IBM

Washington Systems Center

Technical Bulletin

Network Job Entry Formats and Protocols for System/370 Program Products

D. B. Hiller
J. M. Hutchinson
J. E. Richards, Jr.

GG22-9373-02
June 1986

Washington Systems Center
Gaithersburg, MD
Technical Bulletin

Network Job Entry Formats and Protocols for System/370 Program Products

Dean B. Hiller
Research Division

-

James E. Richards, Jr.
NJE Protocols Review Board

-

John M. Hutchinson
Washington Systems Center

GG22-9373-02
June 1986

Third Edition, June 1986

This is a major revision of, and obsoletes, GG22-9373-01. See "Summary of Changes" on page xi for details.

Preface

Abstract

This document explains the Network Job Entry (NJE) formats and protocols used by the following System/370 Program Products which support NJE. Product-specific comments refer to the latest shipped release unless otherwise stated (as of the date of this manual).

- MVS/SP JES2 (Version 1.3.6 / 2.1.5)
- MVS/SP JES3 (Version 2.1.5)
- RSCS Networking (Version 2.1)
- VSE/POWER (Version 2.2)

Network job entry provides access to batch computing facilities from other host systems. It enables users to transfer work and data throughout a distributed network of those systems.

Terminology

"Network job entry" (or "NJE") can be defined as: "A facility for transmitting jobs (JCL and in-stream data sets), sysout data sets, (job-oriented) operator commands and operator messages, and job accounting information from one computing system to another."

"Network job entry", or "NJE", is not a part of "Systems Network Architecture (SNA)", but is an application layer which uses SNA, BSC and CTC transmission facilities.

"Network Job Entry" was originally JES2's name for its implementation of NJE. "Network Job Interface", or "NJI", was the name used by HASP/ASP/JES3/RSCS for those products that interfaced to JES2/NJE.

For other definitions relating to NJE, see "Glossary and Abbreviations" at the end of this publication.

Purpose and Use

This document is intended to consolidate internal data stream, control format, and communications information used in network job entry. Its purpose is to allow users and systems programmers to make full use of job processing facilities available with these products. This document should also be useful in problem determination activities, and in customizing the products for specific environments or applications.

This document is not intended to replace existing product manuals, which are listed in Appendix C, "NJE Bibliography." For most product-related material, the product manuals should be consulted; however, Appendix B, "System Dependent Considerations and Comments" contains detailed descriptions of product deviations from the NJE formats and protocols. That section also contains other relevant product-specific comments. It is recommended that this section be read carefully.

The value of this document stems from its illumination of both the commonalities and differences in current NJE implementations. It is expected that, over time, implementations of NJE will converge to the common standard documented herein.

As an aid to the reader, the index has been augmented by the use of automatic indexing macros. They do not always produce what would be produced manually, but the results have proved to be valuable.

Revision Codes

Changes from prior versions of this document are flagged in the left margin with a revision code. Editorial changes are in general not flagged, however major stylistic changes to the document shall be noted. | Revisions since GG22-9373-01 dated June 1985 are flagged with the symbol "|." See "Summary of Changes" on page xi for more details.

Acknowledgements

This document was written and is maintained by the NJE Protocol Review Board which exists to guide IBM's actions in the area of NJE.

Its first responsibility is to provide IBM with a uniform and consistent set of protocols which can satisfy not only today's needs but can be extended in a consistent fashion to accommodate future requirements.

This document is the product of many people's effort over the last five years. The primary contributors are listed below:

Jim Ashfield, CPD
Vasile Bizovi, DSD
Janis Coltin, DSD
Francisco Chan, IS&TG - Past NJEPRB Chairman
Ron Elliott, DSD
Jon Franks, DSD
James Havender, IS&TG - Past NJEPRB Chairman
Dean Hiller, RES
John Hutchinson, NCMD
Keith Jones, STD
Austin Kilburn, ISG
Mike Laskey, DSD
John Marsh, ISG
Jesse Murphey, DSD
James Richards, Jr, DSD - Current NJEPRB Chairman
Walter Schueppen, DSD
David Stamper, DSD
Cynthia Westerman, STD
Richard Wood, CHQ
Fred Voss, DSD

Contents

Figures

Summary of Changes

Changes for GG22-9373-02 as Updated June, 1986

More minor changes were made to this version of the document. The following items are highlighted for readers of the previous version. (Numbers in parentheses refer to section heading numbers.)

New Features

- V.27 Modem Contention Resolution in RSCS V 2.1 (A.1.6)
- Prepare Mode Option in RSCS V 2.1 (A.3)

New Fields Added

- Data Set Header - OPTB Keys (3.5 and 8.5.6)

Product Release References

- MVS/SP-JES2 now references Versions 1.3.6 and 2.1.5
- MVS/SP-JES3 now references Version 2.1.5
- RSCS - added references to RSCS Networking Version 2.1

Changes for GG22-9373-01 as Updated June, 1985

Corrections and Clarifications

- "NJE Primary" definition corrected for Signon Processing (2.1)
- Signoff record sent by JES2 on CTC (8.8.6)
- JES2 CTC initialization (B.1.8.1)
- Flow figures cleaned up for readability

New Features

- Signon Concurrence feature added in POWER V 2.2 (2.1.1, 8.8.1, 8.8.2 and B)

New Fields Added

- Job Header - JES2 section (3.4 and 8.4.2)
- Job Header - POWER section (3.4 and 8.4.3)
- Data Set Header - POWER section (3.5 and 8.5.5)
- SRCB for CPDS added (5.3)

Product Release References

- RSCS - removed references to RSCS Networking Release 2
- POWER - added references to POWER Networking Version 2.2

Chapter 1. Introduction to the NJE Document

1.1 Audience

This document is intended to be used by anyone concerned with network job entry (NJE) protocols. It may be used by IBM marketing support personnel and IBM users to obtain an overview of how NJE protocols work in section 1.4, "Node Types and Functions" on page 1-2. It may also be used by service personnel or systems programmers for problem determination with NJE systems using Chapter 8, "Control Formats."

For an introduction to these protocols, please see *NJE Concepts and Protocols Overview*, GG66-0224.

1.2 Organization

The document is organized into parts by functional layers. Figure 1-1 shows the relationship of the different sections in the document.

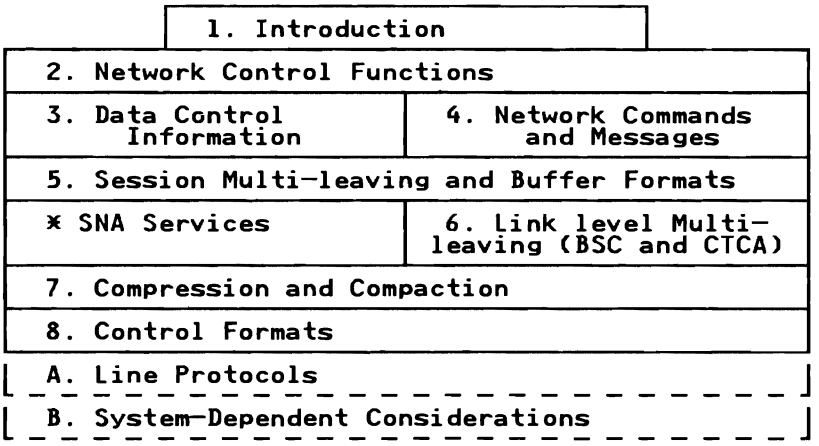

* See "A.4, Systems Network Architecture - LU Type 0" on page A-25 for this information.

Figure 1-1. Document Organization

The control records sent in NJE can be broken down into three types. The first type relates to network management and includes signon, signoff, and network path manager functions. Network management control records are described in Chapter 2, "NJE Control Functions." The second type of control record (Job Header, Dataset Header, and Job Trailer) is used to describe the data being sent through the network. Header records are in Chapter 3, "Data Control Information." The third type of control record is used for sending commands and messages to other network nodes. It is described in Chapter 4, "Network Command and Message Formats."

Chapter 5, "Multi-Leaving (ML) Functions - SNA, BSC, CTCA" describes the next lower layer, i.e. how data and control records are sent from system to system. The data and control traffic is the same regardless of the type of communication link. Chapter 6, "Additional Multi-Leaving (ML) Functions - BSC, CTCA" describes some multi-leaving functions unique to BSC and CTCA. Chapter 7, "Data Compression and Compaction" describes the data record format for BSC and SNA. Chapter 8, "Control Formats" contains a detailed description of every field in each network control block. If the field is not in this section, then it does not exist in the protocol. This section defines limits, default values, size, and type of every field.

NJE supports multiple types of links:

- the half-duplex bisynchronous communications line - BSC
- the channel to channel adapter - CTCA
- the systems network architecture logical unit type 0 - SNA LU0

Any given implementation may support all or a subset of the above types of links. Appendix A, "Line Protocols" describes the basic differences among the communications links.

The final section of the document, Appendix B, "System Dependent Considerations and Comments," contains product specific comments. It is not intended to replace the program logic manual for the product. It is included to allow a single reference for the networking components of the different operating systems.

1.3 Purpose

This document is intended to consolidate internal data stream, control format, and communications information used in network job entry. Its purpose is to allow users and systems programmers to make full use of job processing facilities available with these products. This document should also be useful in problem determination activities, and in customizing the products for specific environments (or applications).

1.4 Node Types and Functions

NJE nodes can support one or more of the three functions: transmit, store and forward, and receive.

The transmit function consists of packaging SYSIN or SYSOUT jobs in NJE Control Records and inserting them into the network.

The receive function recognizes jobs packaged in NJE Control Records and processes them. For example, SYSIN jobs are executed at that node or SYSOUT jobs might be printed.

A node having the store and forward (S&F) function, accepts NJE jobs, stores them and schedules them for forwarding to another node. It is permissible for a store and forward system to add additional sections to the Control Records as they are forwarded. It is not permissible for data or Control Record Sections to be changed, data or Control Record sections to be removed, or Control Record sections to be expanded except as a result of human intervention such as operator command, and so long as the resulting data stream is still a valid NJE data stream.

A second exception is granted for the purpose of updating accounting information in the Job Trailer. See "Print File Transparency" on page B-27 for RSCS specific information and "Print File Transparency" on page B-43 for POWER specific information.

Generally, all nodes must store and forward any data that is imbedded in NJE. When the data reaches its destination, it may or may not be processed as the user intended, depending on the facilities available at that node. NJE protocols allow the destination node to reject files that it cannot process or perform other system-dependent actions.

In NJE terminology, a job refers to both SYSIN and SYSOUT data. Throughout this document the term job will be used if no distinction is required. If a distinction is necessary the term SYSIN or SYSOUT will be used.

There are three data types. The first and most common type is JOB. The JOB type can further be broken down into SYSIN and SYSOUT. SYSIN is that data being transmitted to a system for execution at the system. SYSOUT is generally output from a program intended to be printed, punched, or viewed at a terminal.

The second type of data is called nodal message records (NMR). Nodal message records are actually commands and messages that are sent from one node to another node. In the case of commands, the text is intended to be executed on the receiving node. Messages are intended to be displayed on the receiving system. In general, all systems with command capability must have message capability in order to receive the output from the command. NMRs are not stored and forwarded in the spool system. If a path does not exist to the next node, the NMR is discarded. If a command NMR is discarded, then a new NMR is routed back to the origin to inform the originator of the incomplete path.

The third type of data is network path manager (NPM) records. The network path manager is presently only implemented in JES2 and allows each node in a JES2 complex to keep track dynamically of the path between itself and every other node in the network. In a complex network (one with multiple paths between two or more nodes) the network path manager is able to recover dynamically from a line or node failure and reroute traffic via an alternate path.

Chapter 2. NJE Control Functions

Control Functions in NJE include Signon, Signoff and other path management functions. Signon protocols vary depending on whether or not there is a network path manager. (Currently, only JES2 has a network path manager.)

2.1 Signon Without Path Manager

In NJE on BSC (and CTC) connections, the node which first sends the SOH ENQ is also the node which sends the type 'I' signon record. (See Figure 2-1.) In NJE on SNA sessions, the node with the higher node name (according to the standard EBCDIC collating sequence) sends the type 'I' record. (See Figure 2-2 on page 2-2.) The node which sends the type 'I' record, will be referred to as the NJE "primary" in this discussion. In a system that does not support the Network Path Manager, signon is accomplished by the "primary" sending a type 'I' signon record and the "secondary" responding with a type 'J' response signon record.

This usage of "primary" and "secondary" is to be distinguished from the same terms used in the section A.1.1, "Initialization" on page A-1 where they refer to the relationship of the nodes at the BSC and CTCA link connection level. For SNA, the node relationship at the link connection level is described in A.4.2, "Session Initialization" on page A-25. Here, the terms refer to the relationship of the two nodes as a logical NJE transmitter/receiver pair.

| It is permissible to send a null buffer or a DLE ACK0 before the signon is sent. However, no other records of any type may be sent before the 'I' and 'J' signon records are exchanged. If a system receives a non-null buffer containing anything except a signon record at this time, it should terminate the line and restart the signon process.

'I' and 'J' records are sent uncompressed and uncompacted as indicated by the preceding SCB.

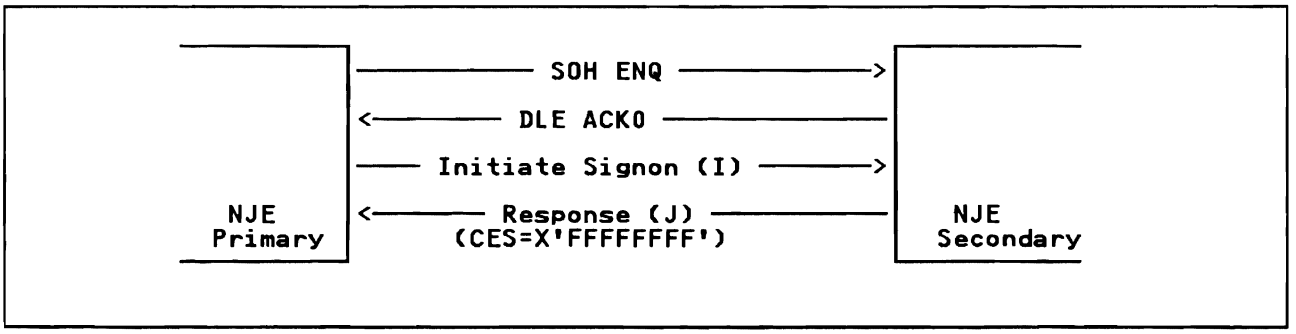

Figure 2-1. BSC/CTC Signon Sequence

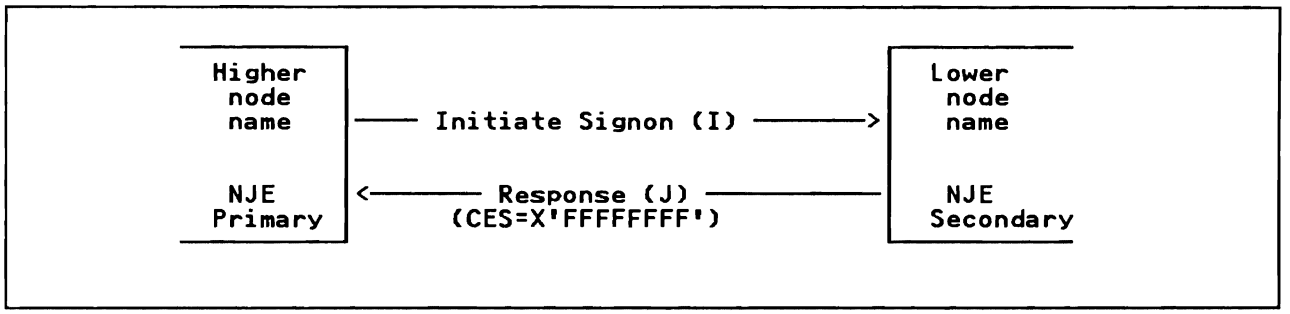

Figure 2-2. SNA Signon Sequence

The format of the type 'I' and 'J' records are described in section 8.8, "Connection and Control Records" on page 8-65. The type 'J' record contains a Connection Event Sequence (CES) of minus one (X'FFFFFFFF'). See B.3.2.1, "Signon" on page B-24 for RSCS specific processing.

2.1.1 Signon Concurrence Feature

This is an optional element of the signon protocol. It allows two systems to determine that each are able to work in an extended mode on the communications line running between them.

There are no operational changes. If anything, migrating to new releases with new features is transparent from an operational point of view until two systems have made it to a given new release.

2.1.1.1 External Interface

A method must be provided to allow the system programmer to define the extended features to be used on each line. Optionally, the networking system could assume all the features are always to be used if the system at the other end of the communications medium is willing.

Any system with extended features should set the bit stating that an extended feature is available and must check the bit in the response from the other side. When a specific bit is set by both systems, then the extended NJE features are used between them.

2.1.1.2 Implementation Example

Each system has a systems supported features word (SSFW) in a common area. This word contains the bit mask describing all the additional features that this system is able to support. In addition, each system adds an authorized features word (AFW) to a control block that is unique to each line. This word contains the features that are to be enabled for the given line.

At signon, the implementing system performs logic equivalent to the following:

- Primary

  1. Copy SSFW into NCCIFEAT before transmitting 'I' record.

- Secondary

  1. Copy NCCIFEAT from 'I' record into AFW.
  2. AND AFW with SSFW storing result in AFW.

3. Copy AFW into NCCIFEAT of 'J' record.

- Primary

1. Copy NCCIFEAT from 'J' record into AFW.

2. AND AFW with SSFW and store result in AFW.

3. Compare AFW with NCCIFEAT. If any difference, terminate the line.

The above assumes the NCCI is expanded before processing is performed. Systems not supporting this feature will have NCCIDL set to a value too small to include the NCCIFEAT bytes. The system receiving such a type 'I' or 'J' record must then assume that none of the optional features are to be used on the session.

2.1.1.3 Protocol Description

The system transmitting the type 'I' signon record (primary) must set the features bits in the type 'I' record showing the features that are both present and desired (the system programmer may have the option of not enabling a given feature on a given line).

When a type 'I' record is received, the secondary must look at the extended features bits and for any features that it is also willing to use, it must transmit the corresponding feature bit in the type 'J' record. The secondary should then use the features. The primary upon receiving the type 'J' record should then examine the features bits and use any features that are specified.

2.2 Path Management Protocols

Note: This section is only supported by JES2.

Each JES2 member of the NJE network has a Network Path Manager. The Path Manager is responsible for connection protocol between the members of the network, promulgating connection information to the other members, maintaining information about which lines should be used to reach a given node, and informing other subcomponents which nodes should be reached over a given line. The Network Path Manager provides routing information for jobs and messages, processes NJE signon and connection/disconnection records from other Network Path Managers, and makes "best choice" decisions for line selection based on resistances specified by system programmers.

Each system with the Network Path Manager provides a mechanism to prevent connection and disconnection records from being sent about certain connections. In JES2 this mechanism is the CONNECT card in the system initialization deck.

Note: Any system without a Network Path Manager must discard NPM records other than those of type 'B', 'I' or 'J'. Any attempt to forward them without having a full Network Path Manager may result in records looping through the network.

2.2.1 Signon Connection Protocol

The protocol for directly connecting two NJE nodes depends upon the capabilities of the two Path Managers, the number of lines between the two systems, and the installation-supplied names of the nodes. Each protocol is described in the following subsections.

A unique connection in an NJE network has 4 basic parts:

1. The identification of the system with the low EBCDIC node name ('low' end).
2. The identification of the system with the high EBCDIC node name ('high' end).
3. The Connection Event Sequence (CES). The CES is the high order four bytes from the time of day clock.
4. The resistance of the connection.

The CES is a binary value that increases each time the 'low' end system initiates or allows a connection. Since the value is ever increasing, Path Managers can decide what information is the most recent, discarding any old connection information. When CES values are assigned the Path Manager insures that the sequence does not go above the current TOD clock value; therefore, the value could possibly overflow in 143 years from the base TOD clock time (same as TOD clock overflow). Because the 'low' end determines the CES, protocols vary depending upon which end initiates a connection. It should be noted that even though a line may be leased no assumption is made that a particular node is at the other end until it identifies itself via the connection protocol.

To keep the CES consistant throughout the network, it is necessary that all systems using the path manager use "GMT" time settings in their TOD clocks, and that these clocks be properly synchronized. This will also affect the accuracy of the reader start time and other SMF information.

2.2.1.1 Full Primary Trunk Protocol ('Low' End Initiates)

This form of signon protocol is used on BSC connections only. Figure 2-3 illustrates "low" end initiated full connection protocol. Note that the 'low' end cannot concur with a primary connection, that is the responsibility of the 'high' end.

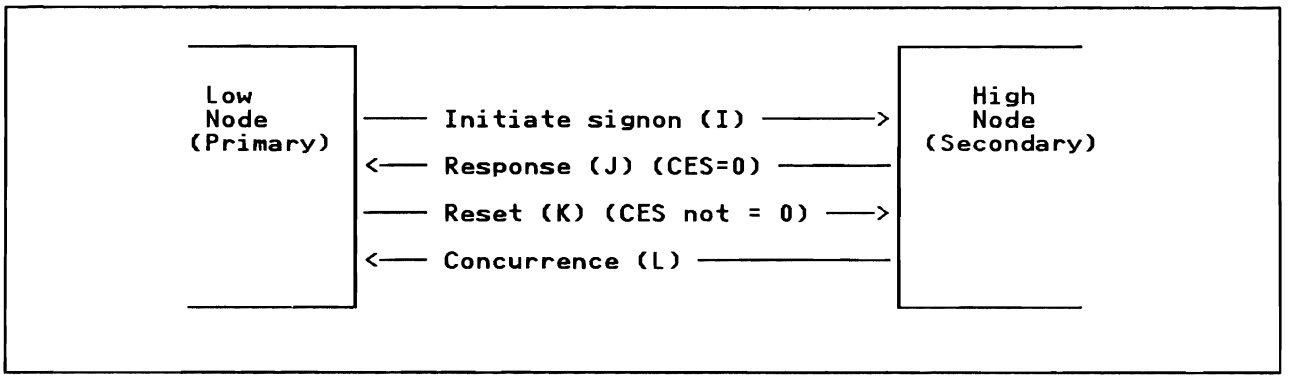

Figure 2-3. Normal Connection (Low Initiates)

2.2.1.2 Full Primary Trunk Protocol ('High' End Initiation)

The 'high' end initiated connection permits a slightly abbreviated protocol.

This signon protocol is always used on SNA sessions, but can also be used on BSC connections. Figure 2-4 illustrates 'high' end initiated full connection protocol.

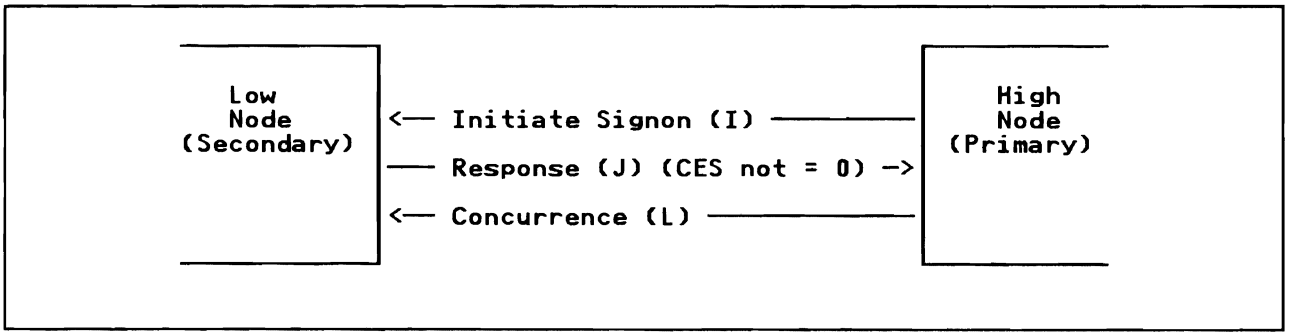

Figure 2-4. Normal Connection (High Initiates)

2.2.1.3 Full Secondary Trunk Protocol

A secondary trunk is a line directly connecting two systems already directly connected when the secondary connection is made and the new line resistance is not less that the original. Since this does not represent a new connection no CES is assigned and no distinction is made between 'low' end or 'high' end. In this case the multi-trunk flag must be set on in the response (J) record. If the new line resistance is less than the original, this becomes the primary trunk with the CES not equal to zero. Figure 2-5 illustrates secondary full connection protocol.

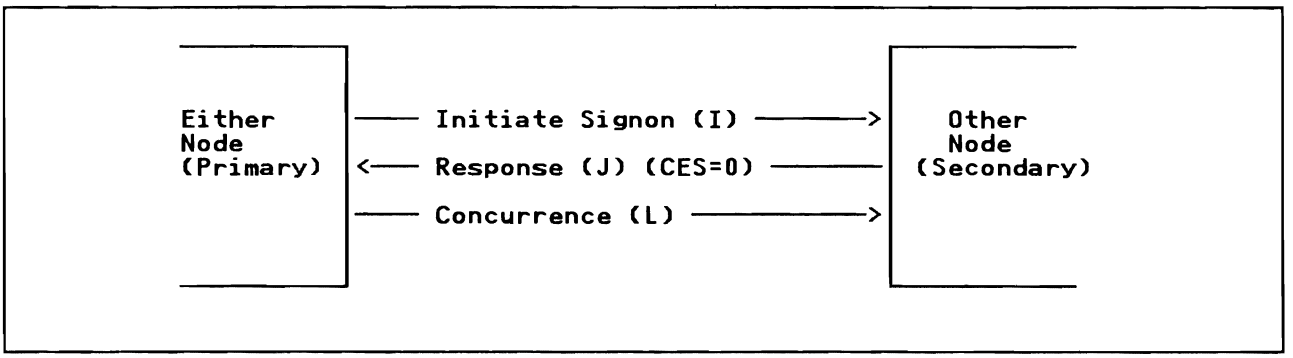

Figure 2-5. Secondary Trunk Connection

2.2.1.4 Full Reset Trunk Protocol

If the 'low' end of a connection determines that the primary trunk of a multi-trunk connection is no longer valid, a reset connection protocol is initiated. The trunk over which the reset control record is transmitted is usually the new primary trunk. The CES value will be set to indicate primary or secondary. Other conditions may cause a reset to be initiated from either end; however, the 'high' end must never require the 'low' end to answer the reset. Figure 2-6 on page 2-6 illustrates primary assignment reset protocol.

