including data that identify the logical and physical route to which the session has been assigned within a network. As can be seen, the session is assigned to the logical (virtual) route called VR2 which uses the physical (explicit) route called ER4. From Figure 10, one can also see that the internetwork session of APPL and T1 originates in the host node identified as HOST1 or subarea 1 and traverses the gateway node called NCPG or subarea 3. The display ends at the gateway node although the session traverses the gateway. An authorized operator can request the display of the next network's configuration data for this specific session. Note that this display does not identify whether there were any other subarea nodes within the session path between subarea nodes 1 and 3.

To identify the complete path used by the session in this network, the NLDM operator could request the next display in the hierarchy. This display, illustrated in Figure 11, identifies each subarea node and transmission group that the session traverses. This display provides the mapping of the logical network to the physical network on a session basis. Again, this display identifies the physical nodes using their user-defined network names traversed within one network for this internetwork session.

The NLDM operator can view the other network's data upon appropriate authorization via NLDM as illustrated in Figure 12. Here we see that the session continues from NCPG to NCP2 via virtual route 6 (VR6) and explicit route 5 (ER5). Note that the logical network names of the SNA subarea nodes are contained on the NLDM network operator display screens. This information is provided since the network operators deal in the logical user-supplied network names for node identification purposes rather than in network addresses. These network names enable the network operators to utilize other communications network management packages, such as the NPDA, to obtain additional data to further isolate a failure to a specific named resource. For instance, after the NCCF network operator views the NLDM display screen shown in Figure 11, the same operator can request display data from NPDA on the communications controller called NCPG from the same console. If only the subarea number 3 were ready, the network operator would first have to "look up" the physical-to-logical name correlation prior to using NPDA or other communications network management tools to obtain data about that communications controller.

Figure 10 Specific session configuration data

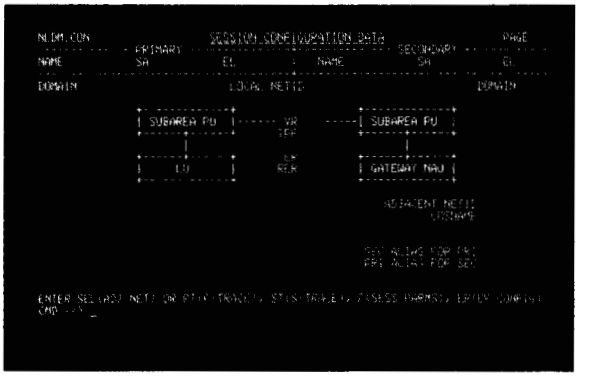

Figure 11 Specific ER configuration

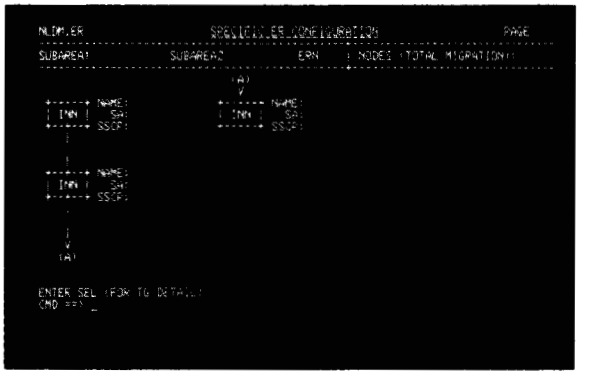

Figure 12 Configuration data of other network

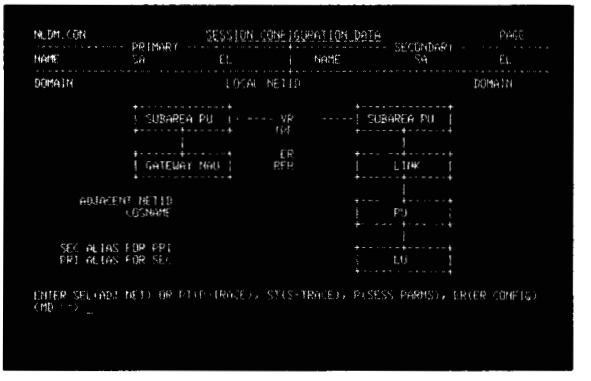

An automatic translation function was needed in NLDM Release 2 to obtain the network names of the nodes since the SNA command, Route Test Reply,

used to collect the configuration data operates on a physical network address basis and not on the logical network name basis. It was further felt that the user should not have to do the "look-up" task to correlate the network addresses to network names.

Prior to describing how NLDM dynamically obtains the translation of the physical network address to its logical network name, a brief description of the mechanism used by NLDM to obtain its local configuration data is needed. Each NLDM dynamically obtains its local configuration data from the session awareness data provided by the SSCP to which it is associated as sessions are initiated. These data include network addresses and network names of resources involved within the sessions. The data allow NLDM to dynamically obtain the physical-to-logical name mapping for all subareas under the control of the SSCP associated with that NLDM. Since these data are available, the only question that remains is how a remote NLDM accesses the data to obtain the physical-to-logical name translation.

At initialization, each NLDM attempts to initiate a cross-domain session with other NLDMs specified by the user via the Cross Domain Resource Management table of NCCF. Note that this table is used because sessions between NLDMs are actually between the logical unit services of NCCFs. NLDM, as well as other communications network management applications such as NPDA, uses NCCF logical unit services to initiate sessions for its use. After an NLDM-to-NLDM session is initiated, the NLDMs exchange configuration data about themselves, including the network identifier in which they reside, the logical and physical SSCP network address and name to which they are connected, the PU network name associated with the SSCP, and the NCCF network name associated with that NLDM. These data are then used in combination with the configuration data obtained from the Route Test reply to identify to which remote NLDM the request for the physical-to-logical name translation should be sent.

Recall that the configuration data obtained on the Route Test reply for each subarea node contain the network address of that node, the transmission group data, and the owning SSCP network address of any links that reside in the transmission groups on that path. To obtain the network-address-to-name translation for a network node identified on the Route Test reply, NLDM Release 2 will do the following processing. The owning SSCP network address associated with the specific resource in question is obtained from the Route Test reply. NLDM uses this owning SSCP network address to find with which remote NLDM this SSCP is associated as obtained from the configuration data exchange at NLDM-to-NLDM session initialization.

With this information, NLDM can issue a request for the network-address-to-name translation of the remote NLDM and wait for a response. If the remote NLDM does not have this information, the requesting NLDM can use the next owning SSCP for the same resource, if available, and repeat this process with another remote NLDM. The received network name can then be placed into the local NLDM configuration data base so that network operator displays will contain the logical network name as well as the physical address. This process is initiated for all subarea node addresses found on the Route Test reply which the receiving NLDM does not have in its local configuration data base. Also, this facility will operate for internetwork sessions as well as with the network identifier appended to the translation request by the sending NLDM.

Connectivity test. NLDM Release 2 also provides the user with the support needed to test the connectivity of a session path from one session end to the other, even for internetwork sessions.

The method employed also involves using an enhancement to the Route Test facility. When a session connectivity test is requested, each NLDM identified as containing the session awareness data for a specific session is requested, in order, to issue a Route Test for the physical route (explicit route) being used by the session within that network. If all the route tests are successful, the requestor is notified. If it fails within any of the interconnected networks, the network operator is notified that a failure has been detected in the physical route and is told which network has detected the error. Notification is also generated to the owner of the node detecting the failure so that corrective action can be taken.

The method is illustrated using the network depicted in Figure 13. The network operator at the NLDM display issues a connectivity test for the session between logical network elements called APPL and T1. This test will cause NLDM1 to request a connectivity test between the subarea node containing APPL (HOST1) and the representation of