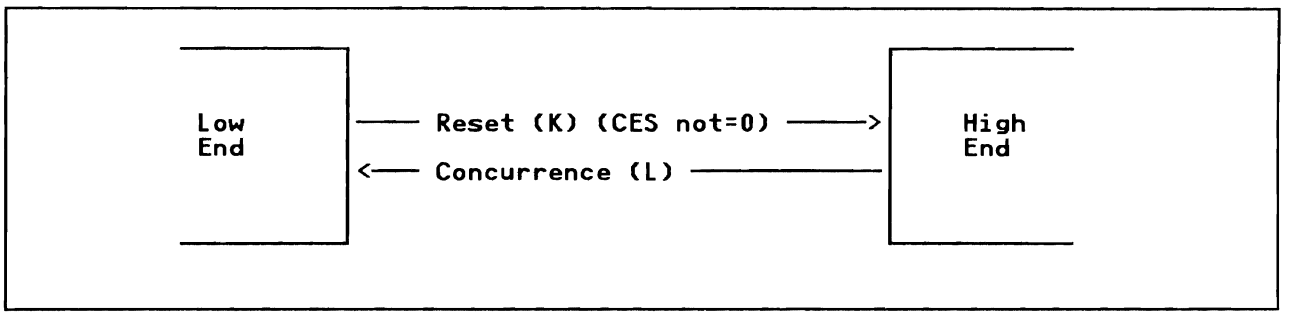

Figure 2-6. Trunk Reset

2.2.2 Pre-Defined Connection Protocol

Pre-defined connections allows connections to occur between a Path Manager system and a system without a Path Manager. It also allows the connection to be known only (private) to the systems connected unless the other systems in the network also define the connection. If one of the connected Path Manager systems pre-defines the connection, the other must also.

Note: Connections between JES2 and any non-JES2 system require pre-defined connections on the JES2 side because non-JES2 systems do not have the complete path management logic. Figure 2-7 illustrates pre-defined connection protocol.

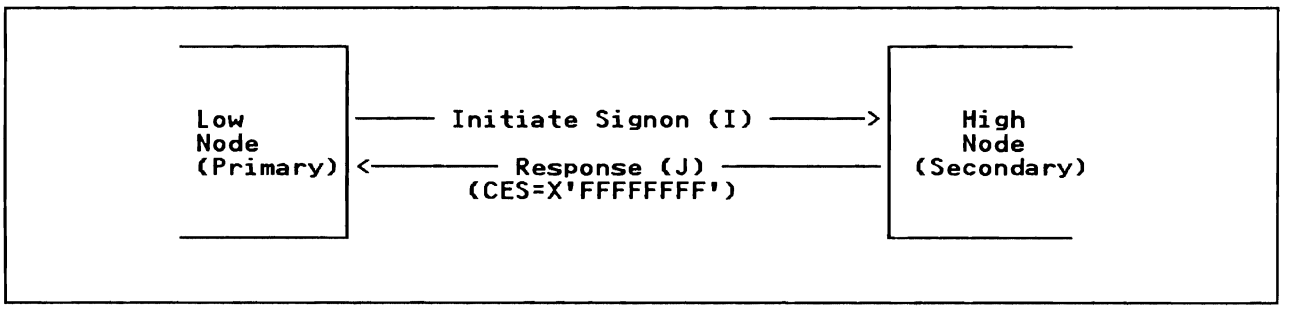

Figure 2-7. Predefined Connection

2.2.3 Connection Status Information

Whenever a dynamic connection is agreed upon, each Path Manager involved will send an Add Connection control record to systems not involved in the connection over all other NJE lines. The add connection control record will be used by receiving Path Managers to determine best paths to nodes within the network. Each Path Manager will forward the add connection control records. If a connection is already known (CES indicates the new control record received is not new), the record is ignored and not forwarded to other network nodes.

Disconnections are promulgated to the members of the network using a Subtract Connection control record. Disconnecting connections may cause nodes formerly reachable via the disconnected line to no longer be available to the system. In this case, dependent connections are automatically determined by each system experiencing the disconnection or receiving the resulting subtraction control records.

Add and Subtract Connection control records may be blocked in the buffer with Reset and Concurrence control records. This will be quite common when a new trunk is established and complete pictures of the network are traded by the systems involved or when a disconnect is received by a JES. See 8.8, "Connection and Control Records" on page 8-65 for the format of the control blocks.

2.2.4 Disconnections

When a NJE line has disconnected, the path manager then clears its own reachable nodes in its tables, validates the queues, and notifies attached nodes that the disconnection has now taken place.

2.3 Signoff (All Systems)

Normal disconnection of an NJE node (using BSC or CTCA) consists of the initiating system sending a final signoff control record. This is a network control record (RCB X'F0', SRCB C'B').

After this record is transmitted, the sending system prepares the line for signon or drains the line as required. The receiving system should prepare for signon after receiving a signoff. Abnormal disconnection of an NJE node consists of the initiating system preparing the line for signon or draining the line as required. Either type of disconnection may occur when all transmitting and receiving functions are idle or when the functions are active. See B.3.2.4, "Signoff" on page B-24 for RSCS specific processing.

Normal disconnection of an NJE node, in SNA uses a TERMSESS if the secondary or CLSDST if the primary.

Chapter 3. Data Control Information

3.1 Overview of a Complete Job

Each job being transmitted for execution (SYSIN) and each collection of output data sets being transmitted for printing/punching (SYSOUT) is preceded by a NJE Job Header record and followed by a NJE Job Trailer record. For SYSOUT transmission, each data set within the collection is preceded by one or more NJE Data Set Header records.

NJE Job Headers, Data Set Headers, Job Trailers, and SYSOUT records (which can be 32,760 bytes long) must be broken into individual records of no more than 256 bytes each for transmission. SYSIN data records are limited to 255 bytes, but are assumed to be 80 byte fixed length records unless a Dataset Header Record Characteristics Change Section is transmitted containing the new characteristics. See section 5.5.2, "Spanned Data Record Format" on page 5-15 for details of how spanned records are sent. These individual records should be blocked/deblocked into transmission buffers.

Before a Job Header record is sent a NJE system must request permission to transmit. If the request is denied by the receiver, nothing is sent. See section 5.4, "Stream Control" on page 5-8 for details. If the request is accepted, the Job Header is sent followed by the records which make up the job (for SYSIN) or by Data Set Header(s) and output records (for SYSOUTs). When all data has been sent, a Job Trailer record is sent and followed by an end of file.

Upon receipt of positive acknowledgment to the end of file, the transmitting system is relieved of responsibility for that job and it can be purged from the spool system of the sending node.

In summary, a NJE system will send SYSIN as shown in Figure 3-1, or Figure 3-2 on page 3-2, depending on the logical record length of the SYSIN stream.

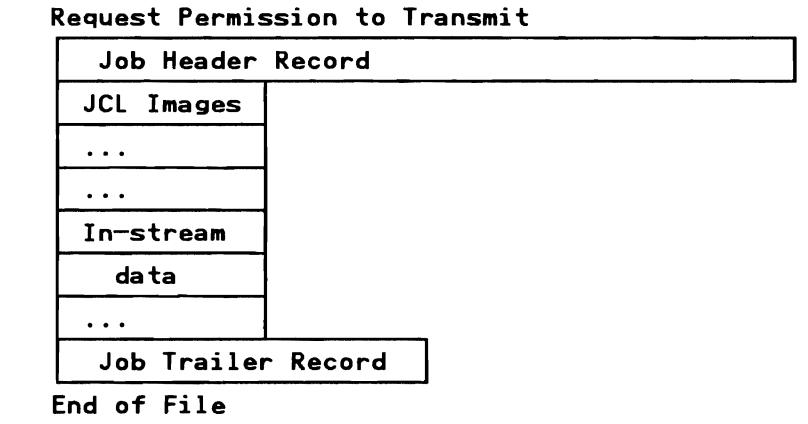

Figure 3-1. Sample MVS SYSIN Data Stream. All Records Fixed, 80-bytes

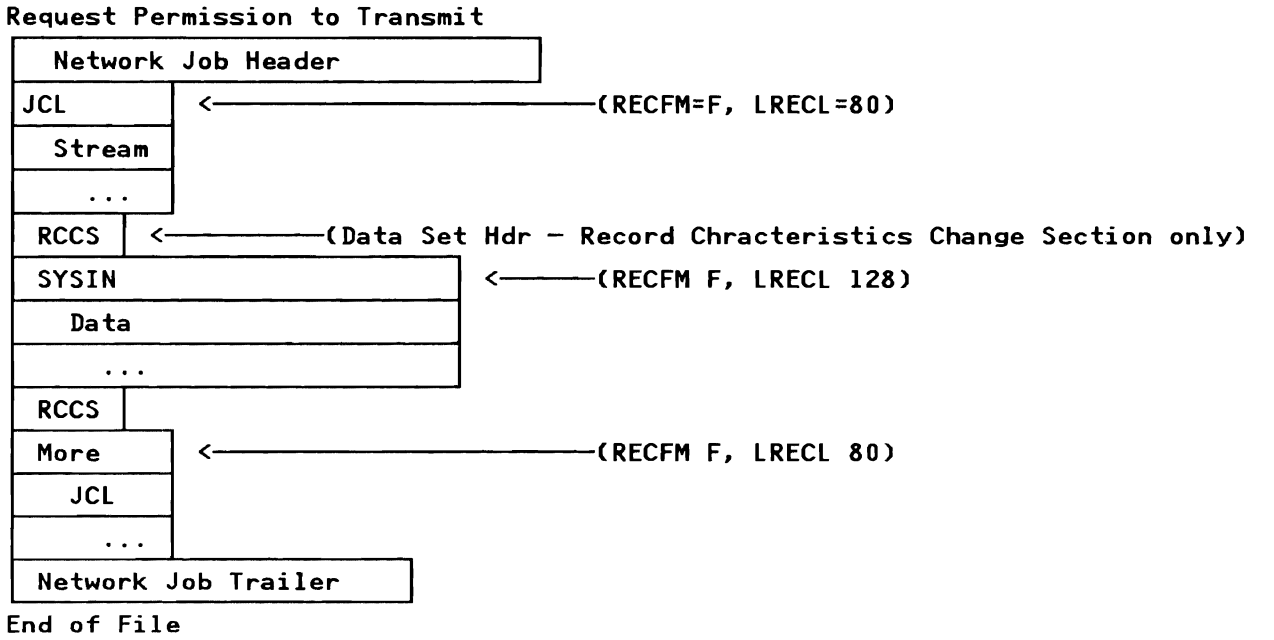

Figure 3-2. Sample MVS SYSIN Data Stream. Multiple SYSIN Data Record Lengths

Systems which do not have job awareness are not concerned with the input job image. They must store and forward such images correctly. However, it is the responsibility of the originator of a job to set this image up correctly so it can be understood by the executing system.

An example of how three output data sets for a job could be sent is shown in Figure 3-3 on page 3-3. In this example, two copies of the third dataset are desired, each copy going to a different destination. This is shown by two Dataset Headers with no data between them. In this case, when a given node determines that data for the two destinations can no longer be sent along the same path, that system logically duplicates the Job Header, Job Trailer, and data to create a second image of the job to be sent along the other path. Multiple consecutive Dataset Headers are not architecturally limited. Currently, only JES2 transmits multiple data set headers in front of a single copy of the data set. See "Multiple Data Set Headers" on page B-17, "Multiple Data Set Headers" on page B-30, and "Multiple Data Set Headers" on page B-44 for system specific processing.

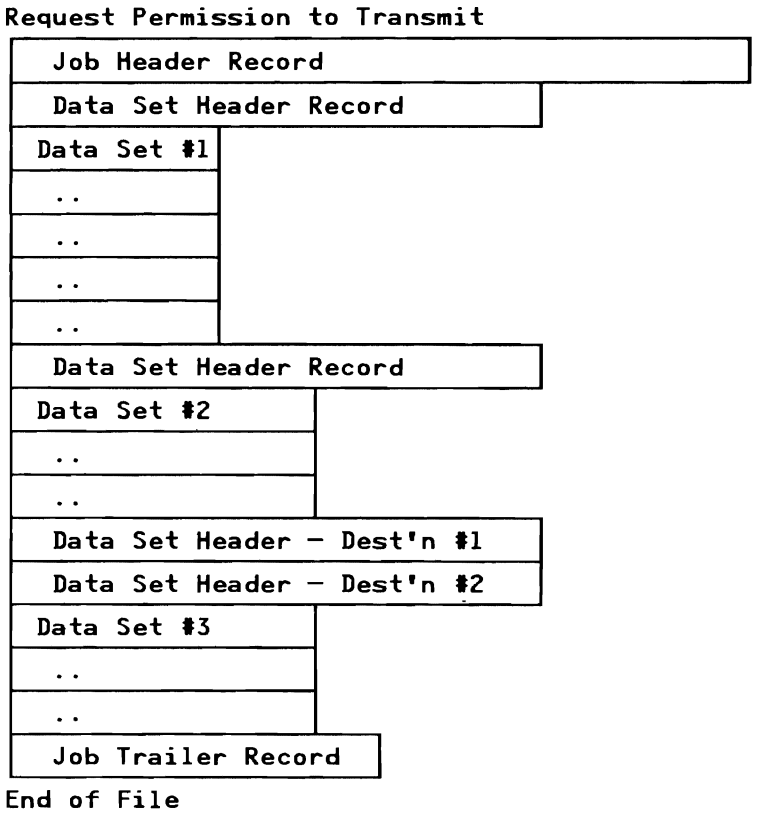

Figure 3-3. Sample SYSOUT Data Stream

3.2 Control Record Sections

The NJE Job Header, Job Trailer, and Data Set Header contain information which is required for job routing, execution, printing, and accounting.

The information in these records is maintained as a job passes from one node to another. Each type of record is identified by a special SRCB.

Note: Each Job Header, Data Set Header, and Job Trailer must be the first record in its transmission block.

3.2.1 Control Record Processing

All three control records (Job Header, Data Set Header, and Job Trailer) are constructed using the same basic format. The format consists of multiple variable-length sections. The first section (the general section) is always present. In the case of the Data Set Header, the Record Change Characteristics Section (RCCS) is also considered a "General" type of section and appears alone as the only section in the header. All other sections are optional. Each section has a prefix showing the length of the section, and the record itself has a prefix which shows the length of the entire record.

The order in which sections appear in a control record is not restricted with the exception of a rule that the first section in a header must be the general section with a modifier of X'00' or X'40'. (See 3.2.4, "Basic Section Flags" on page 3-5 for a discussion of modifiers.) After this, sections may appear in any order and each NJE system must be able to forward all sections correctly as well as recognize and use the sections it

needs no matter where they are in the header. The record characteristics change section of the dataset header must appear without the general section of the dataset header, and must be the first section. The record characteristics change section must appear if SYSIN data of length other than 80 is being transmitted. SYSIN data is limited to 255 bytes.

3.2.2 Control Record Format

A control record begins with the length of the record and the architecture limits it to 32760 bytes. See Appendix B, "System Dependent Considerations and Comments" on page B-1 for specific system limitations. The length of the control record includes the 4 bytes of length, flag, and sequence. The flag and sequence bytes are set to X'00' and are not used unless the record must be segmented for transmission. When a control record is segmented for transmission, the length of the control record is replaced by the length of the first segment, and the sequence count is used to indicate which segment of the control record is being transmitted in the record and if this is the last segment. Each subsequent segment has a four byte header inserted at the beginning.

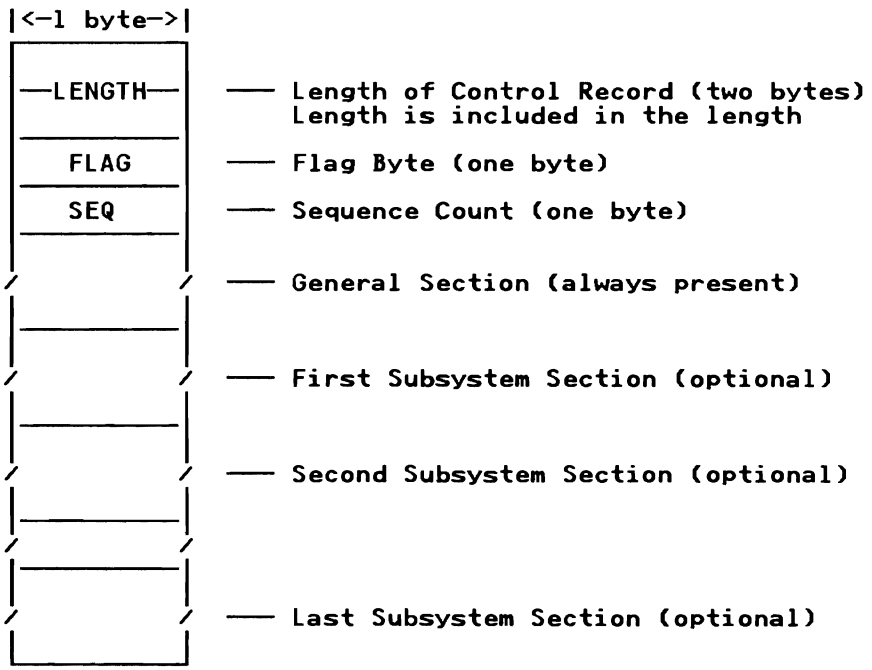

3.2.3 Basic Section Format

Each section of a control record begins with the length of the section. The length of the section includes the two byte length field. The section length is used to locate the start of the next section.

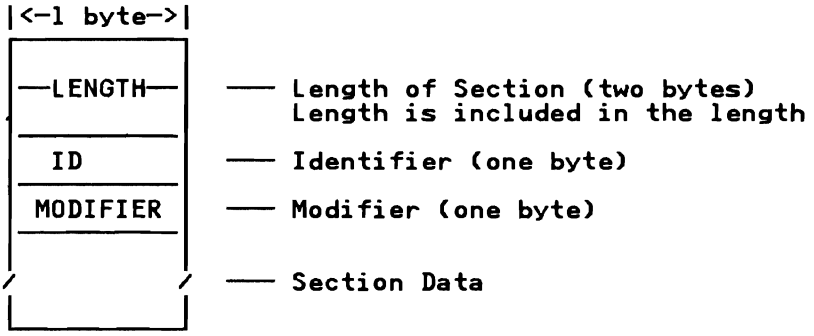

3.2.4 Basic Section Flags

There are two bytes which are used to describe each section of the control record. The first byte, the identifier, is used to show the type of section which follows and also to identify the subsystem (if any) which owns the section. Sections identified as belonging to a specific subsystem may only be changed by that subsystem. Another type subsystem may take information out of such a section if it has already been put in by the proper type subsystem, but that is not recommended as individual subsystem sections are not guaranteed to remain upward compatible. All subsystems can read and write information in the general section. In general, once a control record is generated, it should not be updated by subsequent systems.

The first bit of the first identifier is 0 for any general section and 1 for all other sections. A user section is defined by an ID of B'11nnnnnn' (where the user can specify n) to indicate the user defined sections. An ID of 00 is the general section which must remain standard in order for all systems involved to be able to communicate with each other.

The second flag byte (called the modifier or subtype) can be used to further define the use of a section. For example, the modifier of X'80' is used to represent the 3800 characteristics section of the dataset header. This section is optional unlike the required general section (modifier X'00').

3.3 Header and Trailer Record Segmenting

Each transmitted NJE control record segment is limited to 256 bytes maximum, including 4 bytes for the length/flags/sequence fields at the start of the record. Thus if the 'Overall Length' of the Control Record (Header or Trailer) is more than 256, it must be transmitted in pieces. See B.1.4.2, "Job Header" on page B-4 for JES2 specific information.

The discussion below shows how long headers are transmitted by example.

Control records up to 256 bytes would have the high order bit of the sequence number reset (zero) and would thus be sent as a single segment with sequence number of X'00'.

Suppose we have a 237-byte header:

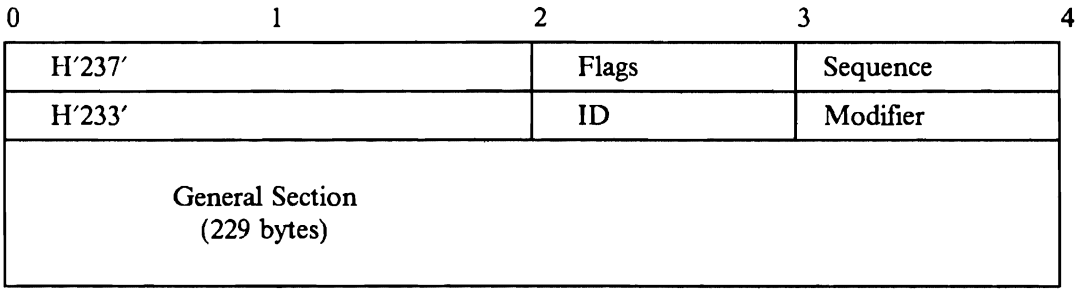

This would be sent just as it appears above, with no segmenting.

Now assume we have a 600-byte header:

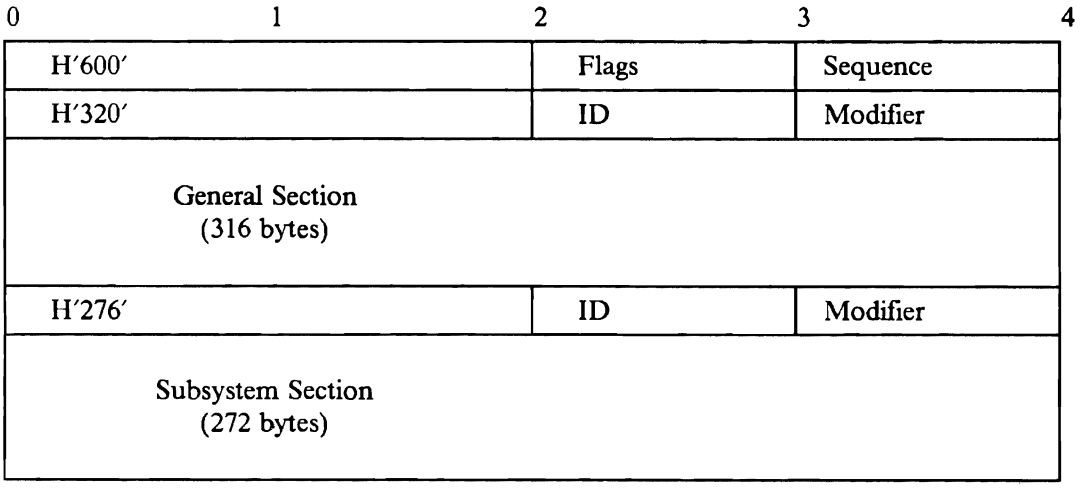

This would be sent as:

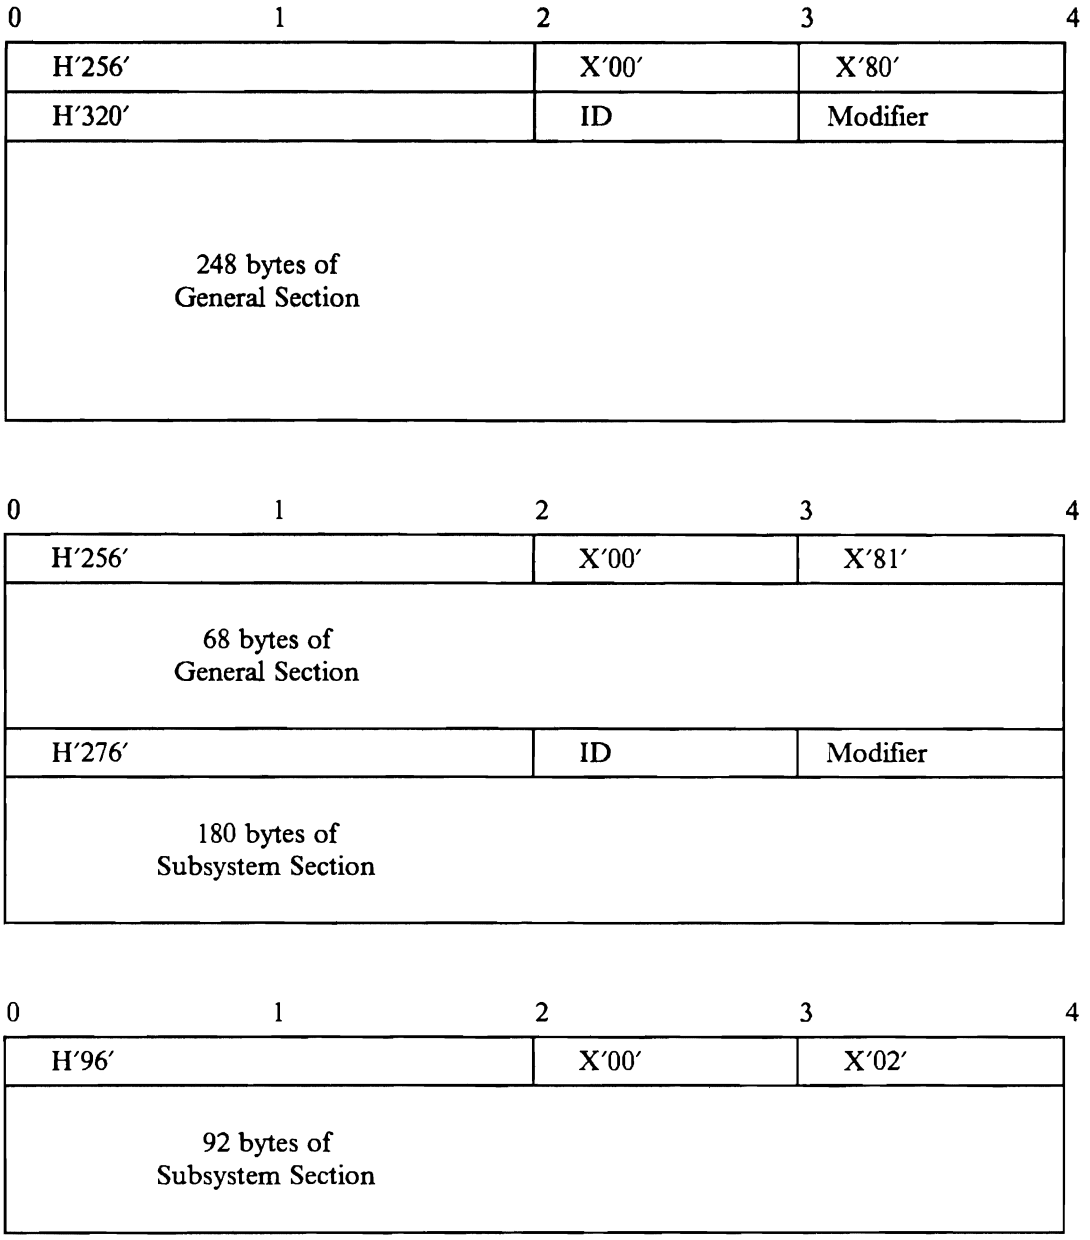

The Header is, in effect, 're-blocked' back to 600 bytes by the receiving system.