Figure 13 Cross-network connectivity test

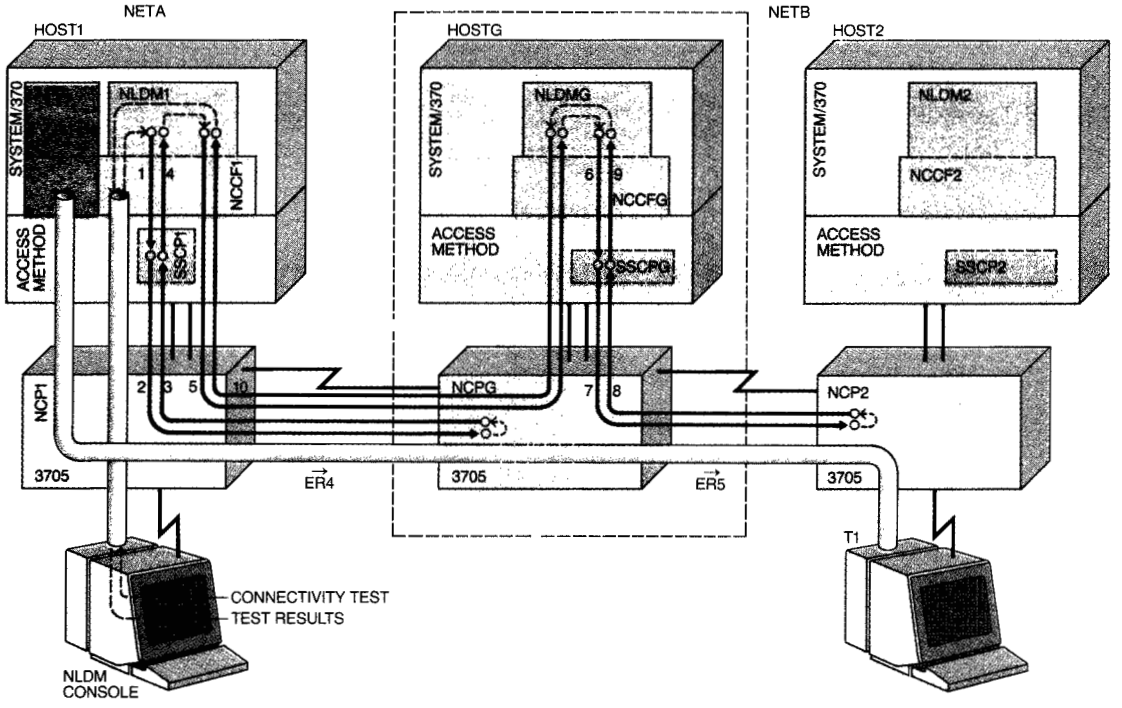

the other session end (T1) in the gateway node (NCPG) on the physical route to which that session is assigned (ER4). The request for the connectivity (Number 1 in Figure 13) will be sent by NLDM1 to its SSCP[21] for processing via the communications network management interface. After the connectivity test traverses network NETA (2) and returns (3), assuming that no failing nodes or links exist within that path, NLDM1 will be notified (4). Since this connectivity test is for an internetwork session, NLDM1 must first ascertain the connectivity of the remainder of the session across the gateway prior to notifying the requesting network operator of its results.

To ascertain the connectivity across the gateway node, NLDM1 will request (5) that NLDMG test the next leg of the session path. From the session awareness data acquired at session initiation, NLDM1 obtains an understanding of which NLDM to ask for the next part of the connectivity test. NLDMG will request that the connectivity test be issued from the gateway NCP (NCPG) to the Boundary Node (NCP2) supporting the session end point (6) via SSCPG.[21] This request will be issued via the communications network management interface to the SSCP that will use the appropriate SNA session to request that NCPG issue the SNA Route Test (7). The request for the connectivity test will contain the physical route number (ER5) and subarea addresses NCPG and NCP2, respectively, as obtained by NLDMG from the session awareness data that it received when the internetwork session was created. After the results from the Route Test from NCP2 have been returned (8,9) to NLDMG from NCPG, NLDMG will respond to NLDM1 (10). The network operator will then be informed as to the success or failure of the requested connectivity test.

The connectivity test can be initiated from any NLDM that has received the session awareness data at session activation. Referring to the network depicted in Figure 13, a network operator attached to either NLDM1, NLDMG, or NLDM2 could initiate the connectivity test. For instance, if a network operator were attached to NLDMG, then NLDMG would issue a Route Test from the gateway node NCPG to each boundary node supporting the session ends (HOST1 in NETA and NCP2 in NETB). Then it would wait for both responses prior to issuing the