Sequence byte (fourth byte in examples above with values of X'80', X'81', X'02') is defined as follows:

1. High-order bit on means more blocks to follow
2. Low-order seven bits are sequence counter starting at zero.

See section "Segmented Headers" on page B-26 for restrictions in earlier RSCS releases.

3.4 Job Header

The job header is used to begin a unit of data sent intact between two nodes. This transmission unit ends with a job trailer. Job headers are used for both input and output files. Input files may be considered "jobs" in the OS/VS concept. They contain instructions which allow execution on a receiving node. Not all NJE systems have a job awareness in that they understand the term "execution" themselves. Such systems may consider jobs as input files. Output files are not executed but are still delimited by job headers and trailers.

The Job Header contains four types of information:

- Identification (job name, JOBID, programmer's name)
- Routing Control (origin and destination node identification, origin and destination userids or remotes)
- Execution Control (job class, estimated line count)
- Network Account Number

The job header is prefixed by a four byte header containing the combined length of all the sections within the header. The prefix looks as follows:

Header Prefix - Required Part

0	NJHLEN	NJHFLAGS	NJHSEQ

The format of the general section of the job header is shown below. Details on what information each subsystem puts into and takes out of each field in the header is shown in section 8.4, "Job Header."

General Section (Identifier X'00') - Required

0	NJHGLEN		NJHGTYPE	NJHGMOD
4	NJHGJID		NJHGJCLS	NJHGMCLS
8	NJHGFLG1	NJHGPRIO	NJHGORGQ	NJHGJCPY
C	NJHGLNCT	RESERVED		
10	NJHGACCT			
14				
18	NJHGJNAM			
1C				
20	NJHGUSID			
24				
28	NJHGPASS			
2C				
30	NJHGNPAS			

34	
38	NJHGETS
3C	
40	NJHGORGN
44	
48	NJHGORGR
4C	
50	NJHGXEQN
54	
58	NJHGXEQU
5C	
60	NJHGPRTN
64	
68	NJHGPRTR
6C	
70	NJHGPUNN
74	
78	NJHGPUNR
7C	
80	NJHGFORM
84	
88	NJHGICRD
8C	NJHGETIM
90	NJHGELIN
94	NJHGECRD
98	NJHGPRGN
	. . .
A8	
AC	NJHGROOM
B0	
B4	NJHGDEPT
B8	
BC	NJHGBLDG
C0	

C4	NJHGNREC

The formats of the other defined job header sections are shown below:

JES2 Subsystem Section (Identifier X'84')

0	NJH2LEN		NJH2TYPE	NJH2MOD
4	NJH2FLG1	RESERVED		
8	NJH2ACCT			
C	NJH2USID			
10				
14	NJH2USR			
18				
1C	NJH2GRP			
20				
24	NJH2SUSR			
28				
2C	NJH2SGRP			
30				

POWER Subsystem Section (Identifier X'86')

0	NJHPLEN		NJHPTYPE	NJHPMOD
4	NJHPFLG1	NJHPDISP	RESERVED	NJHPSYID
8	NJHPUSER			
		. . .		
14				
18	NJHPDSKT		RESERVED	

Job Scheduling Section (Identifier X'8A')

0	NJHELEN	NJHETYPE	NJHEMOD
4	NJHEPAGE		
8	NJHEBYTE		

User Section (Identifier X'C0')

0	NJHULEN	NJHUTYPE	NJHUMOD
4	NJHUCODE		

Note: Beyond the first four bytes, the user section is undefined by NJE.

3.5 Data Set Header

There are three defined general sections and three specific sections which may appear in a Dataset Header:

- General Section (Modifier X'00') - Required for SYSOUT, but not SYSIN
- General Section (Modifier X'40') - Record Characteristics Change Section (SYSIN only), must be sent if SYSIN data of length other than 80 bytes is to be transmitted. (This is sent without a General Section with Modifier X'00'.)
- General Section (Modifier X'80') - 3800 Characteristics Section (optional)
- POWER Section (Identifier X'86')
- RSCS Section (Identifier X'87')
- Output Processing Section (Identifier X'89') - Optional support for advanced printer data streams. (Also called the "Data Stream" section by JES2 and JES3.)

The dataset header begins with the standard four byte control record prefix containing the combined length of all sections.

Dataset Header

0	NDHLEN	NDHFLAGS	NDHSEQ

Details on what information each subsystem puts into and takes out of each field in the header is shown in section 8.5, "Dataset Header" on page 8-17.

The format of the three defined general sections of the data set header are shown below:

Dataset Header - General Section (Modifier X'00')

0	NDHGLEN	NDHGTYPE	NDHGMOD
4	NDHGNODE		
8			
C	NDHGRMT		
10			
14	NDHGPROC		
18			
1C	NDHGSTEP		
20			
24	NDHGDD		
28			
2C	NDHGDSNO	RESERVED	NDHGCLAS

30	NDHGNREC			
34	NDHGFLG1	NDHGRCFM	NDHGLREC	
38	NDHGDSCT	NDHGFCBI	NDHGLNCT	RESERVED
3C	NDHGFORM			
40				
44	NDHGFCB			
48				
4C	NDHGUCS			
50				
54	NDHGXWTR			
58				
5C	RESERVED			
60				
64	NDHGFLG2	NDHGUSCO	RESERVED	
68	NDHGPMDE			
6C				

Dataset Header - Records Characteristics Change Section (Modifier X'40')

0	NDHCLEN		NDHCTYPE	NDHCMOD
4	NDHCFLG1	NDHCRCFM	NDHCLREC	

Dataset Header - 3800 Section (Modifier X'80')

0	NDHALEN		NDHATYPE	NDHAMOD
4	NDHAFLG1	NDHAFLCT	NDHATREF	RESERVED
8	NDHATAB1			
C				
10	NDHATAB2			
14				
18	NDHATAB3			
1C				
20	NDHATAB4			
24				
28	NDHAFLSH			
2C				
30	NDHAMODF			
34				
38	NDHACPYG			
3C				

The formats of all the other defined data set header sections are shown below:

Dataset Header - POWER Subsystem Section (Identifier X'86')

0	NDHPLEN		NDHPTYPE	NDHPMOD
4	NDHPFLG1	NDHPIDEV	NDHPPRIO	NDHPDISP
8	NDHPUSER			
	. . .			
14				
18	NDHPJBSF	NDHPSYID	NDHPNSEP	NDHPOPTN
1C	NDHPPART		RESERVED	
20	NDHPRCFM	RESERVED	NDHPJNUM	
24	NDHPCOMP			
28	NDHPPASS			
2C				
30	NDHPSETP			
	. . .			
70				
74	NDHPSTRT			
78				

Dataset Header - RSCS Subsystem Section (Identifier X'87')

0	NDHVLEN		NDHVTYPE	NDHVMOD
4	NDHVFLG1	NDHVCLAS	NDHVIDEV	NDHVPGLE
8	NDHVDIST			
C				
10	NDHVFNAM			
	. . .			
18				
1C	NDHVFTYP			
	. . .			
24				
28	NDHVPRIO		NDHVVRSN	NDHVREL

Output Processing Section (Identifier X'89')

Note: This is also called the "Data Stream" section by JES2 and JES3.

0	NDHSLEN		NDHSTYPE	NDHSMOD
4	NDHSFLEN		RESERVED	
8	NDHSJDVT			
C				
10	NDHSNSTR			
14	NDHSGPID			
18				
1C	NDHSPRID			
20	NDHSVERS	NDHSPLEN	NDHSDLEN	
24	NDHSVERB			
28				
2C	RESERVED			
30				
34	NDHSFLG2	RESERVED	RESERVED	
38	TEXT BLOCKS			
	.			
	.			
	.			

Note: The Output Processing Text Blocks are variable in length.

User Section (Identifier X'C0')

0	NDHULEN	NDHUTYPE	NDHUMOD
4	NDHUCODE		

Note: Beyond the first four bytes, the user section is undefined by NJE.

3.6 Job Trailer

The job trailer is prefixed by a four byte header containing the combined length of all the sections within the trailer.

Job Trailer

0	NJTLEN	NJTFLAGS	NJTSEQ

The format of the job trailer general section is shown below. Details on what information each subsystem puts into and takes out of each field in the header is shown in section 8.6, "Job Trailer" on page 8-46.

Job Trailer General Section (Modifier X'00')

0	NJTGLEN		NJTGTYPE	NJTGMOD
4	NJTGFLG1	NJTGXCLS	RESERVED	
8	NJTGSTRT			
C				
10	NJTGSTOP			
14				
18	RESERVED			
1C	NJTGALIN			
20	NJTGACRD			
24	RESERVED			
28	NJTGIXPR	NJTGAXPR	NJTGIOPR	NJTGAOPR

The format of other defined sections in the job trailer follows:

Accounting Section (Identifier X'89')

0	NJTSLEN	NJTSTYPE	NJTSMOD
4	NJTSAPAG		
8	NJTSABYT		

User Section (Identifier X'C0')

0	NJTULEN	NJTUTYPE	NJTUMOD
4	NJTUCODE		

Note: Beyond the first four bytes, the user section is undefined by NJE.

Chapter 4. Network Command and Message Formats

4.1 Nodal Message Record

The nodal message record (NMR) is used for transmitting both commands and messages in a NJE network. It is associated with an RCB of X'9A'. Command responses are treated as messages in the Nodal message record. The NMR record consists of a header which contains control information about the command or message and then the message or command text itself. Some of the fields in the header are used differently depending on whether the record is defining a command or message. For example NMROUT refers to the origin for commands and the destination for messages. In addition, the data represented in NMROUT differs based on the bits set in NMRFLAG. A detailed description of the NMR record and fields in it is contained in section 8.7, "Nodal Message Record" on page 8-50.

The rules for sending messages and commands in a NJE network are different than those for sending files. Command and message data do not have to be stored at an intermediate node. Such data consequently does not have to be saved if no paths to the destination node are available. When discarding a command, a message should be generated to the originating system to inform the person issuing the command that it has been discarded. Command and message data may also be discarded without a diagnostic due to disastrous errors to systems along the path to the destination.

While the format of the NMR is shown below, this diagram does not show all the field redefinitions that can occur within the NMR. Section 8.7, "Nodal Message Record" on page 8-50 should be used for a detailed description.

Nodal Message Record

0	NMRFLAG	NMRLEVEL	NMRTYPE	NMRML
4	NMRTONOD			
8				
C	NMRTOQUL	NMROUT		
10				
14		NMRFMNOD		
18				
1C		NMRFMQUL	NMRMSG	
			. . .	
A0				

The NMRMSG field may take many values depending on the settings of NMRFLAG. NMRMSG may optionally begin with the userid of the message originator for messages. For commands, NMRMSG may contain EBCDIC command text, or a formatted command (an encoded command request).

See B.3.3, "Commands & Messages (NMRs)" on page B-25 for a description of RSCS handling of formatted commands.

Chapter 5. Multi-Leaving (ML) Functions - SNA, BSC, CTCA

This section describes those multi-leaving functions that are used in NJE communications. Where differences exist when using SNA versus BSC or CTCA communications, those differences will be explicitly noted.

Multi-leaving provides for the time division multiplexing (interleaving) of independent data streams. In NJE using SNA, the application interleaves job streams on a single SNA session. The communications software (e.g. VTAM) does an independent multiplexing of sessions onto a physical link that is unknown to the application.

By contrast, NJE using BSC or CTCA multileaving also includes the functions of link-level multiplexing and block sequence checking.

This section will discuss the control information for defining streams and the mechanisms available for initiating, terminating, suspending, and resuming streams that are essentially common to NJE using either SNA or BSC or CTCA communications. See Chapter 6, "Additional Multi-Leaving (ML) Functions - BSC, CTCA" for other functions used only for NJE in a BSC or CTCA environment.

5.1 Buffer Formats

The transmission buffer should be at least 300 bytes long. This minimum size is to allow a complete 256 byte record along with the compression bytes, BSC control bytes, and multi-leaving control bytes. There is no defined maximum size. Each system has its own maximum limit. In general, the value must be set based on the communications mechanism. A BSC connection would use a lower buffer size than a CTCA since the time to retransmit and the probability of errors on a BSC connection are greater than on a CTCA.

5.1.1 Non-SNA Buffer Format

For multi-leaving transmissions, a variable number of records can be combined into a transmission buffer. Each record in the buffer is comprised of a series of character strings each prefixed by a string control byte (SCB). To group records of various media together in a single buffer, each record is prefixed by a record control byte (RCB) and a subrecord control byte (SRCB). To control the flow of individual streams, a function control sequence (FCS) is added to each transmission block. Finally, a block control byte (BCB) is added as the first character of each buffer for error detection and correction.

Following is the layout of a typical multi-leaving transmission block for BSC and CTCA communications.

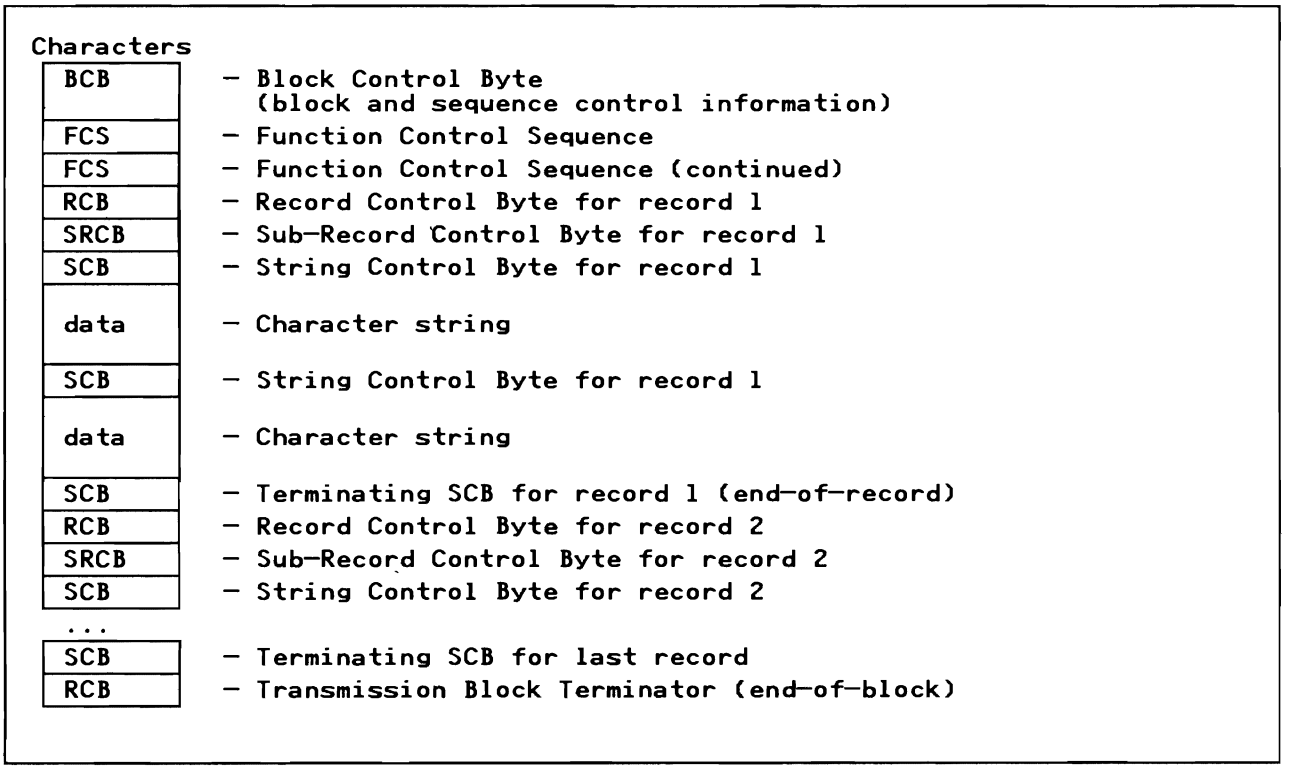

Figure 5-1. Non-SNA Multi-Leaving Transmission Block

5.1.2 SNA Buffer Format - Request/Response Unit (RU)

The SNA 'data buffer' is called an RU (Request/Response Unit). The RU carries control information and data between logical units. A control RU contains a request or acknowledgement, while a data RU contains FMHs, SCBs and data. The RUs are sent as only-in-chain (OIC) elements in exception response mode. A data RU may contain as many NJE records as can fit into the RU. The maximum size of an RU may be 65,535 bytes, but the size used on an NJE session is determined by the BIND parameters. The maximum length for the NJE record is 259 bytes (256 bytes of user data plus 3 byte NJE RID).

In NJE (using SNA), a 3 byte header is prefixed to each record. This prefix is called the Record Identifier (RID). See 5.4.1.1, "Record Identifier (RID)" on page 5-9 for a detailed description of the RID. See "RU Multiplexing" on page B-10 for JES2 specific information.

Following is the layout of a typical multi-leaving transmission block for SNA communications.

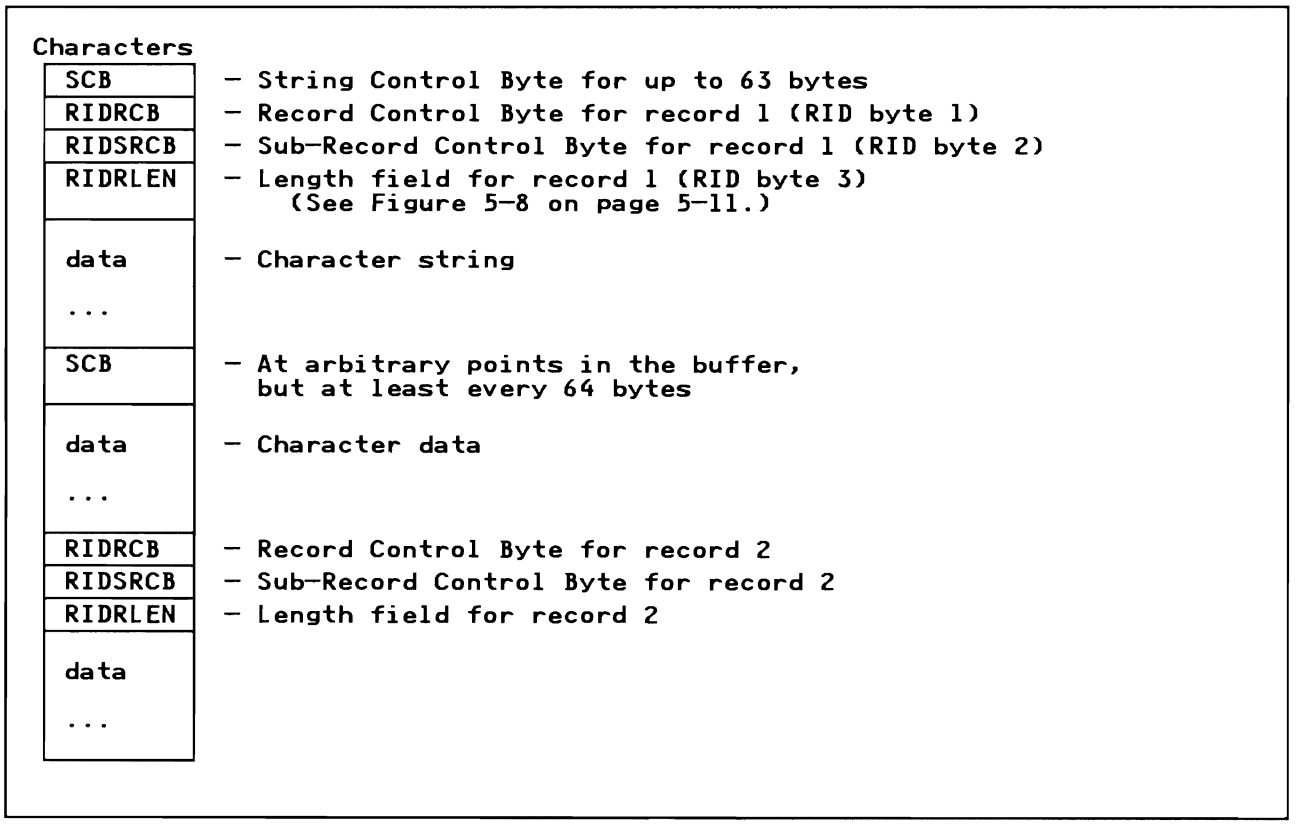

Figure 5-2. SNA Multi-Leaving Transmission Block

Note the contrast with Non-SNA use of SCBs. For SNA buffers, SCBs are used to compress/compact the entire buffer; i.e., the entire buffer is treated as data as far as the SNA transport subsystem is concerned.

The *only restriction* on the placement of SCBs is that an SCB must appear at least every 64 bytes, and describe up to 63 bytes of following data. This restriction is identical for non-SNA.

5.2 Record Control Byte (RCB) Function

The Record Control Byte (RCB) begins each logical record within the network. Throughout this document unless stated otherwise the term record will imply a string of bytes beginning with an RCB. The end of record in SNA is defined by the length in the Record Identifier (RID). In non-SNA transmissions, the end of the record is defined by a null String Control Byte (SCB) for compressed records, or by the data length byte for non-compressed records (signon, signoff, and path manager records).

A transmission block may not contain records with different RCB's. The session must be terminated (all streams) if a transmission buffer is received that contains an unexpected, unrecognized or invalid RCB.

This includes:

- RCBs for streams that have not been started.
- RCBs for different streams in the same buffer.
- Undefined RCB values.

Furthermore, the same error action is taken for valid RCBs that are received out of sequence, for example, a X'B0' receiver cancel with an unstarted stream referenced in the SRCB would be an error situation.

Requests to start a stream *that has already been started* are not handled with these rules. In these cases, the requests are rejected with the X'B0' (permission denied) RCB. The session is not terminated in these cases. Instead, the transmitting system should terminate the stream upon receipt of the X'B0'.

Binary	Hex	Meaning
0000 0000	00	End-of-block (BSC)
rrrr rrrr	01-8F	Reserved
1001 0000	90	Request to initiate stream[1]
1010 0000	A0	Permission to initiate stream [1]
1011 0000	B0	Negative permission or receiver cancel [1]
1100 0000	C0	Acknowledge transmission complete [1]
1101 0000	D0	Ready to receive stream [1]
1110 0000	E0	BCB sequence error
1111 0000	F0	General control record
1001 0001	91	RJE console message
1rrr 0001	A1-F1	Reserved
1001 0010	92	RJE operator command
1rrr 0010	A2-F2	Reserved
1iii 0011	93-F3	RJE input record
1iii 0100	94-F4	RJE print record
1iii 0101	95-F5	RJE punch record
1iii 0110	96-F6	Data set record
1rrr 0111	97-F7	Reserved
1iii 1000	98-F8	SYSIN record
1iii 1001	99-F9	SYSOUT record
1001 1010	9A	operator command/console message
1rrr 1010	AA-FA	Reserved
1rrr 1011	9B-FB	Reserved
1rrr 1100	9C-FC	Reserved
1rrr 1101	9D-FD	Reserved
1rrr 1110	9E-FE	Reserved
1rrr 1111	9F-FF	Reserved

i represents 1 or 0 provided this produces a value within the range shown in the hex column.

iii may be from 1 to 7 and corresponds to the stream number.

r represents 1 or 0 provided this produces a value within the range shown in the hex column. All r values are reserved.

Figure 5-3. RCB Definition

See "Receiver Initiated Processing" on page B-18 for JES3 specific information. See B.3.6, "Stream Support and Control" on page B-32 and B.2.6.1, "Multiple Streams" on page B-18 for system dependent stream support.

[1] The SRCB is set to the RCB of the stream to be initiated.

5.3 Sub Record Control Byte (SRCB) Function

The Sub Record Control Byte (SRCB) is interpreted as follows depending on the RCB value. For a definition of each RCB see 5.2, "Record Control Byte (RCB) Function" on page 5-4.

```
RCB       SRCB

00        None

90        RCB of stream to be initiated

A0        RCB of stream to be initiated

B0        RCB of stream to be cancelled or rejected

C0        RCB of stream which is complete

D0        RCB of stream receiver which is ready

E0        Expected count - BCB Sequence Error (received count is in BCB)

F0        An identification character as follows: [2]
             A = Initial RJE SIGNON [3]
             B = Final RJE/NJE SIGNOFF
             C = Print initialization record [3]
             D = Punch initialization record [3]
             E = Input initialization record [3]
             F = Data set transmission initialization [3]
             G = System configuration status [3]
             H = Diagnostic control record [3]
             I = Initial network SIGNON
                 Must be only record in transmission block
             J = Response to initial network SIGNON
                 Must be only record in transmission block
             K = Reset network SIGNON
             L = Accept (concurrence) network SIGNON
             M = Add network connection
             N = Delete network connection
           O-R = Reserved for future use
           S-Z = Unused

91        1000 0000 (X'80') [3]

92        1000 0000 (X'80') [3]

93-F3     1000 0000 (X'80') [3]

94-F4     Carriage control information (for RJE) as follows: [3]
             1010 00nn - Space immediately 'nn' spaces (not used)
             1011 cccc - Skip immediately to channel 'cccc' (not used)
             1000 00nn - Space 'nn' lines after print
             1000 1100 - Load printer FCB image
             1001 cccc - Skip to channel 'cccc' after print
             1000 0000 - Print and suppress space
```

Figure 5-4. SRCB Definition - Part 1

---

[2] F0 records are not compressed. Except where noted (I and J), multiple F0 records may be blocked in a single transmission. F0 records may not be blocked with other records within a transmission.