Figure 14 NLDM Release 2 network interfaces

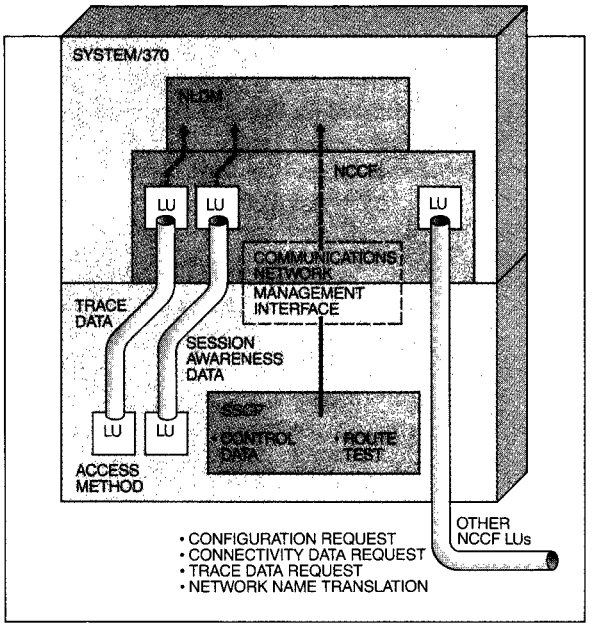

appropriate notification to the network operator. Neither NLDM1 nor NLDM2 would be involved with this test. This same option, using any NLDM receiving the session awareness data, could be used for initiation of the configuration data display.

To allow for the additional control commands from NLDM Release 2 to the access methods and for NLDM-to-NLDM communication, additional changes to the operational characteristics of NLDM Release 1 were made. The NLDM Release 2 network interfaces are illustrated in Figure 14. NLDM Release 2 utilizes an enhanced communications network management interface for control of the starting and stopping of the session awareness and trace data. It has been expanded to allow the route test request/response to be communicated to/from the SSCP. NLDM Release 2 uses two parallel LU-to-LU sessions between itself (NCCF provided) and the access methods. One is for session trace data; the other is to block session awareness data, which prior to this release were communicated to NLDM via the communications network management interface. NLDM also utilizes the NCCF-to-NCCF application-to-application session for the NLDM-to-NLDM initialization configuration exchange, requests for configuration data and connectivity tests for internetwork sessions, session trace data displays, and requests for network-address-to-name translations.

The other NLDM Release 1 function that needed upgrading was the capability to obtain session trace data for internetwork sessions. As previously mentioned, this function requires upgrading since the correlation used by NLDM Release 1 for session trace data was no longer unique for internetwork sessions. The correlation was the session address pair that changes each time the session traverses a gateway node. To allow the network operator to continue to be able to request correlated trace data as previously described in this paper, the NLDM session correlation vehicle was changed from a network address pair to the SNA Procedure Correlator (PCID). The NLDM associated with gateway nodes then swaps the PCID values to match the session identification "as it is known" in the next NLDM associated with the session. This request for rerouting is very similar to the rerouting function performed by gateway SSCPs. The session PCIDs are provided to the NLDM associated with the gateway node via the session awareness information processed at the internetwork session initiation. This methodology allows the network operator to continue to obtain session trace data even for internetwork sessions and still display these data at a single network operator console.

Summary

The two releases of NLDM have enhanced network logical problem determination in two areas. The first release simplifies procedures used to obtain session-related data for logical problem determination purposes. This simplification reduces time spent in problem determination activities. It also makes it less likely that more drastic detection methods, such as link traces and program traps, must be used. The second release of NLDM provides a mechanism that correlates the logical network to the physical network so that physical problem determination methods can be employed to further isolate network failures. Logical errors, especially the undetectable types, have always plagued the network users. Problem determination for these errors has been simplified but not eliminated. These two releases of NLDM have been established as a solid base to continue that effort.

Acknowledgments

The authors wish to acknowledge their colleagues for contributing to the conception and initial

design of the NLDM products: John O'Brien, Barbara Smith, and Tim Sullivan from the Communications Network Management group in Raleigh; Bruce Wilder from the Communications System Staff in Kingston; Jim Gilman and Jim Paterson from the VTAM design group in Kingston; and Glynn Furr and Roy Howard from the NCP design group in Raleigh. Credit is also due to the other designers, SNA architects, product planners, implementers, and testers without whose work these products would not have been completed. Special thanks go to Gary Schultz and Mac McGee from the Communications System Architecture group in Raleigh for their continuous and timely aid in the review of this paper.

The first part of this paper, dealing with the technology and example of NLDM Release 1, is a subset of one presented at the 1983 National Computer Conference in Anaheim, California.[22] Additional data have been added to clarify the rationale for this product and to give more information on the technology.