[3] These codes are assigned to RJE. Their function is not described further in this document.

```
RCB       SRCB

95-F5     1000 1111 (X'8F') [3]

96-F6     Undefined

97-F7     Undefined

98-F8     NJE input control information as follows:
             1000 0000 - Standard record
             1100 0000 - Job header
             1110 0000 - Data set header
             1101 0000 - Job trailer
             1111 0000 - Reserved

99-F9     NJE SYSOUT control information as follows:
             10cc 0000 - Carriage control type as follows:
             ---- ----
             1000 0000 - No carriage control
             1001 0000 - Machine carriage control
             1010 0000 - ASA carriage control
             1011 0000 - CPDS page mode records (with carriage control) [4]

             10cc ss00 - Spanned record control as follows:
             ---- ----
             1000 ..00 - No carriage control
             1001 ..00 - Machine carriage control (also machine
                         opcode in CCW)
             1010 ..00 - ASA carriage control [5]
             1011 ..00 - CPDS page mode records (with carriage control) [4]
             10.. 0000 - Standard record (not spanned)
             10.. 1000 - First segment of spanned record
             10.. 0100 - Middle segment of spanned record
             10.. 1100 - Last segment of spanned record

             11cc 0000 - Control record as follows:
             ---- ----
             1100 0000 - Job header
             1110 0000 - Data set header
             1101 0000 - Job trailer
             1111 0000 - Data set trailer (not used)

9A        Operator Command/Message (NMR)
             1000 0000 (X'80') [6]
```

Figure 5-5. SRCB Definition - Part 2

See "Receiver Initiated Processing" on page B-18 for JES3 specific information. See B.1.2.1, "Network Connection & Control Records" on page B-2, "Path Manager Records" on page B-13, B.3.8.2, "Wait-a-bit Processing" on page B-33, and B.4.2.1, "Network Connection & Control Records" on page B-39 for specific system information.

Control records must be the first record in a transmission block. The trailer record must be the first and only record in a transmission block.

---

4 RSCS Networking does not support this until Version 2.1.

5 See section "Spanned Records" on page B-30 for RSCS deviation.

6 JES2 sends a SRCB of X'00' with NMRs.

Note: The carriage control attributes of the middle and last segment are ignored because only one carriage control character may exist in a spanned record and it must be defined to exist in the first segment.

All SRCB values must be forwarded, even if not understood. An End Node must try to handle all data in the best way possible for that node.

5.4 Stream Control

When a system wants to begin transmitting a unit of work, it prepares a Request to Initiate Stream (RCB X'90') control record. The receiving system then determines if it is willing to receive the unit of work. If reception is to be permitted the receiving system responds with a Permission to Initiate Stream (RCB X'A0') control record. Otherwise the receiving system responds with Negative Permission (RCB X'B0').

No jobs are to be sent until the proper control records are exchanged. The sender may only send a job after the receiver has responded with a permission to initiate stream (RCB X'A0'). Only one job may be sent per request to initiate stream. When a request to initiate a stream that is already active is received, Negative Permission should be returned.

When Negative Permission is returned, the requestor should either wait for an interval of time or until it receives Receiver Initiated (RCB X'D0') for that stream before sending another Request To Initiate Stream.

Note: JES2 and RSCS will hold the job when negative permission is received. POWER and JES3 will drain the transmitter and requeue the job.

5.4.1 Use of RID for SNA Stream Control

All NJE data flows on logical streams, which correspond to transmitter and receiver pairs on the network link. Job and SYSOUT streams are controlled by destination indications in the Record Identifier (RID), which is described in the next section.

Both SYSIN and SYSOUT streams must be initiated before any data can be sent. A request to initiate stream control record (RIDRCB = X'90') is sent to start the SYSIN or SYSOUT stream. If the receiving node has the resources to receive a SYSIN or SYSOUT, a permission granted stream control record (RIDRCB = X'A0') is returned. Permission to allocate a new stream may also be denied (RIDRCB = X'B0'). If permission is granted, the SYSIN or SYSOUT data is sent. Following the EOF stream control record, the sender must wait for the transmitter to send either an acknowledgement of transmission end (RIDRCB = X'C0'), or negative acknowledgement (RIDRCB = C'B0'). These are the same as the RCBs for non-SNA lines.

A SNA stream control record must be the only record in the Request/Response Unit (RU). See 5.1.2, "SNA Buffer Format - Request/Response Unit (RU)" on page 5-3 for additional information about the RU.

5.4.1.1 Record Identifier (RID)

A record identifier, or a RID, is a 3 byte header that is in the front of every NJE (using SNA) record. The RID identifies the type of record that is to follow and the destination of the record. Its purpose is similar to that of the RCB and SRCB in BSC NJE

The RID describes three types of records: network topology records, stream control records and data records. For data records, the RID also contains the length of the data in that record. The maximum data length is 256 bytes. The RID is composed of three bytes. The first byte is the RIDRCB. Its format is shown in Figure 5-6.

Byte	Value	Record type	Meaning
0	X'90'	stream control	Request to allocate SYSIN/SYSOUT stream
	X'A0'	stream control	Permission to allocate stream granted
	X'B0'	stream control	Permission to allocate stream denied or receiver cancel
	X'C0'	stream control	Acknowledge end of transmission
	X'D0'	stream control	Receiver ready
	X'F0'	network topology	An NJE topology record follows
	B'iiii1000'	data	An NJE SYSIN stream data record follows
	B'iiii1001'	data	An NJE SYSOUT stream data record follows
	X'9A'	data	A nodal message record follows

Figure 5-6. RIDRCB: Byte 1 of RID

Note: For all stream control records, the RIDSRCB = RCB of the stream.

The 'iii' in the value field for RIDRCB identifies the particular SYSIN or SYSOUT stream. Valid values are 1-7. A value of zero is not allowed.

The second byte is the RIDSRCB. Its format is dependant on the value of the RIDRCB and is shown in Figure 5-7. The RIDSRCB contains the same values as the SRCB for non-SNA lines.

RIDRCB	RIDSRCB
X'90'	RCB of SYSIN/SYSOUT stream to be allocated
X'A0'	RCB of SYSIN/SYSOUT stream for which permission to allocate has been granted
X'B0'	RCB of SYSIN/SYSOUT stream to be canceled or for which permission to allocate has been denied
X'C0'	RCB of SYSIN/SYSOUT stream for which end of transmission has been acknowledged
X'D0'	RCB of SYSIN/SYSOUT receiver that was initiated
X'F0'	EBCDIC character identifying NJE topology record
X'98'-X'F8'	Transmission end information X'00' - standard SYSIN stream end X'40' - SYSIN stream cancelled by transmitter NJE SYSIN control information (See Figure 5-5 on page 5-7 for details)
X'99'-X'F9'	Transmission end information X'00' - standard SYSOUT stream end X'40' - SYSOUT stream cancelled by transmitter NJE SYSOUT control information (See Figure 5-5 on page 5-7 for details)
X'9A'	X'80' [7]

Figure 5-7. RIDSRCB: Byte 2 of RID

The NJE SYSIN and SYSOUT control information and topology record bit assignments are identical to those for the multileaving SRCB so they have not been repeated here.

[7] JES2 sends a SRCB of X'00' with NMR's.

The third byte of the RID is the RIDRLEN. Its value depends on the value of RIDRCB and is shown in Figure 5-8.

RIDRCB	RIDRLEN
X'90'	N/A(set to zero)
X'A0'	N/A(set to zero)
X'B0'	1 if reason code present, else set to zero
X'C0'	N/A(set to zero)
X'D0'	N/A(set to zero)
X'F0'	Length of network topology plus 3 (for RID)
X'98'-X'F8'	Length - 1 of data record [8]
X'99'-X'F9'	Length - 1 of data record [8]
X'9A'	Length - 1 of nodal message record [8]

Figure 5-8. RIDRLEN: Byte 3 of RID

5.4.2 Use of FCS for Non-SNA Stream Control

The sender should only activate one stream at a time that uses a given FCS stream suspension bit. The receiver should respond with Negative Permission for any attempted invalid requests (e.g. simultaneous use of an FCS bit for two streams).

Note: The restrictions on the number of streams that may be concurrently active are documented in 6.2, "Function Control Sequence (FCS) Function" on page 6-4. Some products implement the same restrictions for SNA stream control, although NJE does not impose them. This is an acceptable product choice.

[8] A value of zero in RIDRLEN is used to indicate that no data follows. The values X'01' - X'FF' represent record lengths of 2 - 256. Since NJE records begin with a one byte length, a record length of 1 need not be represented.

5.4.3 Normal Stream Operation and Termination

Normal operation in a single stream includes the following sequence repeated over and over:

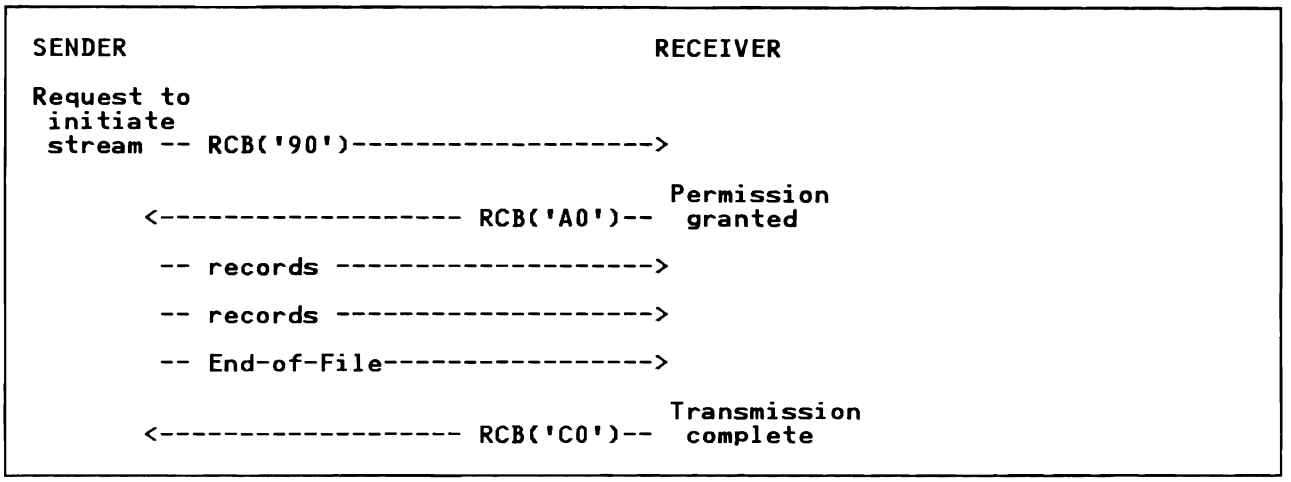

Figure 5-9. Normal Stream Operation

Note: It is also valid to receive an RCB X'B0' in response to an End-of-File when the receiver wishes to abort the file.

End-of-File is discussed in section 7.1, "Non-SNA String Control Byte (SCB) Function" on page 7-1. When transmission complete is received, the sender may purge all copies of the job on his own system. He may not do so before this point. At this point the stream is logically closed.

5.4.4 Abnormal Stream Termination

Either the sender or the receiver may terminate the transmission of a job before the End-of-File or the transmission complete RCB is sent. Such termination is called abnormal termination and may be caused by any problem either the sender or receiver has with a particular job *after* the initiating protocol has proceeded normally (i.e. the permission granted (RCB X'A0') has been sent).

The abnormal termination protocol for a sender initiated termination is as follows.

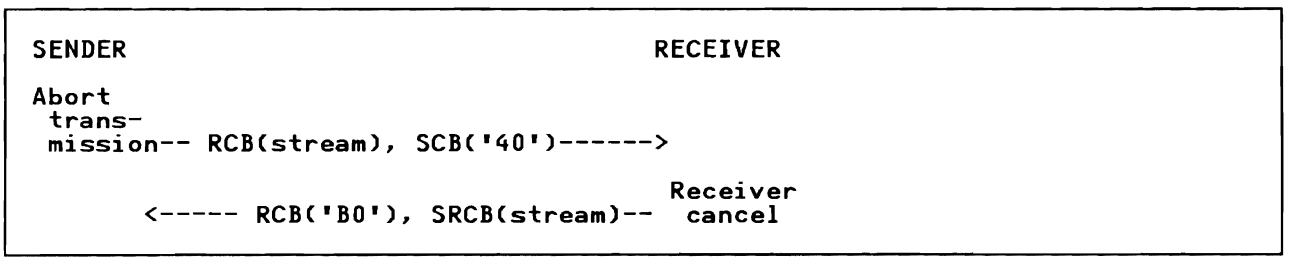

Figure 5-10. Sender Initiated Stream Termination

Note: The SCB is explained in 7.1, "Non-SNA String Control Byte (SCB) Function" on page 7-1. In SNA a RIDSRCB of 40 is used to indicate transmitter cancel rather than an SCB of 40.

No additional data may be sent on this stream until the receiver cancel is received, and that in turn implies the stream has been closed. A new request for permission must be sent before another job is transmitted. If

the abort was caused by system problems as opposed to a user or operator cancel request, a copy of the aborted job should be kept in the sender's spool space for transmission at a subsequent time.

Abnormal termination protocol when the receiver wishes to stop a job in the middle of transmission is as follows.

```
SENDER                                  RECEIVER

                                         Receiver
       <----- RCB('B0'), SRCB(stream)--   cancel

Abort
 trans-
 mission-- RCB('40'), SRCB(stream)----->

----------------- OR ------------------------

SENDER                                  RECEIVER

                                         Receiver
       <----- RCB('B0'), SRCB(stream)--   cancel

        -- RCB (stream), End-of-file--->
```

Figure 5-11. Receiver Initiated Stream Termination (BSC only)

Note: In SNA a RIDSRCB of X'40' is used to indicate transmitter cancel rather than an SCB of X'40'.

When the sender receives the receiver cancel, the following steps must be performed.

1. Stop sending the job for the stream specified in the receiver cancel message.

2. Place the copy of the current job for that stream into a held state.

3. Send an abort transmission to the receiver to show that the transmission of the specified job has been stopped, or send an End-of-File if the job has been completed.

Only after the protocol described above has been followed may the sender attempt to transmit a new job in the aborted stream. A request for permission must begin the attempt to transmit the new job. It is suggested the job be placed into a hold state so that the same job will not be retransmitted. Under most circumstances, if the receiver aborted the job once, it will be aborted on a subsequent retransmission.

When a condition is detected which the receiving system can not handle, a receiver cancel will be sent.

5.5 Data Format

5.5.1 Unspanned Data Record Format

The format of a standard data record for NJE transmission is shown in Figure 5-12 and Figure 5-13. This follows RCB's of 98-F8 and 99-F9.

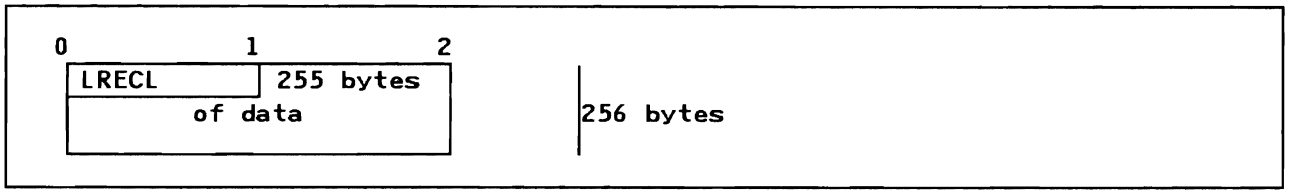

Figure 5-12. Data record without carriage control

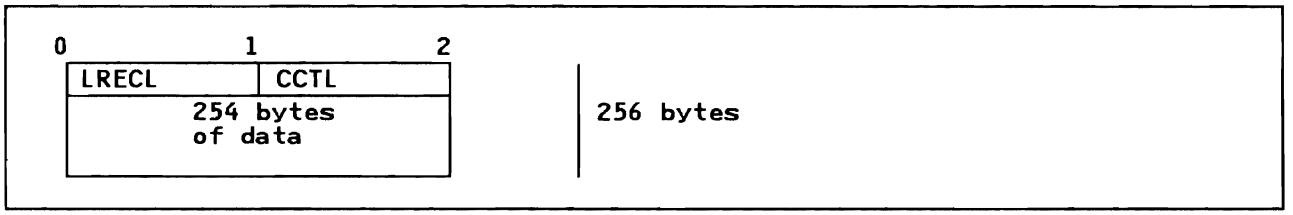

Figure 5-13. Data record with carriage control

The record length (LRECL) includes the carriage control character (CCTL) and data. It does not include the record length field. The maximum size record that can be transmitted without using the spanned record format is 255 bytes including the carriage control character.

If carriage control is machine then any CCW opcode may be sent in that record. S&F systems must forward all CCW opcodes intact. The destination may discard any records that it can not process. See "Print File Processing" on page B-31 for RSCS specific processing.

For a SYSIN stream, the default attribute for each record is fixed format record length 80. If the record length or record format is different, a dataset header record characteristics change section must appear. See section 8.5.2, "Record Characteristics Change Section" on page 8-28 for more information on the dataset header record characteristics change section. SYSIN data records may not contain carriage control.

5.5.2 Spanned Data Record Format

Spanned record support allows records longer than 255 bytes to be transmitted. For transmission, the records are broken into data segments of less than 256 bytes. The maximum size of a spanned record is 32,760 bytes. All segments contain the segment length (SEGL) at the beginning. The total length of the logical record (LRECL) is transmitted following the SEGL in the first segment. (All other segments contain only the SEGL and data.)

The byte used by the SEGL is not included in the SEGL value. The two bytes used by the LRECL is also not included in the SEGL. A sample spanned record transmission is shown in Figure 5-14. A sample spanned record with carriage control is shown in Figure 5-15.

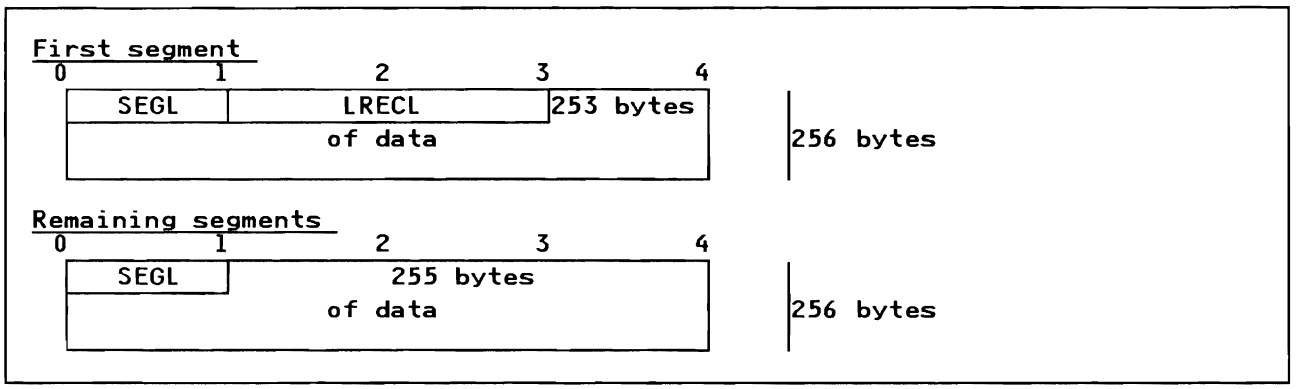

Figure 5-14. Spanned data record

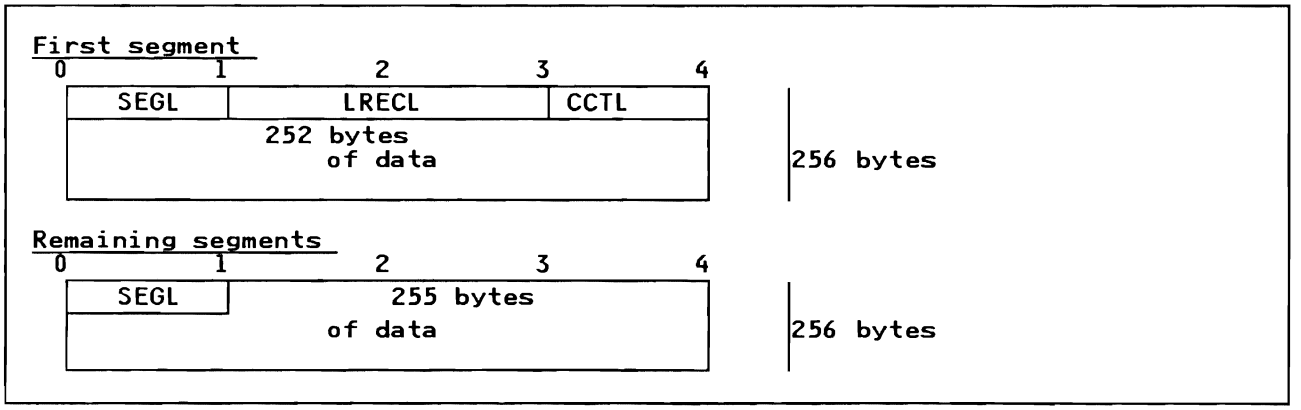

Figure 5-15. Spanned data record with carriage control

Each segment after the first can contain 255 bytes of data. The first may contain at most 253 bytes of data (including the carriage control).

If carriage control is machine then any CCW opcode may be sent in that record. S&F systems must forward all CCW opcodes intact. The destination may discard any records that it can not process. See "Print File Processing" on page B-31 and "Spanned Records" on page B-30 for RSCS specific processing. See "Spanned Records with Carriage Control" on page B-7 for JES2 specific processing. See "Spanned Record Support" on page B-17 for JES3 specific processing.

5.5.3 Trailing Blank Truncation

NJE allows the truncation of trailing blanks prior to transmission. The original LRECL is used to reconstruct the record. Any segmentation done for the purposes of controlling transmission buffers (or RUs) is done after blank truncation of the logical record.

As later sections of the document will show, compression and compaction may also be applied for transmission, but these also have mechanisms independent of the original LRECL described here, that are used for reconstruction of the (possibly truncated) record.

Further efficiency may be obtained by truncating blanks at the end of individual segments of spanned records prior to compression or compaction. In this case, the original SEGL is used to restore the segment. The concatenated segments are then used to build the original record. If required, the reconstructed record is padded with blanks to match the original 2-byte LRECL.

Chapter 6. Additional Multi-Leaving (ML) Functions - BSC, CTCA

Multi-leaving provides for the time division multiplexing of independent data streams on a single link. This section will discuss the additional control information available for initiating, terminating, suspending, and resuming streams, and also for block sequence checking in NJE (using BSC or CTCA).

6.1 Block Control Byte (BCB) Function

```
Binary     Meaning

 1... .... Must be 1
 1xxx .... Control information as follows:
             1000 cccc - Normal block
             1001 .... - Bypass sequence count validation
                         (sometimes called "BCB ignore bit")
             1010 cccc - Reset expected block sequence
                         count to 'cccc'
             1011 .... - Reserved for future use
             11xx .... - Reserved for future use
 1... cccc Modulo 16 sequence count
```

Figure 6-1. BCB Bit Definition

The Block Control Byte (BCB) is used to detect sequence errors in transmission and possible loss of data. After line initialization is complete (SOH ENQ and DLE ACK0 are exchanged), each end of the communication initializes an output BCB counter to 0000. The output BCB count is sent in each transmission block and after acknowledgment the count is incremented by one. The count is maintained modulo sixteen. Normally it is necessary to keep both an input and output BCB. The receiving end checks to see if the BCB it receives (input) is one more than the last BCB it received. If so, data transmission is normal and no data has been lost.

If the BCB received is not what is expected, the receiver must indicate something is wrong by sending an RCB indicating BCB sequence error to the other side (see section 5.2, "Record Control Byte (RCB) Function" on page 5-4 for details). The only exception, to allow error recovery, is when the BCB received is the same as the one received in the immediately previous transmission (i.e. a duplicate BCB): the receiving system assumes that the last transmission block it sent has been lost and its last block must be re-sent rather than indicating a BCB sequence error. The duplicate block it received is discarded.

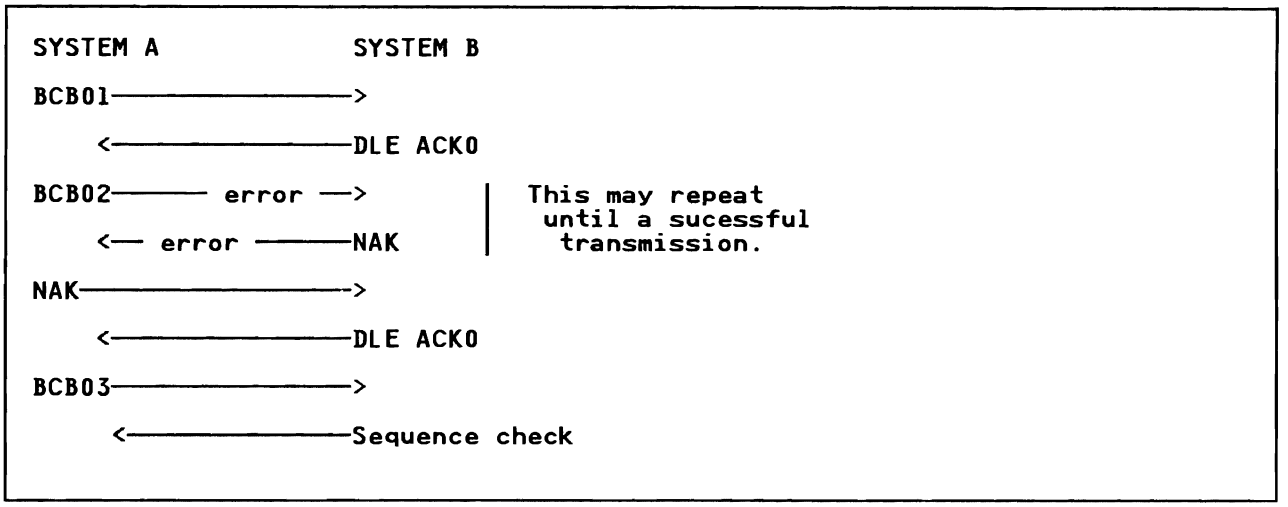

Figure 6-2. Results of not transmitting null records.

Generally when BCB sequence error is received, the side which receives it must terminate the connection because no error recovery is possible. Due to the importance of the BCB in detecting lost blocks, a null block containing a BCB should always be transmitted as an acknowledgment rather than using a DLE ACK0 as the acknowledgment. Figure 6-2 and Figure 6-3 will help to clarify this point. Currently, only POWER transmits Null Buffers rather than DLE ACK0.

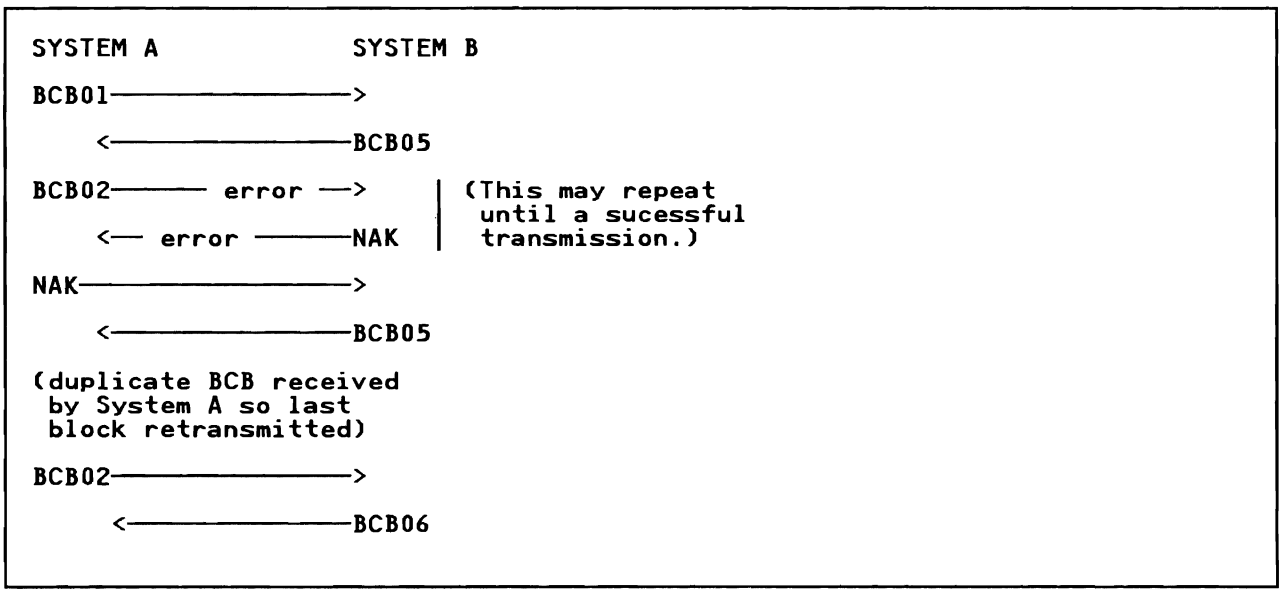

Figure 6-3. Correct recovery with null records.

Figure 6-4 is a state table describing BCB handling.

Current State	Event	Action	New State
S1	E1		S2
	E2	A1	S1
	E3	A3	S1
	E4	A6	S4
S2	E5	A1	S1
	E6	A3	S1
	E7	A5	S1
S3	E1		S2
	E2	A1	S1
	E3	A3	S1
	E4	A6	S4
S4	E8	A2	S3

State	Description
S1	Wait for response
S2	Check BCB
S3	Error recovery wait for response
S4	Check error count

Event	Description
E1	Receive data block
E2	Receive DLE ACK0
E3	Receive NAK
E4	Receive error (not CE-DE status)
E5	Correct BCB
E6	Duplicate BCB
E7	Incorrect BCB
E8	Error count not exceeded

Action	Description
A1	Transmit last non-NAK block of data (Zero error count)
A2	Transmit NAK
A3	Retransmit last non-NAK
A4	Transmit DLE ACK0
A5	Transmit sequence error
A6	Increment error count

Figure 6-4. BCB Handling State Table

See B.2.8.6, "BCB Handling" on page B-21 for JES3 deviations.

6.2 Function Control Sequence (FCS) Function

```
    Binary          Meaning

1... .... 1... .... Must be 1... .... 1... ....
.0.. .... .... .... Normal state
.1.. .... .... .... Suspend all stream transmission (Wait-a-bit)
..rr .... ..rr .... Reserved for future use
.... .... .1.. .... Remote console stream identifier

.... 1... .... .... Function stream identifier for:
                     RJE input stream number 1
                     RJE print stream number 1
                     NJE job transmission stream number 1

.... .1.. .... .... Function stream identifier for:
                     RJE input stream number 2
                     RJE print stream number 2
                     RJE punch stream number 7
                     NJE job transmission stream number 2
                     NJE SYSOUT transmission stream number 7

.... ..1. .... .... Function stream identifier for:
                     RJE input stream number 3
                     RJE print stream number 3
                     RJE punch stream number 6
                     NJE job transmission stream number 3
                     NJE SYSOUT transmission stream number 6

.... ...1 .... .... Function stream identifier for:
                     RJE input stream number 4
                     RJE print stream number 4
                     RJE punch stream number 5
                     NJE job transmission stream number 4
                     NJE SYSOUT transmission stream number 5

.... .... .... 1... Function stream identifier for:
                     RJE input stream number 5
                     RJE print stream number 5
                     RJE punch stream number 4
                     NJE job transmission stream number 5
                     NJE SYSOUT transmission stream number 4


.... .... .... .1.. Function stream identifier for:
                     RJE input stream number 6
                     RJE print stream number 6
                     RJE punch stream number 3
                     NJE job transmission stream number 6
                     NJE SYSOUT transmission stream number 3

.... .... .... ..1. Function stream identifier for:
                     RJE input stream number 7
                     RJE print stream number 7
                     RJE punch stream number 2
                     NJE job transmission stream number 7
                     NJE SYSOUT transmission stream number 2

.... .... .... ...1 Function stream identifier for:
                     RJE punch stream number 1
                     NJE SYSOUT transmission stream number 1
```

6.2.1 Multi-Leaving Stream Control

NJE allows up to 7 SYSIN transmitters, 7 SYSOUT transmitters, 7 SYSIN receivers and 7 SYSOUT receivers per network line. The total number of SYSIN and SYSOUT streams per line in either direction cannot exceed 8. This means that the sum of the SYSIN receivers and SYSOUT receivers on a line is less than or equal to 8 (similar logic applies for transmitters).

Furthermore, only certain combinations of stream numbers are allowed since the FCS bit encodings map to different stream numbers for SYSOUT and job transmission (SYSIN). Thus, no two streams that use the same FCS bit may be active at the same time.

| *Note:* RSCS Networking Release 3 and JES3 Networking support only one stream for input job transmission and one stream for output job transmission. (Stream number 1 is supported in each case).

Individual streams may be suspended by setting to zero the appropriate bit for that stream in the FCS.

| *Note:* RSCS Networking Release 3 will ignore that bit and send data for that stream anyway. JES3 always sends an FCS with all streams enabled.

Whenever a stream is suspended, no data records for that stream should be transmitted. If any data is received for a stream that was suspended when the last FCS was sent, the receiving system (Wait-a-bit sender) has the option either of accepting the data anyway or of aborting that stream by sending a reject function (RCB of X'B0') for that stream.

If data is received when Wait-a-bit All is on, the receiver may take either of two actions: either a NAK is sent because a lost data condition was detected or an SOH ENQ is sent to restart all activity on that line. The lost data condition could be caused by the data buffer not being large enough to receive the block that was sent. The data buffer for the Wait-a-bit response should be large enough to receive a complete Null Buffer (DLE,STX,BCB,FCS,FCS,RCB(0),ETB). A minimum of ten (10) bytes is recommended.

The system receiving a Wait-a-bit should delay responding for a time in the range of one to two seconds if the Wait-a-bit was sent with a null record. A delay of greater than two seconds is not recommended as a timeout will occur if the response is not received within three seconds. If the Wait-a-bit was sent with data, the system should respond immediately. The immediate response allows the system requesting the delay to continue to transmit data as fast as possible. The approach of always delaying causes transmission delays whenever a system has data to transmit, but can not receive data.

The response may be either a DLE ACK0 or a Null Buffer. It is not permissible to transmit a data record in response to a Wait-a-bit. If a DLE ACK0 is sent, the system receiving the DLE ACK0 will respond based on the last FCS that was sent as it always would upon receipt of a DLE ACK0. It is recommended that Null Buffers be sent rather than DLE ACK0. Currently, only POWER transmits Null Buffers rather than DLE ACK0. (See 6.1, "Block Control Byte (BCB) Function" on page 6-1 for the reasons behind this recommendation.)

Chapter 7. Data Compression and Compaction

Compression is a method of reducing the length of records for transmission by removing blanks and duplicate characters. Compaction is a method of reducing the length of records for transmission by representing certain 8 bit characters with only 4 bits.

For BSC and CTCA NJE transmission, data compression is always used; compaction is not used.

For SNA data transmission, compression is also required, and the use of compaction is optional and is negotiated via NJE and SNA protocols. See A.4.2, "Session Initialization" on page A-25 for details. SCBs are always used, whether or not the data is actually compressed; i.e., even non-compressed or compacted data must be interspersed with SCBs having the two high order bits set (see Figure 7-1).

A bit-encoded String Control Byte (SCB) is used to specify compression and compaction parameters, but the bit encodings are different for SNA than for Non-SNA transmission.

Signon records are not compressed or compacted and are preceded with an SCB indicating no compression.

7.1 Non-SNA String Control Byte (SCB) Function

The String Control Byte is used to compress multileaving data. Each record is compressed before being placed into the transmission buffer. The format of the SCB is defined in Figure 7-1.

```
  Binary    Meaning

0000 0000   End-of-record
            If first SCB, this also indicates end-of-file
0100 0000   Abort transmission
100b bbbb   'bbbb' blanks are to be inserted
101d dddd   The single character following this SCB is to
            be duplicated 'ddddd' times
11cc cccc   The 'cccccc' characters following this SCB are
            to be inserted (maximum 63)
```

Figure 7-1. SCB Definition

End of File Indication Using BSC or CTCA: The End-of-File is a zero length data record. It may be in a transmission block by itself.

The transmitted EOF stream is as follows:

```
DLE,STX,BCB,FCS,FCS,RCB(stream),SCB(0),DLE,ETB  - BSC

DLE,STX,BCB,FCS,FCS,RCB(stream),SCB(0)          - CTCA

DLE,STX,BCB,FCS,FCS,RCB(stream),SCB(0),DLE,ETB  - CTCA/PREPARE
```

Note: When receiving these streams, remember that the BSC hardware strips off the trailing DLE but passes the ETB.

7.2 SNA String Control Byte (SCB) Function

The SNA SCB defines the beginning and end of compacted and compressed data. An SCB must begin each RU following the exchange of FM headers (see A.4.2.2, "Function Management Headers" on page A-29 for more information). The one-byte SCB consists of a 2 bit description field, followed by a 6 bit count field. This count is the number of characters that are described by this SCB. Therefore, one SCB may replace up to 63 characters, or identify up to 63 intervening uncompressed characters before the next SCB. In all cases, a count of 0 is a reserved value.

The SCB format is shown in Figure 7-2.

Description	Count	Usage
00	cccccc	No compressed characters follow
10	cccccc	Repeat blanks
11	cccccc	Repeat non-blank character which follows
01	cccccc	Compacted characters follow

Figure 7-2. String Control Byte (SCB) Format for SNA

Compression: Compression may be optionally indicated in the BIND for LUTYPE 0, however, compression is required for NJE using SNA. Two or more blanks and three or more non-blanks are compressed. For example, if 5 blanks were being compressed, the SCB would be B'10000101'. If 5 'A's were being compressed, the SCB would be B'11000101' followed by the character 'A' (or X'C1').

Compaction: Compaction allows certain 8-bit character sequences to be represented in network transmission as 4 bits. The characters which are compacted are called master characters. When two master character are adjacent in a data stream, they are compacted from their normal 8-bit representation into 4 bits. Non-master characters may also be defined; these characters are not compacted, but when they are adjacent to master characters, they will not interrupt the compaction SCB which is in effect. All other characters are considered non-compactable and will be transmitted in their true 8-bit form. Obviously the non-compactable characters should be those which are least frequently used in the data stream.

Master and non-master characters are transmitted via FMH3 at session initialization. See "Compaction Table Format" on page A-32 for more information.

NJE End of File Indication Using SNA: Note that EOF in NJE using SNA does not use the SCB. Instead, EOF is indicated by the sequence RCB(stream),SRCB(0). The EOF record may appear alone in the transmission buffer.

Chapter 8. Control Formats

8.1 Control Formats Representation

All fields in control formats will be represented in the following form:

SYMBOL disp len type default range

Text describing purpose of this field.

JES2 JES2 specific comments.

JES3 JES3 specific comments.

RSCS RSCS specific comments.

POWER POWER specific comments.

disp The hex displacement of this field within the control format. This displacement includes the section type, subtype, and length fields.

len The length of this field in decimal. If the length is preceded by a '.' then it is the number of bits rather than the number of bytes.

type This will usually be binary or character.

default This is the default value for the field. Usually this will be zero or blanks.

range This is the range of values that may be used within this field. If not specified, any bit value is valid.

Any field that does not have system specific text is set and used as described by the general descriptive text.

8.2 Control Record Header

All Control Records begin with a four byte header. It is defined as follows:

NXXLEN 0 2 Binary 8-32K

This is the length of the control record. It includes all the data in each section and the four bytes for the header. Some implementations limit this value to 4096 bytes or less. See

Appendix B, "System Dependent Considerations and Comments" on page B-1 for specific system limitations.

NXXFLAGS 2 1 Binary 0

This is an unused flag byte.

NXXSEQ 3 1 Binary 0

This is an unused sequence indicator until the control record is segmented for transmission. At that time, it contains the position of this segment within the record and the high order bit is a flag to indicate that more segments follow.

8.3 Control Record Section Header

All sections within a control record begin with a four byte header. This four byte header contains the length of the section and the section identifier and section modifier. Within each section the identifier and modifier field for that section will be listed.

NXXXLEN 0 2 Binary 4-32764

This is the length of the control record section. It includes all the data in the section and the four bytes for the header. See Appendix B, "System Dependent Considerations and Comments" on page B-1 for specific system limitations.

NXXXTYPE 2 1 Binary

This byte defines the major type of this section. All general sections are represented by a byte of 0. Each operating system has a defined byte for its section.

NXXXMOD 3 1 Binary 0

This byte modifies the section identifier. It allows each section to have up to 256 different formats. For example, the general section of the dataset header has an extension for 3800 data that is defined by a modifier of X'80'.

8.4 Job Header

For an overview of the Network Job Header, please see 3.4, "Job Header" on page 3-8.

8.4.1 General Section

This section is identified by an identifier field of X'00' and a modifier field of X'00'. All fields of the job header should remain unchanged for the life of the job unless specifically stated as modified.

NJHGJID 4 2 Binary none 1-65535

This is a number associated with the JOB assigned at the origin node. It remains with the job throughout its life on the network.

JES2 Jobs, started tasks, time sharing users are in the range 1 to 9999. Used for SMF 26 records. This number may be used for displaying status when the job resides on a JES2 node.

JES3 This field is set from the job number. It is in the range 1 to 9999. It is used for SMF 26 records.

RSCS This field is set from the origin spool file identifier. This field is kept as the origin spool file identifier on received jobs. The value is in the range of from 1 to 9900.

POWER This field is set from the Job Number if the job was entered by POWER, otherwise it is kept as the origin job number in various account records for all received jobs. It may be from 1 to 32767.

NJHGJCLS 6 1 Character A A-Z,0-9

This is the default class associated with the job, if not specified elsewhere.

JES2 Set from Jobcard, Jobparm, JECL, and input device class.

JES3 Set to "A" for jobs originating at JES3. Otherwise unused.

RSCS This field is not used by RSCS. It is set to "A" for jobs originating at RSCS.

POWER This is used as the Job Class.

NJHGMCLS 7 1 Character A A-Z,0-9

This is the message class associated with the job.

JES2 Set from Job Card.

JES3 Set to "A" for jobs originating at JES3. Otherwise unused.

RSCS This field is not used by RSCS. It is set to "A" for jobs originating at RSCS.

POWER This is not used by POWER. It is set to "A" for jobs originating at POWER.

NJHGFLG1 8 1 Bits 0

This flag byte defines the following bits.

NJHGF1PR 8 .1 80 bit 0

This is the recompute selection priority flag.

JES2 For Job Transmitter: Turned on if user specifies priority on either /*PRIORITY JECL or Job Card.

For Job Receiver: If on, JES2 uses the priority in NJHGPRIO as the job's execution priority. If off, NJHGPRIO value is not used and the job is given the installation's default priority.

JES3 Defaulted, not used.

RSCS Defaulted, not used. (RSCS always uses the value in NJHGPRIO.)

POWER Defaulted, not used. (POWER always uses the value in NJHGPRIO.)

NJHGPRIO 9 1 Binary 0 0-15

This is the selection/transmission job priority. 0 is lowest, and 15 is highest.

JES2 It is used as selection priority. Job receiver: If NJHGF1PR is on, JES2 will use as execution priority. If off, JES2 will ignore this field.

Job transmitter: JES2 will set this field as execution priority.

JES3 Defaulted, not used. Transmits FIFO.

RSCS RSCS translates its priorities 99-0 to 0-15 on transmission and translates them back again on reception. Store and Forward jobs are never altered, even if their priority is changed while on VM/370 systems. (RSCS ignores the setting of NJHGF1PR.)

- 0 to 99; 90-99 to 0
- 1 to 92; 84-89 to 1
- 2 to 85; 78-83 to 2
- 3 to 78; 72-77 to 3
- 4 to 71; 66-71 to 4
- 5 to 64; 60-65 to 5
- 6 to 57; 54-59 to 6
- 7 to 50; 48-53 to 7
- 8 to 44; 42-47 to 8
- 9 to 37; 36-41 to 9
- 10 to 31; 30-35 to 10
- 11 to 27; 24-29 to 11
- 12 to 19; 18-23 to 12
- 13 to 12; 12-17 to 13
- 14 to 6; 6-11 to 14
- 15 to 0; 0-5 to 15

POWER This is translated to the POWER priority (0-9) on received jobs and translated from the POWER priority on transmitted jobs. It is never changed for Store and Forward jobs, even if the operator alters the value on the POWER node. (POWER ignores the setting of NJHGF1PR.)

- 0 to 0; 0 to 0
- 1 to 1; 1 to 1
- 2 to 2; 2 to 3
- 3 to 2; 2 to 3
- 4 to 3; 3 to 5
- 5 to 3; 3 to 5
- 6 to 4; 4 to 7
- 7 to 4; 4 to 7
- 8 to 5; 5 to 8
- 9 to 5; 5 to 8
- 10 to 6; 6 to 10
- 11 to 7; 7 to 12
- 12 to 7; 7 to 12
- 13 to 8; 8 to 13
- 14 to 8; 8 to 13
- 15 to 9; 9 to 15

NJHGORGQ A 1 Binary 0

This is the origin system qualifier in a loosely coupled multi-processor complex. This is used as the qualifier in any nodal message records relating to this job.

JES2 This is the system ID for MAS.

JES3 This is used in a JES3 complex as the index to a JES3 local processor for TSO submitted jobs. For jobs originating at JES3, if the job is from a TSO user, the proper index value is set, which will never be X'00'. A zero qualifier indicates the job is not from TSO. On input, the field is saved for generating TSO notify messages.

RSCS Not used, Set to X'01'.

POWER Set from the POWER SYSID field. This may be X'40' or X'F1' through X'F9'. Defines the shared spooling system when running in a shared spooling complex.

NJHGJCPY B 1 Binary 1

This is the number of output copies for the job as indicated in the pre-execution SYSIN stream.

JES2 Set from the /*JOBPARM COPIES= parameter. Used on the SYSOUT destination node: multiplied by the data set header copy count for each SYSOUT data set in the job.

JES3 Not used. Not set for SYSIN jobs. For SYSOUT, set to X'01'.

RSCS Defaulted, used on SYSIN jobs to set the spool copy count. Not used for SYSOUT.

POWER Set to zero, not used.

NJHGLNCT C 1 Binary 0

This is the default lines per page for SYSOUT files. A value of X'00' or X'FF' causes the system not to count lines (or use a local default value, depending on the system).

All other values are treated as an explicit number of lines on a page before a page eject is generated by the printing subsystem.

JES2 Set from the /*JOBPARM or JOB statement.

X'00' (the default): Use the default value of the 'printing' (i.e., destination) subsystem.

X'FF' Do not force any page ejects, but let the skipping be solely determined by the carriage control (if present) in the SYSOUT data.

Used by the SYSOUT receiver to set JCTLINCT as follows:

1. If NJHGLNCT is X'FF', JCTLINCT is set to X'00' - no system-generated skipping.
2. If NJHGLNCT is X'00', JES2 used the local system default - LINECT.

Used for actual output processing only if NDHGLNCT is zero.

JES3 Defaulted, not used.

RSCS Set to X'FF' for print and defaulted for punch. Not used.

POWER Defaulted. For received SYSOUT it is used as the default lines per page when the output does not contain carriage control and NDHGF1OV is *not* set. For values of X'00' and X'FF', POWER does not insert skips. (Skips in the SYSOUT data set are honored.)

RESERVED D 3 Binary 0

NJHGACCT 10 8 Character blanks

This field contains the network accounting information for the job.

JES2 For transmission, JES2 sets this field from the /*NETACCT card or converts the local account number using the NETACCT translation tables. Upon receiving this field, the job would be assigned a local account number if the network account number is found in the NETACCT translation tables. The default is set to zero. Blanks are accepted as a default if received.

JES3 Set from the //*NETACCT control statement. The default is blanks.

RSCS Defaulted, not used.

POWER Defaulted, not used.

NJHGJNAM 18 8 Character

This is the name of the job. It remains intact throughout the life of the job. It may be blanks.

Note: The execution node may use the value on the job card or a local default instead of this field, for the purposes of local processing.

JES2 Used as the jobname. Blanks will be accepted if received. JES2 may send a blank jobname if a JOB card with a blank name field is followed by a /*XMIT card.

JES3 Used as the jobname. Defaulted to "NJEJOB" if not specified at the origin node.

RSCS Set to "RSCS" followed by the 4-digit spool file number. Used for file name of received files.

POWER Used as the jobname. Initialized to the POWER job name. If null on arrival it's set to 'JOB' + NJHGJID. The value within POWER must never be null.

NJHGUSID 20 8 Character blanks

This field is to contain the userid of the person to notify at the NJHGORGN node.

JES2 Set from the userid specified in the NOTIFY parameter of the Job statement or the userid specified in the /*NOTIFY JECL statement.

JES3 This is set if submitted by a TSO user or if specified on the //*NETACCT control statement. It is the origin userid unless specified on //*NETACCT. In this case, it is the notify id. In either case, this is the userid to which notification should be sent. The default is blanks if no userid applies.

RSCS It is used for the origin userid and set from the origin userid.

POWER It is used for the Notify Userid. It is set from the Notify Userid. If specified, then a Notify message is sent when the SYSIN or SYSOUT is transmitted, and if the SYSIN executes on POWER, then a notification is sent at execution completion time. If SYSOUT is received which is destined for the receiving node, then the userid is notified.

NJHGPASS 28 8 Character 0

This is the password for the job. It is used to validate that the user is allowed to run the job.

JES2 Used instead of transmitted Job card, if present. Set from the JOB card unless sent with /*XMIT statement. Blanked out for SYSOUT transmission and not used if received for SYSOUT.

JES3 Used, not set.

RSCS Defaulted, not used.

POWER It is used as the job password and set from the job password. It is used for job protection.

NJHGNPAS 30 8 Character 0

This is the new password. If the old password is valid and this is specified, the password will be changed to this value.

JES2 Set from the JOB card unless sent with /*XMIT statement. Used instead of transmitted Job card, if present.

Set to blanks for SYSOUT transmission. Not used if received for SYSOUT.

JES3 Used, not set.

RSCS Defaulted, not used.

POWER Set to blanks, not used.

NJHGETS 38 8 Binary 0

This is the job entry date/time stamp on the origin node in 370 store clock format. It should reflect GMT.

JES2 Set and used for SMF.

JES3 Set with store clock on job submission. Used for SMF if not zero for received jobs.

RSCS Set to GMT from spool file, used as origin time.

POWER Set but not used.

NJHGORGN 40 8 Character origin node name

This may be the origin node name. It is used as the node for job notification.

JES2 Set from the origin node name. It may be overridden by the /*NOTIFY JECL statement. For SYSOUT reception, only used for SMF.

JES3 Set to the origin node. May be overridden for SYSOUT from locally executed jobs (i.e. the job contained a //*MAIN statement to change the apparent origin). Used for SMF and also for notify messages.

RSCS Set only to the origin node. Used as the origin node id.

POWER Set from the origin node name. It will be overridden if NTFY is specified. Used as the origin node id and also for notify node name.

NJHGORGR 48 8 Character blanks

This is the origin userid or remote name (workstation name).

JES2 Set as the origin remote name.

JES3 Set and used for SYSIN jobs. Used but not set for SYSOUT. The default is zeros.

RSCS This is used as the origin userid if NJHGUSID is not specified. It is set to the origin userid.

POWER Set and used as origin userid or remote name. The format for a remote name is 'Rnnn'.

NJHGXEQN 50 8 Character blanks

This is the execution node name. This will be changed if the operator reroutes a SYSIN job.

JES2 Set to the execution node for jobs using /*ROUTE XEQ, /*XEQ or /*XMIT JECL. For SYSOUT reception, only used for SMF.

JES3 Set to the execution node for SYSIN jobs. For locally submitted and executed jobs, set to the origin for SYSOUT. JES3 has no SYSIN rerouting capability.

RSCS This is used for the destination node for SYSIN files. It is set to the destination node for SYSIN files, and the origin node for SYSOUT files.