Cited references and notes

1. R. J. Sundstrom and G. D. Schultz, "SNA's first six years: 1974–1980," *ICCC 80 Proceedings* (1980), pp. 578–585.
2. *NPDA Version 3 General Information Manual*, GC34-2110, IBM Corporation; available through IBM branch offices.
3. S. Huon and R. Smith, "Network problem-determination aids in microprocessor-based modems," *IBM Journal of Research and Development* **25,** No. 1, 3–16 (January 1981).
4. *NLDM General Information Manual*, GC30-3081, IBM Corporation; available through IBM branch offices.
5. *Network Program Products General Information Manual*, GC27-0657, IBM Corporation; available through IBM branch offices.
6. J. H. Benjamin, M. L. Hess, R. A. Weingarten, and W. R. Wheeler, "Interconnecting SNA networks," *IBM Systems Journal* **22,** No. 4, 344–366 (1983, this issue).
7. J. B. Ford, "Enhanced problem determination capability for teleprocessing," *IBM Systems Journal* **17,** No. 3, 276–287 (1978).
8. R. A. Weingarten, "An integrated approach to centralized communications network management," *IBM Systems Journal* **18,** No. 4, 484–506 (1979).
9. T. P. Sullivan, "Communications Network Management enhancements for SNA networks: An overview," *IBM Systems Journal* **22,** Nos. 1/2, 129–142 (1983).
10. A transmission group consists of a set of parallel links defined between two subarea nodes that appear as a single logical link.
11. ACF/NCP/VS Version 2, Release 1 supports a new function called the Generalized PIU Trace. This trace allows the network operator to request a continuous trace on a specific ACF/NCP/VS resource.
12. The IBM Study group was chaired by M. B. McGahee of the Information Systems Group.
13. *NCCF General Information Manual*, GC27-0429, IBM Corporation; available through IBM branch offices.
14. *ACF/TCAM, Version 2 General Information: Introduction*, GC30-3057, IBM Corporation; available through IBM branch offices.
15. *ACF/VTAM Version 2 General Information*, GC27-0608, IBM Corporation; available through IBM branch offices.
16. *ACF/NCP/VS Version 2 General Information Manual*, GC30-3058, IBM Corporation; available through IBM branch offices.
17. *Systems Network Architecture Format and Protocol References Manual: Architecture Logic*, SC30-3112, IBM Corporation; available through IBM branch offices.
18. J. P. Gray and T. B. McNeill, "SNA multiple-system networking," *IBM Systems Journal* **18,** No. 2, 263–297 (1979).
19. J. G. Waclawsky, *ACF Release 3 Problem Determination Guide*, Technical Bulletin, G229-0381, IBM Corporation; available through IBM branch offices.
20. Support of NLDM Release 2 full function is contained in ACF/NCP Version 3, NCCF Version 2, and ACF/VTAM Version 2, Release 2. NLDM Release 2 provides limited support for various other levels of ACF/NCP Version 2, ACF/TCAM Version 2, and ACF/VTAM Version 2.
21. The SSCP request that the Physical Unit Service of the subarea node initiates is the actual SNA Route Test commands.
22. R. A. Weingarten and E. E. Iacobucci, "Logical problem determination in an SNA network," *National Computer Conference Proceedings*, Anaheim, CA (May 1983).

Reprint Order No. G321-5201.

Robert A. Weingarten *IBM Corporate Headquarters, Old Orchard Road, Armonk, New York 10504.* Mr. Weingarten joined IBM in 1969 in the former IBM New York Development Center, where his assignment was on the OS/360 linkage editor. He joined the U.S. Army in 1970. Upon returning to IBM in 1972, he worked on the DOS RPG II compiler. In 1974, he transferred to Kingston, New York, where he was involved in various aspects of the definition of Systems Network Architecture, including high-level systems design, systems requirements gathering and planning, and system design management for the Advanced Communication Function access methods and communications network management. In 1982, he was the control program design and development manager in the scientific and engineering processor development area. Since March 1983, he has been assigned as a consultant on the Engineering, Programming, and Technology corporate staff, concentrating on communication programs. Mr. Weingarten received his B.S. and M.S. degrees in electrical engineering from New York University in 1967 and 1969, respectively.

Edward E. Iacobucci *IBM Communication Products Division, P.O. Box 12195, Research Triangle Park, North Carolina 27709.* Mr. Iacobucci is an advisory programmer with Network Management products, currently engaged in advanced network management application design. His present interests and responsibilities center around the definition of software tools which may be employed to assist in the management of large SNA networks. Until 1982, he worked in Communications Network Management design, where he led design teams for two Network Management products. Prior to joining IBM, Mr. Iacobucci had worked for Black and Decker, where he specialized in discrete time simulation models of distribution networks. Mr. Iacobucci received his B.S. in systems engineering from the Georgia Institute of Technology in 1975.