POWER Set to the execution node for SYSIN.

NJHGXEQU 58 8 Character blanks

This is the execution userid for VM/370 only.

JES2 Set to the userid from the /*XEQ and /*XMIT JECL.

JES3 Set to the execution userid for SYSIN. Not set for SYSOUT. Defaults to zeros. Set by the //*ROUTE XEQ statement.

RSCS This is used for the destination userid for SYSIN files. It is set to the destination userid for SYSIN files and to the origin userid for SYSOUT files.

POWER Set to the execution userid for SYSIN. Not used.

NJHGPRTN 60 8 Character blanks

This is the default print destination node for print type files that are not specifically routed elsewhere at execution time.

JES2 Set by /*ROUTE PRINT card to default print destination node, or to the PRNODE parameter on the reader through which the job was read. Defaults to the origin node. Used at the execution node for default sysout routing and propagated to NDHGNODE field in the data set header(s).

JES3 Set to the origin node for SYSIN jobs. Not set for SYSOUT from locally submitted and executed jobs. For incoming SYSIN jobs, if the default print node is not null, and is not equal to the origin node, a JES3 //*FORMAT statement may be inserted by the user into the JCL to cause the default for all print output to be the default print node.

RSCS Set to the origin node for SYSIN and the destination node for SYSOUT.

POWER Set to the origin node or to the value specified in the LDEST parameter.

NJHGPRTR 68 8 Character blanks

This is the default print output userid or workstation.

JES2 Set by /*ROUTE PRINT card to default print output userid or remote workstation, or to the PRDEST parameter on the reader through which the job was read. Used at the execution node for default sysout routing and propagated to NDHGRMT field in the data set header(s).

JES3 Set to the remote origin id for SYSIN jobs. Default is zeros. Not set for SYSOUT from locally submitted and executed jobs. For incoming SYSIN jobs, if the default print node and remote are not null, and the default print remote is not equal to the origin remote, a JES3 //*FORMAT statement may be inserted by the user into the JCL to cause the default for all print output to be the default print node and remote.

RSCS For SYSOUT, set to destination userid or remote workstation. If destination is 'SYSTEM', set to X'00'. For SYSIN, set to origin userid or X'00' if file originated from a workstation.

POWER Set from the origin userid or remote or to the value specified in the LDEST parameter.

NJHGPUNN 70 8 Character blanks

This is the default punch destination node for punch type files that are not specifically routed elsewhere.

JES2 Set by /*ROUTE PUNCH card to default punch destination node, or to the PUNODE parameter on the reader through which the job was read. Defaults to the origin node. Used at the execution node for default sysout routing and propagated to NDHGNODE field in the data set header(s).

JES3 Set to the origin node for SYSIN jobs. Not set for SYSOUT from locally submitted and executed jobs. For incoming SYSIN jobs, if the default punch node is not null, and is not equal to the origin node, a JES3 //*FORMAT statement may be inserted by the user into the JCL to cause the default for all punch output to be the default punch node.

RSCS Set to the origin node for SYSIN and the destination node for SYSOUT.

POWER Set to the origin node for jobs or to the value specified in the PDEST parameter.

NJHGPUNR 78 8 Character blanks

This is the default punch output userid or workstation.

JES2 Set by /*ROUTE PUNCH card to default punch output userid or remote workstation, or to the PUDEST parameter on the reader through which the job was read. Used at the execution node for default sysout routing and propagated to NDHGRMT field in the data set header(s).

JES3 Set to the remote origin id for SYSIN jobs. Default is zeros. Not set for SYSOUT from locally submitted and executed jobs. For incoming SYSIN jobs, if the default punch node and remote are not null, and the default punch remote is not equal to the origin remote, a JES3 //*FORMAT statement may be inserted by the user into the JCL to cause the default for all punch output to be the default punch node and remote.

RSCS For SYSOUT, set to destination userid or remote workstation. If destination is 'SYSTEM', set to X'00'. For SYSIN, set to origin userid or X'00' if file originated from a workstation.

POWER Set from the origin userid or remote id or to the value specified in the PDEST parameter.

NJHGFORM 80 8 Character blanks

This is the default job form for any output created by the job at execution time. It should be copied into the dataset header at execution time.

JES2 Default forms are specified at the destination node by the JES2 parameter &STDFORM. The default for &STDFORM is 'STD'.

JES3 Not used. Initialized to zeros.

RSCS Not used. Initialized to zeros.

POWER Not used.

NJHGICRD 88 4 Binary 0

This is the input card count.

JES2 This is the input card count as counted by the reader at the input node. This is not the number of records transmitted. This field is used for SMF accounting.

JES3 Defaulted, not used.

RSCS Release 3 - Not used, set to TAGRECNM for SYSIN only.
Version 2.1 - Used as the record count for SYSIN. Set to TAGRECNM for SYSIN only.

POWER Not used, set to the transmitted record count for the input job.

NJHGETIM 8C 4 Binary 0

This is the estimated job execution time limit.

JES2 Set from the JOB card or /*JOBPARM card. Used for SMF.

JES3 Not set, used for SMF 26.

RSCS Defaulted, not used.

POWER Defaulted, not used.

NJHGELIN 90 4 Binary 0

This is the estimated output print lines.

JES2 Set from the JOB card or /*JOBPARM card. Used for SMF.

JES3 Not set for jobs originating at JES3 nodes. Set after execution to the actual line count. Used for SMF 26.

RSCS Not used or set.

POWER Defaulted, not used.

NJHGECRD 94 4 Binary 0

This is the estimated punched card output.

JES2 Set from the JOB card or /*JOBPARM card. Used for SMF.

JES3 Not set for jobs originating at JES3 nodes. Set after execution to the actual card count. Used for SMF 26.

RSCS Defaulted, not used.

POWER Defaulted, not used.

NJHGPRGN 98 20 Character blanks

This is the programmer's name.

JES2 Set from the JOB card. Used for SMF.

JES3 Set from the //*NETACCT control statement. Default is blanks. Used for SMF.

RSCS Release 3 - Used in the message acknowledging receipt of the file. Set to the origin userid for SYSOUT only. Set to X'00' for SYSIN.
Version 2.1 - Not used. Set as in Release 3.

POWER Defaulted; used for separator pages.

NJHGROOM AC 8 Character blanks

Programmer's room number.

JES2 First four characters set from the JOB card or /*JOBPARM card. Used for SMF.

JES3 Set from the //*NETACCT control statement. Default is blanks. Used for SMF.

RSCS Release 3 - Set from VM/370 distribution code. Not used.
Version 2.1 - Set from VM/370 distribution code. Used to set VM/370 distribution code.

POWER Defaulted; used for separator pages.

NJHGDEPT B4 8 Character blanks

Programmer's department number.

JES2 Defaulted, not used.

JES3 Set from the //*NETACCT control statement. Default is blanks. Used for SMF.

RSCS Defaulted, not used.

POWER Defaulted; used for separator pages.

NJHGBLDG BC 8 Character blanks

Programmer's building number.

JES2 Defaulted, not used.

JES3 Set from the //*NETACCT control statement. Default is blanks. Used for SMF.

RSCS Defaulted, not used.

POWER Defaulted; used for separator pages.

NJHGNREC C4 4 Binary 0

Record count on SYSOUT transmission.

JES2 Set by the SYSOUT transmitter as pure record count (not adjusted for multiple copies). This is the sum of record counts for each dataset after this job header and does not include the job headers nor the data set headers. Used by the SYSOUT receiver when issuing the JOB received message.

JES3 Defaulted, not used.

RSCS Not used. Set from TAGRECNM for SYSOUT only. Not set for SYSIN.

POWER Used as the line count of the output file; a spanned (extended) record counts as 1 regardless of the number of segments.

8.4.2 JES2 Section

This section is identified by an identifier field of X'84' and a modifier field of X'00'.

NJH2FLG1 4 1 Bits 0

Job level flags.

NJH2FJOB 4 .1 03 bit

Equate when job is not a batch job (i.e., either of the two following bits).

NJH2FSTC 4 .1 01 bit

Job is a started task.

NJH2FTSU 4 .1 02 bit

Job is a time-sharing user.

NJH2ACCT 8 4 Character 0

This is the originator's JES2 account number.

NJH2USID C 8 Character blanks

User SMF field - Defaulted, not used. This was added with APAR OZ74836 for the Spool Offload facility, and is not used for NJE transmissions.

The following four fields (-USR, -GRP, -SUSR and -SGRP) were added to JES2/1.3.4 with APAR OZ81051. They are not filled in for NJE transmissions, but only used for spool offload operations.

NJH2USR 14 8 Character binary zeroes

JCL User ID - Defaulted, not used.

NJH2GRP 1C 8 Character binary zeroes

JCL Group ID - Defaulted, not used.

NJH2SUSR 24 8 Character binary zeroes

Submitter's User ID - Defaulted, not used.

NJH2SGRP 2C 8 Character binary zeroes

Submitter's Group ID - Defaulted, not used.

8.4.3 POWER Section

This section is identified by an identifier field of X'86' and a modifier field of X'00'.

NJHPFLG1 4 1 Bits 0

This flag byte is not used.

NJHPDISP 5 1 Character D

This is the job disposition as specified in the * $$ JOB statement.

RESERVED 6 1 Binary 0

NJHPSYID 7 1 Character blank

Used as the target system qualifier for a shared spool environment

NJHPUSER 8 16 Character blanks

Contains user information copied from the * $$ JOB statement.

NJHPDSKT 18 2 Binary 0

Diskette address.

RESERVED 1A 2 Binary 0

Not used.

8.4.4 Job Scheduling Section of the Job Header

Fields are necessary to contain values for estimated pages and bytes transmitted. These fields are grouped together in this section with identifier X'8A' and modifier X'00'.

NJHEPAGE 4 4 Binary 0

This is the estimated 'begin page' structured fields for page-mode sysout data sets; i.e., the number of data records that begin with the sequence X'D3A8AF'.

JES2 Set from the /*JOBPARM PAGES= parameter, and used for SMF.

JES3 Set from the //*MAIN PAGES= parameter, and used for SMF. Set to actual pages after execution.

NJHEBYTE 8 4 Binary 0

This is the estimated number of output bytes. This includes all user SYSOUT data bytes.

JES2 Set from the /*JOBPARM BYTES= parameter, and used for SMF.

JES3 Set from the //*MAIN BYTES= parameter, and used for SMF. Set to actual bytes after execution.

8.4.5 User Section

This section is identified by an identifier field of B'11xxxxxx' and a modifier field of the users choice. This section is not used by any IBM products, but will be passed through the network.

Note: Beyond the four byte header, this definition is the responsibility of the installation. The section length is limited by the design to 32764. The combined length of all sections in the header is limited to 32764, although product implementations may be more restrictive. The length must be reflected in the NJHULEN field in the front of the control record section header. See 8.3, "Control Record Section Header" on page 8-2.

A sample four byte field is shown below only as an example:

NJHUCODE 4 4 Character blanks

This could be used for the SHARE/GUIDE/SEAS installation code, or some other unique identifier.

8.5 Dataset Header

For an overview of the Network Data Set Header, please see 3.5, "Data Set Header" on page 3-12.

8.5.1 General Section

This section is identified by an identifier field of X'00' and a modifier field of X'00'.

NDHGNODE 4 8 Character

This is the destination node name for this dataset. It is defaulted at job execution time from either NJHGPRTN or NJHGPUNN.

NDHGRMT C 8 Character blanks

This is either the destination userid or the destination remote workstation. It is defaulted at job execution time from either NJHGPRTR or NJHGPUNR.

JES2 (Used by TSO/E Interactive Data Transmission Facility for userid.)

JES3 (Used by TSO/E Interactive Data Transmission Facility for userid.)

RSCS Release 3 - Set to the destination userid unless the DEST keyword on the TAG is specified. In that case, the value specified for DEST is placed in this field. Used for destination userid if not defaulted.
Version 2.1 - Set to the destination userid, or X'00' when 'SYSTEM' is the destination. Used for destination userid if not defaulted.

POWER Route to RMT, userid or program.

NDHGPROC 14 8 Character blanks

This is the name of the JCL procedure that was being executed when this dataset was created.

JES2 Set in SP1.3.3. Defaults to zero in previous releases.

RSCS Set from FILENAME, not used.

POWER Defaulted, not used.

NDHGSTEP 1C 8 Character blanks

This is the step name that was executing when this dataset was created.

JES2 Set in SP1.3.3. Defaults to zero in previous releases.

RSCS Set from the FILETYPE. It is not used.

POWER Defaulted, not used.

NDHGDD 24 8 Character blanks

This is the DDNAME that was used to reference this dataset when the dataset was created.

JES2 Set in SP1.3.3. Defaults to zero in previous releases.

JES3 If blank on input, a default ddname of '*NJEWKnn' is used, where nn is an internal index value.

RSCS Release 3 - Set to VM/370 distribution code, not used.
Version 2.1 - Defaulted, not used.

POWER Defaulted, not used.

NDHGDSNO 2C 2 Binary 0

This is a counter that is incremented each time within the same job that a SYSOUT data set is allocated. It is to maintain uniqueness of datasets while still allowing them to be spun off.

JES2 Used and set from the PDBDSKEY field in JES2. The maximum value is 9999 until the fix for APAR OZ93770, after which the maximum is 32767.

JES3 Defaulted, not used.

RSCS Release 3 - Set to X'0001', not used.
Version 2.1 - Defaulted, not used.

POWER Defaulted, not used.

RESERVED 2E 1

NDHGCLAS 2F 1 Character A A-Z,0-9

This is the SYSOUT class.

JES2 If the class is not alpha-numeric, JES2 sets the class to A.

JES3 Set to the SYSOUT class. For incoming data sets, used if the class and type (indicated by NDHGF2PR and NDHGF2PU) do not conflict with the local system's use of class/type.

If a conflict exists, a different class will be used for the appropriate type. An example of a conflict is if the data set header indicates class 'K', type punch, but class 'K' is a print class on the destination system.

If the NDHGF2PR and NDHGF2PU flags are not present (i.e. data set originates from a down-level node), the class will only be used if it is 'B'. Otherwise, the class will default to 'A'.

RSCS Always used for the VM/370 output class and always set from this class unless overridden by SYSOUT keyword on the TAG.

POWER Set from output class, used as output class. (Only A-Z are valid - received files set to "A" if not in range A-Z.)

NDHGNREC 30 4 Binary 0

This is the count of the number of records in the data set.

JES2 Used as the line count of the data set. A spanned record counts as 1, regardless of the number of segments.

JES3 Set but not used.

RSCS This is used in a message to the user. It is set from the record count in the spool file block.

POWER This is set and used as the line count.

NDHGFLG1 34 1 Bits 0

This is a flag byte containing the following flags.

NDHGF1SP 34 .1 80 bit 0

This is the spin data set flag.

JES2 Set by JCL. This flag is not set for TSO/E Interactive Data Transmission Facility.

JES3 Set for spin data sets and TSO/E Interactive Data Transmission Facility, not used.

RSCS Defaulted, used to separate datasets into individual spool files.

POWER Set if data set is segmented. Used for determining job number: If flag is on, job number is set from job header (NJHGJID), otherwise a new job number is assigned.

NDHGF1HD 34 .1 40 bit 0

This bit tells the destination node to hold the data set.

JES2 Set via JCL. Used to place the data set on the system output hold queue.

JES3 Defaulted, used to place output in hold.

RSCS Release 3 - Defaulted, not used.
Version 2.1 - Used to spool file held at destination. Set when specified by originator on TAG command (HOLD keyword).

POWER Used. Set if POWER disposition is 'H' or 'L'.

NDHGF1LG 34 .1 20 bit 0

This is the job log indicator bit.

JES2 Indicates this data set is the JES2 job log.

JES3 Defaulted, not used.

RSCS Defaulted, not used.

POWER Defaulted, not used.

NDHGF1OV 34 .1 10 bit 0

This is the page overflow indicator. When set, output is to print over the fold. When not set, the output should skip to a new page at the appropriate spot for the final output device.

JES2 Defaulted, not used.

JES3 Set if OVFL=OFF specified. Used for the same purpose.

RSCS Defaulted, not used.

POWER Used, but not set.

NDHGF1IN 34 .1 08 bit 0

This bit requests the interpret feature be used for punched output if punched on a device with that feature installed.

JES2 Set from JCL.

JES3 Used and set.

RSCS Defaulted, not used.

POWER Defaulted, not used.

NDHGF1LC 34 .1 04 bit 0

JES2 Sets this flag to indicate that NDHGLNCT is set.

JES3 Defaulted, not used.

POWER Defaulted, not used.

RSCS Defaulted, not used.

NDHGF1ST 34 .1 02 bit 0

The job statistics are in the job log dataset.

JES2 Used to avoid multiple instances of the job statistics in the job log.

JES3 Defaulted, not used.

RSCS Defaulted, not used.

POWER Defaulted, not used.

NDHGF1DF 34 .1 01 bit 0

This bit specifies whether the interpret bit (NDHGF1IN) was explicitly set (1) or not (0).

The two bits, NDHGF1IN and NDHGF1DF, are used as follows:

```
NDHGF1IN    0  1  0  1
NDHGF1DF    0  0  1  1
                          Meaning:
            |  |  |  |
            |  |  |  +---- Interpret
            |  |  +------- Do not interpret
            |  +---------- Interpret
            +------------- Use default
```

JES2 Defaulted, not used.

JES3 Used to determine whether NDHGF1IN was explicitly set or defaulted.

RSCS Defaulted, not used.

POWER Defaulted, not used.

NDHGRCFM 35 1 Binary 0

This specifies the record format of the data set.

11.. .... is undefined format
10.. .... is fixed format
01.. .... is variable format
.... .00. is for no carriage control
.... .10. is for ASA control characters
.... .01. is for machine control characters

JES2 This field is set from JCL (RECFM in DCB). The DCB defines other settings which are not defined here; JES2 allows all flag settings and will send and receive them across the network.

JES3 Carriage control set to either ASA, machine or none (bits 5 & 6 off). Undefined format set (bits 0 & 1 on). Only carriage control bits are used on input.

RSCS Set to X'42' for print files and X'80' for punch files. For print files it is used to decide what record length to use in determining whether to make an incoming file a virtual 1403 file or 3211 file. See the discussion of NDHGLREC for further details.

For punch files, it is used, along with NDHGLREC, to determine the proper record length to use in detecting files containing punch records greater than 80 characters. It is also used in determining whether to forward punch records of a

store and forward file with or without carriage control. For further details, see B.3.5.1, "Store and Forward Transparency" on page B-27.

This field is modified for certain types of data for store and forward when RSCS acts as an intermediate node. For details see B.3.5.1, "Store and Forward Transparency" on page B-27.

POWER POWER version 2.1: Not used, set to X'42' (variable). Store and forward files are converted to machine carriage control.

POWER version 2.2: Not used, set to X'Cx' (undefined). At S&F nodes, files with ASA carriage control will not have the CC converted to machine carriage control. All values of the field will be supported.

NDHGLREC 36 2 Binary 0 1-32760

This is the maximum logical record length of any record that will appear within the data set. It includes the carriage control character if carriage control is specified for the data set.

JES2 LRECL is taken from the DCB. It includes the carriage control character.

JES3 Set and used if valid (not zero). Default is set according to the type of data set and the presence of carriage control as follows:

```
PRINT WITH CC:     133
PRINT WITHOUT CC: 132
PUNCH WITH CC:      81
PUNCH WITHOUT CC:   80
```

RSCS Used as follows:

1. For incoming print files, this determines whether they are to be defined as 1403 or 3211 files. If the record length specified in this field is greater than 132 (133 if NDHGRCFM specifies any type carriage control), files are made 3211. If less than this they are made 1403 files.
2. Used to obtain necessary storage for processing spanned records. This field must be filled in with the unspanned size of the maximum record in the file. Otherwise, RSCS will reject the file containing the spanned record with a receiver cancel X'B0' when it actually encounters the record.
3. Release 3 - If a punch file has this field set to a value greater than 80 and NDHGRCFM indicates no carriage control, then the file will be rejected by RSCS with a receiver cancel. A punch file is similarly rejected if this field is set to a value greater than 81 and NDHGRCFM indicates either machine or ASA carriage control.
   Version 2.1 - Punch records greater than 80 (greater than 81 with carriage control) are truncated at the destination node.

Set from TAGRECLN (with one added to each record to account for the CCTL byte) as follows:

1. Set to 80 for punch files.
2. Virtual 3800 files which contain any load CCWs have this field set to 8192 (regardless of the actual maximum record length in the file.)

Release 3 - This field is modified for certain types of data for store and forward when RSCS acts as an intermediate node.
Version 2.1 does not modify this field as an intermediate node.

For details see B.3.5.1, "Store and Forward Transparency" on page B-27.

POWER Used, set to maximum record length of any record in the data set. Can be maximum of 32K-1.

NDHGDSCT 38 1 Binary 0 0-255

This is the data set copy count and defines the number of copies of this data set to output at the destination.

JES2 This is the number of copies of this data set. It will be multiplied by the job copy count. This field will be used if the data set is printed on an impact printer. If it is printed on a non-impact printer, the copy group count is used. The maximum value for this field is 255.

JES3 Set and used. If received value is 0, a 1 is used.

RSCS Set from the copy count specified on the SPOOL command. Used to set the spool copy count.

POWER Set from COPY parm of * $$ LST or * $$ PUN statement. Used as copy count. Defaulted to 1.

NDHGFCBI 39 1 Binary 0 -31 to 31

This is the index byte for use when loading the FCB on a 3211. It gives the user the ability to print each line with the data shifted right or left by up to 31 characters. A negative value implies left indexing, a positive value implies right indexing. Right indexing adds leading blanks. Left indexing removes data characters.

JES2 Both right and left indexing are supported.

JES3 Defaulted, not used.

RSCS Not used, set if the INDEX keyword is specified on the TAG, only positive values are allowed.

POWER Defaulted, not used.

NDHGLNCT 3A 1 Binary 0

This is the lines per page for SYSOUT files. A value of X'00' or X'FF' causes the system not to count lines or use a local default value, depending on the system.

All other values are treated as an explicit number of lines on a page before a page eject is generated by the printing subsystem.

JES2 Set from JCL.

The special values X'00' and X'FF' are used as follows:

X'00' (the default): Use the default value of the 'printing' (i.e., destination) subsystem.

X'FF' Do not force any page ejects, but let the skipping be solely determined by the carriage control (if present) in the SYSOUT data.

JES3 Defaulted, not used.

RSCS Defaulted, not used.

POWER Defaulted, not used.

RESERVED 3B 1

NDHGFORM 3C 8 Character blanks

This is the name of the form to use when printing or punching this data set.

JES2 JES2 uses only the first four bytes, followed by blanks until SP1.3.3. Forms are set to zero if not specified. If received, either zero or blanks indicate the standard (STD.) forms are desired. After SP1.3.3 the blanks will indicate 'STD'.

JES3 Set to zero if forms not specified.

RSCS Set from CP FORM unless overridden by TAG. The default is only set when the CP SPOOL command specifies NULL for FORM name.

This field is always used for CP FORM name when not defaulted. When the header field contains the default (blanks), the CP FORM name is set to the installation's default.

POWER Only the first four bytes are used or set as POWER only has four character forms names. For a 3800 SYSOUT data set, a value of X'00' means use hardware defaults, while a value of X'40' means use software defaults. When job originated on the a non-VSE system (i.e., POWER section not present), X'00' is taken as software default.

NDHGFCB 44 8 Character blanks

This is the name of the FCB to use if this data set is printed.

JES2 JES2 uses only the first four bytes, followed by blanks. FCB is set to zero if not specified. If received, either zero or blanks indicate the default FCB (****) is desired. The first four bytes are prefixed with:

'FCB1' for 1403
'FCB2' for 3211 and 3203-5
'FCB3' for 3800

JES3 Set and used, but the default is zero.

RSCS Release 3 - Used for CP 3800 FCB only if 3800 section is present. Set from CP 3800 FCB only. It may also be specified on the CP TAG command for a 3211 or 3800 FCB. RSCS will not use this field for a 3211 FCB because such an FCB can not be specified on any CP command. Only the first four characters are used.

Version 2.1 - Used for FCB on spool file. Set from FCB on SPOOL command. Only the first four characters are used. May be overridden by the CP TAG command (up to eight characters).

POWER Used as the FCB name. Set from the FCB name if one was specified. With the fix for DY33952, the last four characters may be set to '$$$$' for device independence. For a 3800 SYSOUT data set, a value of X'00' means use hardware defaults, while a value of X'40' means use software defaults. When job originated on the a non-VSE system (i.e., POWER section not present), X'00' is taken as software default.

NDHGUCS 4C 8 Character blanks

This is the name of the Universal Character Set (UCS) to use if this data set is printed. It is more commonly known as the print train.

JES2 JES2 uses only the first four bytes, followed by blanks. UCS is set to zero if not specified. If received, either zero or blanks indicate the default UCS (****) is desired. NDHATAB1 is used if UCS is specified as either '****', blanks or zeros. The four characters from this field are prefixed by JES2 with:

'UCS1' for 1403
'UCS2' for 3211
'UCS3' for 3203-5

JES3 Only the first 4 bytes are set and used. Defaulted to zeros. The four characters from this field are prefixed by JES3 with:

'UCS1' for 1403
'UCS2' for 3211
'UCS3' for 3203-5

RSCS Not used, only set if specified by the UCS keyword in the TAG.

POWER Used as the UCS name. Set to the UCS name if one is specified in the * $$ LST statement.

NDHGXWTR 54 8 Character blanks

This is the name of the external writer to process this data set. External writers are used by the TSO/E Interactive Data Transmission Facility and VM/CMS SENDFILE to allow files to be sent to individual MVS users.

JES2 This is the name of the external writer, followed by blanks. JES2 sets a default of zero. Either zero or blanks are accepted as defaults for input.

RSCS Not used, set only if the EXTWTR keyword is specified in the TAG.

POWER Not used, set to the name of the external writer as specified in the DEST parameter.

RESERVED 5C 8

NDHGFLG2 64 1 Bits 0

This is a flag byte containing the following flags.

JES3 If the general section is not large enough to contain this flag byte, JES3 determines the type of data set from NDHGCLAS.

RSCS Release 3 - If the general section is not large enough to contain this flag byte, RSCS uses NDHGCLAS to determine the type of device. In this case only, CLASS B files are made to be punch and all other classes are made to be print. Version 2.1 always assumes the presence of this byte.

NDHGF2PR 64 .1 80 bit 0

This bit is set if the data set is to be printed.

JES2 Set by the origin node if the output class is a print class. Ignored if received.

JES3 Used to denote print output, set if print output.

RSCS Used to cause output to be spooled to a printer, set if input file is a print file.

POWER Used to denote print output, set if print output.

NDHGF2PU 64 .1 40 bit 0

This bit is set if the data set is to be punched.

JES2 Set by the origin node if the output class is a punch class. Ignored if received.

JES3 Used to denote punch output, set if punch output.

RSCS Used to cause output to be spooled to a punch, set if input file is a punch file.

POWER Used to denote punch output, set if punch output.

RESERVED 64 .1 20 bit

NDHGUCSO 65 1 Bits 0

This field contains the UCS option byte.

JES2 Defaulted, not used.

JES3 Defaulted, not used.

RSCS Defaulted, not used.

POWER Used and set if UCS specified.

NDHGUCSD 65 .1 80 bit 0

If this bit is on it means that the UCS will be loaded with the block data check option.

JES2 Defaulted, not used. JES2 always blocks data checks.

JES3 Defaulted, not used.

RSCS Defaulted, not used.

POWER Used and set. Defaulted to 0 (do not block data check)

NDHGUCSF 65 .1 40 bit 0

If this bit is on it means that the UCS will be loaded with FOLD option.

JES2 Defaulted, not used. JES2 loads the FOLD option on a character set if bit X'40' of the first byte (byte 0) is on in the UCS image.

JES3 Defaulted, not used.

RSCS Defaulted, not used.

POWER Used and set. Defaulted to 0 (do not FOLD).

RESERVED 66 2

NDHGPMDE 68 8 Character Blanks

This is a 1-8 character name (padded on the right with blanks) that specifies the process mode set by the user. It does not necessarily imply the presence or lack of stream mode records. It only indicates a request by the user for the preferred type of output processing. This field should be used for SYSOUT selection and scheduling.

Current values defined are:

LINE
PAGE
SOSI1
SOSI2

JES2 Set from PRMODE on the OUTPUT JCL statement. Any user-specified name may be present. If PRMODE is not specified on the user's JCL, then the data itself is examined to determine PRMODE.

JES3 Set from PRMODE on the OUTPUT JCL statement. If PRMODE is not specified on the user's JCL, then the data itself is examined to determine PRMODE.

RSCS Defaulted, not used.

POWER Defaulted, not used.

8.5.2 Record Characteristics Change Section

This section is identified by an identifier field of X'00' and a modifier field of X'40'. This is only used on SYSIN data and only present if the SYSIN data is other than RECFM F and LRECL 80. When present, it is sent as the only section in the data set header, without the general section with modifier X'00'. JES3 does not use or send these sections.

NDHCFLG1 4 1 Bits 0

This flag byte is not used.

NDHCRCFM 5 1 Binary 0

This is used to change the record format in the middle of a data set.

11.. .... is undefined format
10.. .... is fixed format
01.. .... is variable format

JES3 Unused.

RSCS Release 3 - Not set or used. RSCS never creates this section of the header.

Version 2.1 - Used at a store and forward node to determine if a SYSIN file should be stored as coded NOPs (done whenever NDHCRCFM indicates the file isn't fixed format). RSCS never creates this section of the header.

POWER RECFM always undefined.

NDHCLREC 6 2 Binary 0

This specifies a new maximum record length for records in the data set.

JES3 Unused.

RSCS Release 3 - Not set or used. RSCS will store and forward this section of the header. However, if the actual SYSIN data set contains records longer than 80 bytes, the data set will be rejected with receiver cancel. Any SYSIN records shorter than 80 bytes will be padded with blanks so that they will be 80 bytes long when store and forwarded.

Version 2.1 - Used at a store and forward node to determine if a SYSIN file should be stored as coded NOPs (done whenever NDHCLREC indicates record length not equal to 80).

POWER Changes if appropriate. Maximum POWER will send is 128.

8.5.3 3800 Section

This section is identified by an identifier field of X'00' and a modifier field of X'80'.

NDHAFLG1 4 1 Bits 0

This is a flag byte containing the following flags.

NDHAF1J 4 .1 80 bit 0

This is set if OPTCD=J is specified. It is equivalent to saying the file contains a Table Reference Character (TRC) in the first byte of every record (second if carriage control is present).

JES2 Used and set if OPTCD=J specified.

JES3 Used and set if OPTCD=J specified.

RSCS Used as follows:

1. When this flag is on the TRC bytes are stripped off each record and select translate table CCWs are inserted to handle the TRCs. This occurs both when a file is destined for the RSCS node and when it is store and forwarded. Store and forward files have the TRC bytes re-inserted and the selects removed when forwarded.
2. For Release 3, the TRC bytes are only re-inserted when the file is forwarded to a non-VM/SP NJI system.
3. For Release 3, if this flag is on and RSCS is running on a VM/SP2 (or above) system, the incoming file is made a virtual 3800 print file.
4. If RSCS Release 3 is running on a VM/SP1 system, the flag is not used.
5. On Version 2.1, the file is always made a virtual 3800 print file when this flag is on.

Set as follows (for Release 3 and Version 2.1):

1. This flag is turned on for all virtual 3800 files. Any virtual 3800 file will have TRC bytes inserted (in all records except those representing CCWs for intermediate operations and all spanned records). Any select translate table CCWs in these files will be removed. This action will occur for any virtual 3800 file regardless of the level of VM/SP that RSCS is running on.
2. For all other files the flag is only turned on if OPTCD=J is specified on the TAG command.

More information on how 3800 files are handled by RSCS may be found in "3800 Files" on page B-29.

POWER Used; set when a Spool Access Support (SAS) user indicates that the first character is a TRC.

NDHAF1BR 4 .1 40 bit 0

This is set if the 3800 burster is to be used.

JES2 Set from JCL.

RSCS Not used. Set if BURST = Y specified in TAG.

POWER Used, set from * $$ LST statement or from default printer setup.

NDHAF1BN 4 .1 20 bit 0

This is set if the 3800 burster is not to be used.

JES2 Defaulted, not used.

JES3 Set if STACKER = C is specified on the JES3 //*FORMAT statement. Indicates continuous forms stacking.

RSCS Not used. Set if BURST = N specified in TAG.

POWER Used, set from * $$ LST statement or from default printer setup.

NDHAFLCT 5 1 Binary 0

This is the 3800 flash count. This defines how many copies of the data set should be flashed. If not specified, but FLASH was specified, then all copies are flashed.

JES2 Set from JCL.

RSCS Set via the CP spool command or may be overridden by the CP TAG command. Used for flash count with the spool file.

POWER Used, set from * $$ LST statement or from default printer setup.

NDHATREF 6 1 Binary 0 Range 0-3

This is the table reference character. Specifies which of the four translate table entries should be used when printing the copy modification.

JES2 Set from JCL.

RSCS Used. Set by CP SPOOL command or may be overridden by the MODTRC keyword on the TAG command.

POWER Used, set from * $$ LST statement or from default printer setup.

NDHATAB1 8 8 Character blanks

This is the first translate table name. The translate table defines an index to a certain font which is defined for the 3800 printer. It means that by using different translate tables different character sets may be used within one print output. The translate table to be used is defined in the record when OPTCD = J is specified. Only the first four characters are used and set.

JES2 Used and set by JCL or JECL. Set to zero if defaulted by the user. If received, zero or blanks indicate default. If set to '****', blanks or zeros, then NDHGUCS is used.

RSCS Used, set from the CHARS operand on the CP spool command. May be overridden by the CHARS keyword in the TAG.

POWER Used and set as CHAR1. For a 3800 SYSOUT data set, a value of X'00' means use hardware defaults, while a value of X'40' means use software defaults. When job originated on the a non-VSE system (i.e., POWER section not present), X'00' is taken as software default.

NDHATAB2 10 8 Character blanks

This is the second translate table name. Only the first four characters are used and set.

JES2 Used and set from JCL or JECL. If not specified by user, then default set to zero. If received, zeros or blanks indicate default.

RSCS Used, set from the CHARS operand on the CP spool command. May be overridden by the CHARS keyword in the TAG.

POWER Used and set as CHAR2. Defaulted to X'00'.

NDHATAB3 18 8 Character blanks

This is the third translate table name. Only the first four characters are used and set.

JES2 Used and set from JCL or JECL. If not specified by user, then default set to zero. If received, zeros or blanks indicate default.

RSCS Used, set from the CHARS operand on the CP spool command. May be overridden by the CHARS keyword in the TAG.

POWER Used and set as CHAR3. Defaulted to X'00'.

NDHATAB4 20 8 Character blanks

This is the fourth translate table name. Only the first four characters are used and set.

JES2 Used and set from JCL or JECL. If not specified by user, then default set to zero. If received, zeros or blanks indicate default.

RSCS Used, set from the CHARS operand on the CP spool command. May be overridden by the CHARS keyword in the TAG.

POWER Used and set as CHAR4. Defaulted to X'00'.

NDHAFLSH 28 8 Character blanks

This is the flash cartridge identifier. This overlay will be printed on every page before the data is placed into position. It is used to produce "pre-printed" forms. Only the first four bytes are used.

JES2 Used and set from JCL or JECL. If not specified by user, then default set to zero.

JES3 Used. Set if flash specified, last four bytes are all zeros. All zeros if flash not specified.

RSCS Used and set from CP flash name, may be overridden by the FLASH keyword on the TAG.

POWER Used and set as FLASH. For a 3800 SYSOUT data set, a value of X'00' means use hardware defaults, while a value of X'40' means use software defaults. When job originated on the a non-VSE system (i.e., POWER section not present), X'00' is taken as software default.

NDHAMODF 30 8 Character blanks

This is the copy modification id. It contains the name of a module which is always placed on every page of the output data set when it is being printed. Only the first four bytes are used.

JES2 Used and set from JCL or JECL. If not specified by user, then default set to zero.

JES3 Used. Set if copy mod specified, last 4 bytes are zero. All zeros if copy mod not specified.

RSCS Used and set from the CP MOD name, may be overridden by the MODIFY keyword on the TAG.

POWER Used and set as copy modification. In the DOS/VSE 3800 the user can specify an additional character arrangement table to be used for the copy modification, which does not need to be specified in the CHARS parameter. POWER will default to the first CHAR if so. For a 3800 SYSOUT data set, a value of X'00' means use hardware defaults, while a value of X'40' means use software defaults.

NDHACPYG 38 8*1 Binary 0

This is the copy groups. It is really eight one byte fields. The fields define how many times each page of the data set will get copied when the data set is sent to the 3800. The first field refers to the first transmission, the second field to the second transmission, etc. In this case NDHGDSCT is NOT used. The number of transmissions is determined by how many copy groups are defined.

JES2 Used and set. If received, the total must be less than 255 or JES2 will set it to 255. It is not multiplied by the job copy count.

RSCS The sum of all these bytes is used as the CP copy count and the CP copy group flag is turned on. The CP copy count goes into the first byte if the CP copy group flag is on. This may be overridden by the COPYG keyword on the TAG.

POWER Used as copy groupings. Specified in * $$ LST statement or by program. Sum of all copy groups, or a single copy group, cannot exceed 255.

8.5.4 RSCS Section

This section is identified by an identifier field of X'87' and a modifier field of X'00'.

NDHVFLG1 4 1 Bits 0

This flag byte is not used.

NDHVCLAS 5 1 Character A-Z,0-9

This is the CP spool file class.

NDHVIDEV 6 1 Binary

This is the origin CP device type. See CP DEVTYPES macro for valid values.

NDHVPGLE 7 1 Binary

This is the 3800 virtual page length.

NDHVDIST 8 8 Character blanks

This is the CP distribution code.

NDHVFNAM 10 12 Character blanks

This is the CP file name.

NDHVFTYP 1C 12 Character blanks

This is the CP file type.

NDHVPRIO 28 2 Binary 50 0-99

This is the RSCS transmission priority as specified in the TAG.

NDHVVRSN 2A 1 Binary

Version number of RSCS system which created the header.

NDHVRELN 2B 1 Binary

Release number of the RSCS system which created the header.

NDHVTAGR 2C 136 Character

RSCS Version 2.1 only - Tag record as specified on the TAG command.

8.5.5 POWER Section

This section is identified by an identifier field of X'86' and a modifier field of X'00'.

NDHPFLG1 4 1 Bits 0

This flag byte is not used.

NDHPIDEV 5 1 Binary 0

VSE device type.

NDHPPRIO 6 1 Character

Output priority. Defaulted to the priority as specified in the POWER generation.

NDHPDISP 7 1 Character D

Output disposition. May be "D" (Delete), "K" (Keep), "H" (Hold), or "L" (Leave).

NDHPUSER 8 16 Character blanks

Contains any user information which might have been specified. Used for printing on separator page.

NDHPJBSF 18 1 Binary 0

Job suffix number. This is created whenever we have output which is segmented. It is similar to the spin data set number.

NDHPSYID 19 1 Character zero

Contains the system qualifier for a shared spool configuration. Used for printing on separator page.

NDHPNSEP 1A 1 Binary 0

Number of separator pages which should be printed for this output.

NDHPOPTN 1B 1 Bits 0

Contains the COPYSEP parameter as specified for the output.

NDHPPART 1C 2 Character blank

Contains the partition-id of the partition in which the job executed. Used for printing on separator page.

RESERVED 1E 2 Character blank

NDHPRCFM 20 1 Bits 0

Specifies the record format (e.g., SCS, BMS, 3270).

RESERVED 21 1 Character blank

NDHPJNUM 22 2 Binary 0

Job number.

NDHPCOMP 24 4 Character blank

Compaction table name used for output destined to RJE/SNA

NDHPPASS 28 8 Character blank

The password which has been specified for this output.

NDHPSETP 30 68 Character blank

Default SETPRT parameter list.

NDHPSTRT 74 8 Character blank

Time (in STCK format) when output spooling started.

8.5.6 Output Processing Section of the Data Set Header

Note: This is also called the "data stream" section by JES2 and JES3.

A separate section of the data set header has been defined for transmitting output processing parameters and control formats. The identifier for this output processing section is X'89'. A modifier of X'00' indicates that Output Parameter Text Blocks (OPTBs) are being transmitted in this section, and is currently the only modifier defined or allowed in this section.

8.5.6.1 Relationship to the 3800 Section

The 3800 section of the data set header is still built if any 3800 attribute is specified by the user. Some fields are duplicated in the Output Parameter Text Blocks (OPTBs). However, if any of the 3800 fields is changed by operator command, those changes need only be reflected in the 3800 section, not in the OUTPUT OPTBs. The 3800 section fields override the OPTBs at the SYSOUT destination.

8.5.6.2 Field Definitions

The output processing section is as follows. Note that all lengths include the length of the length field.

NDHSFLEN 4 2 Binary

This is the length of the fixed area of the section, down through the field "NDHSGPID", but not including NDHSOPTB or any of the OPTBs.

RESERVED 6 2 Binary X'0000'

NDHSJDVT 8 8 Binary X'00's

The JDVT level (identifier) at the execution node. A value of zero indicates the default JDVT.

JES2 Set to binary zeros, used to determine JDVT for SJF processing.

JES3 Set to binary zeros, used to determine JDVT for SJF processing.

RSCS Not used or set.

POWER Not used or set.

NDHSNSTR 10 4 Binary

The number of stream mode 'begin page' structures are in the data set; i.e., the number of data records that begin with the sequence X'D3A8AF'. This field is not adjusted for multiple copies.

JES2 Sets and uses this field.

JES3 Sets and uses this field.

NDHSGPID 14 8 Character Blanks

Output Group Name. This name may be provided by the user; otherwise it is generated at

job execution time. This name is used by the receiving node to determine how the job's data sets are grouped together for printing.

JES2 Set from the user's output JCL, or generated at job execution time.

JES3 Defaulted, not used.

| **NDHSOPTB** 1C Mixed

| **NDHSPRID** 1C 4 Character 'SJPF'

| Required constant that identifies the prefix within MVS systems.

| **NDHSVERS** 20 1 Binary X'02'

| Required constant that identifies the version of the prefix. This constant should be set, but
| may be ignored when received by non-MVS systems.

| **MVS** **01** = MVS/SP 1.3.1 (FMID JBB1327)
| **02** = MVS/SP 1.3.3 (FMID JBB1329)

| **NDHSPLEN** 21 1 Binary X'1C'

| Prefix length. Used to point to the beginning of variable length OPTBs. That is, start of
| OPTBs = NDHSOPTB + value of NDHSPLEN. Fixed at X'1C' for the version 2
| OPTB.

| **NDHSDLEN** 22 2 Binary X'0000'

| Data length following the prefix.

| **NDHSVERB** 24 8 Character 'OUTPUT '

| Constant required for MVS systems compatibility. Identifies the JCL statement normally
| used to specify the following parameter values in the MVS environment.

| **RESERVED** 2C 8 Character blanks

| **NDHSFLG2** 34 1 Bits X'80'

| Flag byte. Reserved.

| **RESERVED** 35 1

| **RESERVED** 36 2 Binary X'0000'

| The end of the last field above is the start of the variable area for the Output Parameter Text Blocks
| (OPTBs). Its structure reflects its origin as an internal JES2 and JES3 text unit.

8.5.6.3 Syntax of Output Parameter Text Block (OPTB)

The OPTB is structured as a fixed area followed by a sequence of text data units. The number of and sequencing of text data units is arbitrary, hence the appearance of any specific text data unit is optional. When text data units do not appear that would have supplied values needed at the destination to process the data, installation or product defined defaults may be used.

Some defaults have been suggested in the keyword definition table.

OPTB Structure Definition: The structure of an OPTB can be defined in terms of a set of rules for building them. The rules have been written in the "BNF" format, where elements of the structure are indicated in brackets, such as <ELEMENT>, and the definition of the element appears to right, such as <ELEMENT> := definition. The sequence of elements allowed is indicated by explicit sequences of bracketed items. Alternative sequences are delimited with the symbol "|" which means "OR".

```
<DSHOPT> := Output processing section of the data set header
         := <FIXED><OPTBS>

 <FIXED> := Fixed portion of the output processing section
            Starting with NDHSLEN through NDHSGPID

 <OPTBS> := <OPTB>

<OPTB>   := <PREFIX><TDUS>

<PREFIX> := Fixed portion of the OPTB starting with NDHSPRID
            through the RESERVED field at offset X'36'.

<TDUS>   := <TDU>|<TDUS><TDU>

<TDU>    := <KEY><COUNT><PARMS>

 <KEY>   := Two byte key defined in Keyword definition Table

 <COUNT> := Two byte count of number of parms: 1-255

 <PARMS> := <PARM>|<PARMS><PARM>

<PARM>   := <LL><DATA>

 <LL>    := Two byte data length: 1-128

 <DATA>  := Sequence of exactly <LL> bytes
```

| ***Text Data Unit Structure:*** Text data units are identified by a 2-byte key that is registered and unique within NJE. Each text data unit is defined to include a specific type and maximum number of data elements that represent keyword parameter values. The number of these parameters included in the text data unit is specified in a 2-byte count field that follows the key and precedes the parameter list.

The figure below illustrates this structure.

KEY	COUNT	LENGTH1	data1	LENGTH2	data2	...	...
key	cc	ll	data	ll	data	...	...

Where:

key Two byte registered keyword identifier. All values not defined are reserved for future use, and are should not be specified.

cc Two byte count of the number of values supplied for the keyword parameter. Range is from 1 - 16383. A count of 0 is used to indicate either a missing positional parameter or a defaulted parameter. In this case, no data elements should follow the count field.

ll Two byte length of the keyword parameter value. Range is from 1 - 16383. A length of 0 indicates a null value.

| *Note: For compatibility with MVS systems, the parameter length must be restricted by the protocols to 128 bytes. Lengths greater than 128 will be correctly stored and forwarded by JES, but cannot be used at an MVS destination.*

data Keyword values associated with the key.

8.5.6.4 Parameter Prefix

Description of the fixed area based on SJF prefix.

8.5.6.5 Keyword Definition Table

The suggested defaults listed in the following table may be used when installation or other user-supplied values are not available.

Key	Max Repeat	Length Range	Keyword	Data Values Allowed
0001	1	1	Reserved	Reserved - private JES2/JES3
0002	4	4	Reserved	Reserved - private JES2/JES3
0003	1	2	CKPTLINE	Integer: range 0 - 32767
0004	1	2	CKPTPAGE	Integer: range 1 - 32767
0005	1	2	CKPTSEC	Integer: range 1 - 32767
0006	1	1	Reserved	Reserved - private JES2/JES3
0007	1	8	COMPACT	alphameric
0008	1	1	CONTROL	X'80' - Force single space X'40' - Force double space X'20' - Force triple space X'10' - Use first character in line as CC.
0009	1	1	Reserved	Reserved - private JES2/JES3
000A	8	1	Reserved	Reserved - private JES2/JES3
000B	8	1	Reserved	(Not used at output nodes)
000C	1	17	Reserved	Reserved - private JES2/JES3
000D	1	4	Reserved	Reserved - private JES2/JES3
000E	1	4	Reserved	Reserved - private JES2/JES3
000F	1	1	Reserved	Reserved - private JES2/JES3
0010	1	8	Reserved	Reserved - private JES2/JES3
0011	1	8	Reserved	Reserved - private JES2/JES3
0012	1	1	INDEX	Integer: range 1 - 31
0013	1	1	Reserved	(Not used at output nodes)
0014	1	1	LINDEX	Integer: range 1 - 31
0015	1	1	Reserved	Reserved - private JES2/JES3
0016	1	4	Reserved	Reserved - private JES2/JES3
0017	1	1	Reserved	Reserved - private JES2/JES3

Key	Max Repeat	Length Range	Keyword	Data Values Allowed
0018	1	8	Reserved	Reserved - private JES2/JES3
0019	1	1	PRTY	Integer: range 0 - 255
001A	1	1	Reserved	Reserved - private JES2/JES3
001B	1	4	Reserved	Reserved - private JES2/JES3
001C	1	8	WRITER	alphameric or national
001D	1	1-6	FORMDEF	name: alphameric or national
001E	-	-	Reserved	Reserved for future use
001F	1	1-6	PAGEDEF	name: alphameric or national
0020	-	-	Reserved	Reserved for future use
0021	-	-	Reserved	Reserved for future use
0022	1	4	THRESHLD	Integer: range 1 - 99999999
0023 -- 1FFF			Reserved	Not used in NJE

8.5.6.6 Keyword Semantics and Implementation Notes

The following format is used to describe the keys in the output processing section.

KEYWORD key length type range

Except where noted otherwise, the following semantics are supported on both VM, MVS and VSE Systems.

Note: It is acceptable for implementations to ignore keys for which no support exists; however, it is not acceptable to flag as an error, any key which is defined herein.

ALPHAMERIC The characters A through Z, a through z, and 0 through 9.

CKPTLINE 0003 2 Binary 0-32767

A value from 0 to 32767 that is the maximum number of lines contained in a logical page. This value is used to determine when to take checkpoints for printed output or SNA data sets. Installation defaults may be used when this parameter is not specified by the user.

VSE Not supported.

VM Not supported.

CKPTPAGE 0004 2 Binary 1-32767

A value from 1 to 32767 that is the number of logical pages to be printed or transmitted before the next output data set checkpoint is taken. This value represents the number of pages transmitted as a single SNA chain to a SNA work station.

VSE Not supported.

CKPTSEC 0005 2 Binary 1-32767

A value from 1 to 32767 that is the number of seconds that may elapse between printer checkpoints. Installation defaults may be used when this parameter is not specified by the user.

VSE Not supported.

VM Not supported.

COMPACT 0007 1-8 Character

A symbolic name, from 1 to 8 alphameric characters long, used to determine the compaction table when sending the SYSOUT data set described by this control statement to a SNA remote terminal. This specification overrides any remote-device-defined compaction table. Installation defaults may be used when this parameter is not specified by the user.

VSE Not supported.

VM Not supported.

CONTROL 0008 1 Binary X'10'-'80'

Type of forms control used.

X'80' = **SINGLE** indicates forced single spacing.

X'40' = **DOUBLE** indicates forced double spacing.

X'20' = **TRIPLE** indicates forced triple spacing.

X'10' = **PROGRAM** indicates that the carriage control character is the first character of each logical record in the data set.

Installation defaults may be used when this parameter is not specified by the user.

VSE Not supported.

VM Not supported.

FORMDEF 001D 1-6 Character

Specifies the 1- to 6-character (alphameric or national) member name of of the SYS1.IMAGELIB partitioned data set containing information that the 3800 printer model 3 uses to print a data set. Note that the first two characters of the member name are predefined by installation conventions. and are prefixed to the name specified here.

The members can contain the following information:

- Overlays that are to be invoked during output processing

- Location on the page where the overlays are placed
- Suppressions that can be activated for specified page formats

VSE Not supported.

INDEX 0012 1 Binary 1-31

Specifies a value (from 1 to 31) that indicates the data set indexing print position offset (to the right) for the 3211 printer.

VSE Not supported.

VM Not supported.

LINDEX 0014 1 Binary 1-31

A value (from 1 to 31) that indicates the data set indexing print position offset (to the left) for the 3211 printer.

VSE Not supported.

VM Not supported.

NATIONAL The characters #, @, and $.

PAGEDEF 001F 1-6 Character

Specifies the 1- to 6-character (alphameric or national) name of a member of SYS1.IMAGELIB containing the information used that the 3800 printer model 3 uses to print a data set. Note that the first two characters of the membername are pre-defined by installation conventions and are prefixed to the name specified here.

The member can contain the following information:

- Logical page size and width
- Fonts you want to use to print the page
- Page segments you want to use to print the page
- Definition of multiple page types or formats
- Definition of lines within a page; for example:
  - line origin
  - carriage controls
  - spacing instructions
  - font specifications
  - suppression control
  - constant data generation
- Definition of multiple logical pages on a physical page and line-skipping instructions.

VSE Not supported.

PRTY 0019 1 Binary 0-255

A value from 0 to 255 that represents the priority of the output data set for output queuing. 0 is the lowest, and 255 is the highest.

MVS Supported as documented.

VM Not supported.

VSE Not supported.

THRESHLD 2033 4 Binary 1-99999999

Specifies the maximum size for the sysout data set, before a new unit of work is created on a data set boundary. The size is based on the number of records times the number of copies.

JES2 Not supported.

JES3 Supported as documented

VM Not supported.

VSE Not supported.

WRITER 001C 8 Character

The 1 to 8 character name of an installation-written program, in the system library, that is to write the output data set.

VSE Not supported.

VM Not supported.

8.5.6.7 Error Handling

Destination nodes that do not process Output Parameter Text Blocks (OPTBs) may either reject the SYSOUT (X'B0' RCB) or accept the SYSOUT and subsequently ignore the data, print it (with unpredictable results), or hold the SYSOUT. A message describing the problem must be sent to the NOTIFY node and USERID in the Job Header.

Destination nodes that process Output Parameter Text Blocks (OPTBs) must adhere to the following rules in the case of errors:

Invalid text block format: A system dependent action may be taken, such as rejecting the file or substituting an entire set of defaulted parameters. A message must be issued to either the destination or originating node.

Keyword parameter values out of range: Specified defaults must be used unless overridden by the installation. No message is required.

Invalid keyword identifier: Invalid keywords must not be processed. They may be ignored, or the file may be rejected. A message should be issued at the destination node but is not required unless the file is rejected.

8.5.7 User Section

This section is identified by an identifier field of B'11xxxxxx' and a modifier field of the users choice. This section is not used by any IBM products, but will be passed through the network.

Note: Beyond the four byte header, this definition is the responsibility of the installation. The section length is limited by the design to 32764. The combined length of all sections in the header is limited to 32764, although product implementations may be more restrictive. The length must be reflected in the NDHULEN field in the front of the control record section header. See 8.3, "Control Record Section Header" on page 8-2.

A sample four byte field is shown below, only as an example:

NDHUCODE 4 4 Character blanks

This could be used for the SHARE/GUIDE/SEAS installation code, or some other unique identifier.

8.6 Job Trailer

For an overview of the Network Job Trailer, please see 3.6, "Job Trailer" on page 3-18.

8.6.1 General Section

This section is identified by an identifier field of X'00' and a modifier field of X'00'.

NJTGFLG1 4 1 Bits 0

There are no bits defined in this flag byte.

NJTGXCLS 5 1 Character A A-Z,0-9

This is the actual execution class of the job.

JES2 This is the execution class. The default is 'A'.

JES3 Not used, set to zero.

RSCS Not used.

POWER Not used.

RESERVED 6 2

NJTGSTRT 8 8 Binary 0

This is the actual time the job started execution in 370 TOD clock format (GMT).

JES2 Set and used for SMF type 26 records.

JES3 Defaulted, not used.

RSCS Defaulted, not used.

POWER Not used, set from STCK value.

NJTGSTOP 10 8 Binary 0

This is the actual time the job completed execution in 370 TOD clock format.

JES2 Time job completed, but will be zero for spin datasets. Used in SMF type 26 records.

JES3 Defaulted, not used.

RSCS Defaulted, not used.

POWER Not used, set from STCK value, zero for spin output.

RESERVED 18 4

NJTGALIN 1C 4 Binary 0

This is the total number of print lines for this job or job segment at all locations. This is not multiplied by the number of copies.

JES2 Used in SMF type 26 records. For spin output, this field is set to zeros if the sysout is sent before the job finishes execution.

JES3 Defaulted, not used.

RSCS Not used, set to the number of records in the file (from TAGRECNM) if it is a print file.

POWER Not used. This is set to the number of records produced by the total job. For spin output, all trailers except the last have zero.

NJTGACRD 20 4 Binary 0

This is the total number of card images produced for this job or job segment at all locations. This is not multiplied by the number of copies.

JES2 Used in SMF type 26 records. For spin output, this field is zero.

JES3 Defaulted, not used.

RSCS Not used, set to the number of records in the file (from TAGRECNM) if it is a punch file (SYSOUT) or SYSIN.

POWER Not used. This is set to the number of records produced by the total job. For spin output, all trailers except the last have zero.

RESERVED 24 4

NJTGIXPR 28 1 Binary 0 0-15

This is the initial requested execution selection priority.

JES2 Used in SMF type 26 record. Not set. (Low-order 4 bits used and shifted to high-order nibble of JCTIPRIO by sysout receiver.)

JES3 Defaulted, not used.

RSCS Not used. Set to RSCS transmission priority as shown in NJHGPRIO.

POWER Set from the PRI= parameter on the * $$ JOB statement.

NJTGAXPR 29 1 Binary 0 0-15

This is the actual execution selection priority used.

JES2 Set, used in SMF type 26 record.

JES3 Defaulted, not used.

RSCS Not used. Set to RSCS transmission priority as shown in NJHGPRIO.

POWER Defaulted, not used.

NJTGIOPR 2A 1 Binary 0 0-255

This is the initial job priority for output processor selection.

JES2 Set to user-specified or JES2-computed job priority at the time when the job was selected by the SYSOUT transmitter. Only the high-order 4 bits (values 0 thru 15) are used and shifted down to low-order bits. Set to 1 for spin data sets if job is still executing.

Used in SMF type 26 record.

Not used. If greater than 15 when received, then 15 is used as the selection priority.

JES3 The greater value of NJTGIOPR or NJTGAOPR is used for the job's priority for received output (unless it's greater than 15, in which case 15 is used). Set from the job's priority for outgoing SYSOUT.

RSCS Not used. Set to RSCS transmission priority as shown in NJHGPRIO.

POWER Set from the PRI= parameter on the * $$ LST or * $$ PUN statement.

NJTGAOPR 2B 1 Binary 0 0-255

This is the actual output selection priority used.

JES2 Defaulted, used in SMF 26 record.

JES3 The greater value of NJTGIOPR or NJTGAOPR is used for the job's priority for received output (unless it's greater than 15, in which case 15 is used). Set from the job's priority for outgoing SYSOUT.

RSCS Not used. Set to RSCS transmission priority as shown in NJHGPRIO.

POWER Defaulted, not used.

8.6.2 Accounting Section

This section has been created with an identifier of X'89' and a modifier of X'00' to transmit the actual output page mode and total byte count fields across the network.

Note that these statistics are for the entire job, not just the SYSOUT data sets created within this stream.

NJTSAPAG 4 4 Binary

This is the actual number of 'begin page' structured fields.

JES2 Set and used for SMF.

JES3 Not set or used.

NJTSABYT 8 4 Binary

This is the actual number of bytes.

JES2 Set and used for SMF.

JES3 Not set of used.

8.6.3 User Section

This section is identified by an identifier field of B'11xxxxxx' and a modifier field of the users choice. This section is not used by any IBM products, but will be passed through the network.

Note: Beyond the four byte header, this definition is the responsibility of the installation. The section length is limited by the design to 32764. The combined length of all sections in the header is limited to 32764, although product implementations may be more restrictive. The length must be reflected in the NJTULEN field in the front of the control record section header. See 8.3, "Control Record Section Header" on page 8-2.

A sample four byte field is shown below, only as an example:

NJTUCODE 4 4 Character blanks

This could be used for the SHARE/GUIDE/SEAS installation code, or some other unique identifier.

8.7 Nodal Message Record

Please see 4.1, "Nodal Message Record" on page 4-1 for an overview of the NMR. This control block is used in three different modes. The correct format should be used when looking up information in this section. See 8.7.1, "Unformatted Commands," 8.7.2, "Formatted Commands" on page 8-54 , or 8.7.3, "Messages" on page 8-59, depending on the type of NMR.

FLAGC	TYPEF	"Type" of NMR
1	0	Unformatted Command
1	1	Formatted Command
0	N/A	Message

Figure 8-1. NMR Types. This describes the type of NMR depending on the setting of NMRFLAGC and NMRTYPEF.

8.7.1 Unformatted Commands

Unformatted commands are sent in NMRs and can be distinguished from other NMRs because they have the NMRFLAGC ("Command") bit on, and the NMRTYPEF ("Formatted") bit off.

NMRFLAG 0 1 Bits 0

This byte defines the following flags.

NMRFLAGC 0 .1 80 bit 1

This bit must be set to 1 to define a command as opposed to a message.

NMRFLAGW 0 .1 40 bit 0

This bit says NMROUT has the origin remote number.

JES3 For output commands, set if FROM a remote. Not used for input commands.

RSCS Used, not set.

POWER Used and set if NMROUT is a remote station.

NMRFLAGT 0 .1 20 bit 0

This bit says NMROUT has the origin userid.

JES3 Not used or set for commands.

RSCS Used and set.

POWER Not used, not set.

NMRFLAGU 0 .1 10 bit 0

This bit says NMROUT has originating console identifier.

JES3 For output commands, set if from a JES3 console. Not used for input.

RSCS Not used, not set. Not forwarded in Release 3.

POWER Not used

NMRFLAGR 0 .1 08 bit 0

This bit says the console is only remote authorized.

JES3 Not used for input. Set for all output commands.

RSCS Not used, not set. Not forwarded in Release 3.

POWER Not used, not set.

NMRFLAGJ 0 .1 04 bit 0

This bit says the console is not job authorized.

JES3 Not used or set.

RSCS Not used, not set. Not forwarded in Release 3.

POWER Not used, not set.

NMRFLAGD 0 .1 02 bit 0

This bit says the console is not device authorized.

JES3 Not used or set.

RSCS Not used, not set. Not forwarded in Release 3.

POWER Not used, not set.

NMRFLAGS 0 .1 01 bit 0

This bit says the console is not system authorized.

JES3 Not used or set.

RSCS Not used, not set. Not forwarded in Release 3.

POWER Not used, not set.

NMRLEVEL 1 High 4 bits Binary 0

This is the importance level. See Figure 8-4 on page 8-64 for the bit definitions.

JES3 Not used. Set to 7.

RSCS Not used. Set to 7. Not forwarded unless 7 in Release 3.

POWER Not used. Set to 7.

NMRPRIO 1 Low 4 bits Binary 0

This is the output priority.

JES3 Not used. Set to 7.

RSCS Not used. Set to 7. Not forwarded unless 7 in Release 3.

POWER Not used. Set to 7.

NMRTYPE 2 1 Bits 0

This byte defines the following flags.

NMRTYPEX 2 .4 F0 bits 0

These are reserved bits.

NMRTYPE4 2 .1 08 bit 0

This bit does not apply for commands.

NMRTYPET 2 .1 04 bit 0

This bit does not apply for commands.

NMRTYPEF 2 .1 02 bit 0

This is the formatted command bit. It must be zero.

NMRTYPED 2 .1 01 bit 0

This bit does not apply for commands.

NMRML 3 1 Binary 1 1-132

This is the length of the information in NMRMSG.

NMRTO 4 9 Mixed

This field is made up of the following two fields.

NMRTONOD 4 8 Character blanks

This is the destination node for the command.

NMRTOQUL C 1 Binary 0

This is the qualifier for the node. This is the number of the member of the JES2 MAS node or the index to the JES3 system to which the command is being sent.

JES2 Set to zero unless specified.

JES3 For output commands, set to X'00' or the destination system qualifier specified on the *SEND command. Not used for input commands (all commands are processed on the global).

RSCS Defaulted, not used. Not forwarded in Release 3.

POWER Defaulted, not used.

NMROUT D 8 Character zeros

This is the origin userid, the origin remote id, or the origin console id depending on the setting of NMRFLAG.

JES3 Set to either the origin remote or the origin JES3 console (the format for consoles is the same as messages - NMRDESC = 0000, NMRROUT = JES3 console number, NMRDOMID = 00000000). Used for the command origin - either a console or remote - and will be used for command responses.

RSCS This contains the origin userid. It is always used as the origin userid if NMRFLAGT or FLAGW is on.

POWER This can contain a remote identification (Rnnn) or an ICCF userid. It is defaulted if the central operator issued the command. Used as the target remote/userid for command responses.

NMRFM 15 9 Mixed

This field is made up of the following two fields.

NMRFMNOD 15 8 Character blanks

This is the origin node.

NMRFMQUL 1D 1 Binary 0

This is the qualifier for the origin node. This is the number of the member of the JES2 MAS node or the index to the JES3 system from which the command is being sent.

JES2 Set to MAS member number.

JES3 For input commands, designates the origin system qualifier, and will be used for command responses. Set to X'00' for all output commands.

RSCS Not used or set. Not forwarded in Release 3.

POWER Used as NMRTOQUL for command responses.

NMRMSG 1E 132 Mixed

This is the actual data for the command.

8.7.2 Formatted Commands

Formatted commands are sent in NMRs and can be distinguished from other NMRs because they have the NMRFLAGC ("Command") bit on, and the NMRTYPEF ("Formatted") bit on.

JES2 is the only system that currently sends and accepts formatted commands.

RSCS and JES3 only support formatted commands for input. They are changed to equivalent RSCS or JES3 commands prior to processing. Therefore, in the following section, the RSCS and JES3 description will only indicate the usage of the fields, since RSCS and JES3 never send out formatted commands. See section "Formatted Commands" on page B-25 for more information on RSCS specific processing, and section "NMR Command Length Restriction" on page B-15 for JES3 specific processing.

POWER does not send these commands. If received they are thrown away without any notification.

NMRFLAG 0 1 Bits 0

This byte defines the following flags.

NMRFLAGC 0 .1 80 bit 1

This bit must be set to 1 to define a command as opposed to a message.

NMRFLAGW 0 .1 40 bit 0

This bit says NMROUT has a JES2 remote number.

JES3 Not used.

RSCS Not used. Not forwarded in Release 3.

NMRFLAGT 0 .1 20 bit 0

This bit says NMROUT has a userid.

JES3 Not used.

RSCS Used. Not forwarded in Release 3.

NMRFLAGU 0 .1 10 bit 0

This bit says NMROUT has originating console identifier.

JES3 Not used.

RSCS Not used. Not forwarded in Release 3.

NMRFLAGR 0 .1 08 bit 0

This bit says the console is only remote authorized.

JES3 Not used.

RSCS Not used. Not forwarded in Release 3.

NMRFLAGJ 0 .1 04 bit 0

This bit says the console is not job authorized.

JES3 Not used.

RSCS Not used. Not forwarded in Release 3.

NMRFLAGD 0 .1 02 bit 0

This bit says the console is not device authorized.

JES3 Not used.

RSCS Not used. Not forwarded in Release 3.

NMRFLAGS 0 .1 01 bit 0

This bit says the console is not system authorized.

JES3 Not used.

RSCS Not used. Not forwarded in Release 3.

NMRLEVEL 1 High 4 bits Binary 0

This is the importance level. See Figure 8-4 on page 8-64 for the bit definitions.

JES3 Not used.

RSCS Not used. Not forwarded in Release 3.

NMRPRIO 1 Low 4 bits Binary 0

This is the output priority.

JES3 Not used.

RSCS Not used. Not forwarded in Release 3.

NMRTYPE 2 1 Bits 0

This byte defines the following flags.

NMRTYPEX 2 .4 F0 bits 0

These are reserved bits.

NMRTYPE4 2 .1 08 bit 0

This bit does not apply for commands.

NMRTYPET 2 .1 04 bit 0

This bit does not apply for commands.

NMRTYPEF 2 .1 02 bit 0

This is the formatted command bit. It must be on.

NMRTYPED 2 .1 01 bit 0

This bit does not apply for commands.

NMRML 3 1 Binary 1 1-132

This is the length of the information in NMRMSG, includes NMRFR even if not specified.

NMRTO 4 9 Mixed

This field is made up of the following two fields.

NMRTONOD 4 8 Character blanks

This is the destination node for the command.

NMRTOQUL C 1 Binary 0

This is the qualifier for the node. This is the number of the member of the JES2 MAS node or the index to the JES3 system to which the command is being sent.

JES2 Set to zero.

JES3 Not used.

RSCS Not used. Not forwarded in Release 3.

NMROUT D 8 Character zeros

This is the origin userid, the origin remote id, or the origin console id depending on the setting of NMRFLAG.

RSCS Used as the origin userid.

NMRFM 15 9 Mixed

This field is made up of the following two fields.

NMRFMNOD 15 8 Character blanks

This is the origin node.

NMRFMQUL 1D 1 Binary 0

This is the qualifier for the origin node. This is the number of the member of the JES2 MAS node or the index to the JES3 system from which the command is being sent.

JES3 This designates the origin system qualifier and will be used for command responses.

RSCS Not used. Not forwarded in Release 3.

NMRMSG 1E 132 Mixed

This is the actual data for the command and is defined as follows.

NMRFNORM 1E 20 Binary

This is the area for a normal formatted command.

JES2 Initialized to zeros.

NMRFRTE 1E 36 Binary

This is the area for a route command.

JES2 Initialized to zeros.

NMRFOP 1E 1 Binary

This is the operation for the formatted command. The operations are defined as:

1. NMRFOPD - display job
2. NMRFOPC - cancel job
3. NMRFOPA - release job
4. NMRFOPH - hold job
5. NMRFOPR - route job

JES2 Used and set as defined with the exception that STCs and TSUs in execution will not be routed.

JES3 Display command converted to JES3 inquiry command. Cancel, hold, and release converted to JES3 modify command with the appropriate keyword. Route commands ignored.

RSCS Although not sent by RSCS, it is received and translated into an RSCS command. Display command converted to RSCS QUERY command. Cancel command converted to RSCS PURGE command. Hold and release commands converted to RSCS CHANGE command with appropriate keyword. Route command converted to RSCS TRANSFER command. See "Formatted Commands" on page B-25 for details.

NMRFFLG 1F 1 Bits 0

This byte contains flags or opcode modifiers.

JES2 Initialized to zeros.

NMRFFLGO 1F .1 80 bit 0

If on for Cancel, purge output. If on for Route, route job output as opposed to routing a job for execution.

NMRFFLGD 1F .1 40 bit 0

This is the cancel job execution with dump flag. This field is mutually exclusive with NMRFFLGO when NMRFFLGO is on for cancel.

JES3 Not used.

NMRFJID 20 2 Binary

This is the initial job number of the job to be processed.

JES2 Initialized to zeros. This field is set only if specified via $G operator command.

JES3 Not used.

NMRFORGN 22 8 Character blanks

This is the origin node for the job to be processed.

JES2 Set from $G operator command. Defaults to submitting node.

NMRFJNAM 2A 8 Character blanks

This is the job name of the job to be processed.

JES2 Set from $G operator command.

JES3 Used as job name or number on JES3 commands.

RSCS First four characters used as spoolid the command is to act upon.

NMRFD 32 8 Character blanks

This is the destination for the route command.

JES2 Set directly from $G operator command with no syntax checking.

RSCS Used for new node on TRANSFER command.

NMRFR 3A 8 Character blanks

This is the remote name for the route command if not implied by NMRFD.

JES2 Set directly from $G operator command with no syntax checking.

JES3 Not used.

RSCS Used for new user-id or remote name on TRANSFER command.

8.7.3 Messages

Messages are sent in NMRs and can be distinguished from other NMRs because they have the NMRFLAGC ("Command") bit off.

NMRFLAG 0 1 Bits 0

This byte defines the following flags. When the X'70' bits are off, NMROUT contains an MCS routing descriptor.

NMRFLAGC 0 .1 80 bit 0

This bit must be set to 0 to define a message as opposed to a command.

JES2 Messages are sent via $DM operator command or in response to a NMR command.

NMRFLAGW 0 .1 40 bit 0

This bit says NMROUT has the destination remote number.

RSCS Set if destination has a second level address that is not blank. This address could be a userid or a remote name.

NMRFLAGT 0 .1 20 bit 0

This bit says NMROUT has the destination userid.

JES3 Set for TSO notify messages. For input messages, indicates that the destination is a TSO userid.

RSCS Set if destination has a second level address that is not blank. This address could be a userid or a remote name.

NMRFLAGU 0 .1 10 bit 0

This bit says NMROUT has the destination console identifier.

JES3 For input messages, indicates that the destination is a JES3 console. Not set for output messages.

Note: JES3 sets the previous 3 bits off to indicate a message to a console.

RSCS Defaulted, not used.

POWER Defaulted, not used.

NMRFLAGR 0 .1 08 bit 0

Not used for messages.

NMRFLAGJ 0 .1 04 bit 0

Not used for messages.

NMRFLAGD 0 .1 02 bit 0

Not used for messages.

NMRFLAGS 0 .1 01 bit 0

Not used for messages.

NMRLEVEL 1 High 4 bits Binary 0

This is the importance level. See Figure 8-4 on page 8-64 for the bit definitions.

JES3 Not used. Set to 7.

RSCS Not used. Set to 7 for notify messages and 3 for command responses. Not forwarded in Release 3.

POWER Not used. Set to 7.

NMRPRIO 1 Low 4 bits Binary 0

This is the output priority.

JES3 Not used. Set to 7.

RSCS Not used. Set to 7. Not forwarded in Release 3.

POWER Not used. Set to 7.

NMRTYPE 2 1 Bits 0

This byte defines the following flags.

NMRTYPEX 2 .4 F0 bits 0

These are reserved bits.

NMRTYPE4 2 .1 08 bit 0

This bit says the first 8 bytes of NMRMSG contain the userid sending the message unless NMRTYPET is zero. If NMRTYPET is zero then NMRMSG+8 contains the userid. See Figure 8-2 on page 8-64 for more detail.

JES3 Not used or set.

NMRTYPET 2 .1 04 bit 0

This bit says the beginning of NMRMSG does not contain a time stamp. If this bit is set the actual text starts at NMRMSG or NMRMSG+8 depending on the setting of NMRTYPE4. See Figure 8-2 on page 8-64 for more detail.

JES3 Not used. Set for all output messages.

RSCS Always set, not used. Not forwarded.

POWER Always set for message transmission.

NMRTYPEF 2 .1 02 bit 0

This bit is not used for messages.

NMRTYPED 2 .1 01 bit 0

This is the Delete Operator Message (DOM) bit. DOM is a function of Multiple Console Support (MCS). This bit is not used or set by any system.

NMRML 3 1 Binary 1 1-132

This is the length of the information in NMRMSG. This count includes the userid if it is present. It does not include the 8-byte time stamp if it is present.

NMRTO 4 9 Mixed

This field is made up of the following two fields.

NMRTONOD 4 8 Character blanks

This is the destination node for the command.

NMRTOQUL C 1 Binary 0

This is the number of the member of the JES2 MAS node or the index to the JES3 system to which the message is being sent. This should be set from the input command or the qualifier in the Job Header (NJHGORGQ) for job related messages.

JES3 For output messages, set to X'00' or the system qualifier specified on the *MESSAGE command. For responses to commands, set from the NMRFMQUL of the command. For input messages destined for TSO (NMRFLAGT on), used as an index to the proper JES3 Local processor unless zero, which is an invalid value (in this case the message will be sent to the default NJE message class).

RSCS Not used, not set. Not forwarded in Release 3.

POWER Not used. Set from NMRFMQUL of input command, otherwise defaulted.

NMROUT D 8 Character zeros

This is the destination userid, the destination remote id, or the destination console id

depending on the setting of NMRFLAG. This field is further defined below. See Figure 8-3 on page 8-64 for how this field is used.

NMRUCM D 1 Binary 0

This contains the MCS console identifier when NMRFLAGU is set.

JES3 Not used or set.

RSCS Not used or set.

POWER Not used or set.

NMRUCMA E 1 Binary 0

This contains the MCS console area when NMRFLAGU is set.

JES3 Not used or set.

RSCS Not used or set.

POWER Not used or set.

NMRLINET F 2 Binary 0

This contains the line spacer for a multi-line WTO when NMRFLAGU is set.

JES3 Not used or set.

RSCS Not used or set.

POWER Not used or set.

NMRDESC D 2 Binary 0

This contains the MCS descriptor codes when NMRFLAGT and NMRFLAGU are not set.

JES3 Not used or set for messages.

RSCS Release 3 - If no destination userid is specified then it is set to X'0800'. Not forwarded.

Version 2.1 - Not used or set, but forwarded.

POWER If no destination userid is specified then it is set to X'0800'.

NMRROUT F 2 Binary 0

This contains the MCS routing codes when NMRFLAGT and NMRFLAGU are not set.

JES3 Not set for messages. For input messages, this field may contain a JES3 console number.

RSCS If no destination userid is specified then it is set to X'4100'. Not used. Not forwarded in Release 3.

POWER If no destination userid is specified then it is set to X'4100'.

NMRDOMID 11 4 Binary 0

This contains the DOM identifier when NMRTYPED is set.

JES3 Not set for messages. If zero on input, and NMROUT is not ALL zero, NMRROUT is assumed to contain a JES3 console number.

RSCS Not used or set.

POWER Not used or set.

NMRRMT D 8 Character

This contains the remote name when NMRFLAGW is set. JES2 uses the form 'Rnnn'.

RSCS Not set. Used as VM userid.

NMRUSER D 8 Character

This contains the destination userid when NMRFLAGT is set.

NMRFM 15 9 Mixed

This field is made up of the following two fields.

NMRFMNOD 15 8 Character blanks

This is the origin node.

NMRFMQUL 1D 1 Binary 0

This is the qualifier for the origin node. This is the number of the member of the JES2 MAS node or the index to the JES3 system from which the command is being sent.

JES3 Set to X'00' for output messages. Not used for input messages.

RSCS Not used or set. Not forwarded in Release 3.

POWER Not used or set.

NMRMSG 1E 132 Mixed

This is the actual data for the message. If both a time stamp and userid are present (NMRTYPET zero and NMRTYPE4 one) this field could be 148 bytes.

NMRECSID 1E 8 Character blanks

This contains the origin userid if NMRTYPE4 is set. Otherwise, this field does not exist. If this field exists, then the message text begins after this field. The Time Stamp is 8 bytes long.

JES3 Not used or set.

RSCS Used and set for messages originating from a VM user.

TYPE4	TYPET	MSG
0	0	Time stamp, Message text
0	1	Message text
1	0	Time stamp, Userid, Message text
1	1	Userid, Message text

Figure 8-2. NMRMSG format. This describes the data in NMRMSG depending on the setting of NMRTYPE4 and NMRTYPET.

FLAGW	FLAGT	FLAGU	OUT
0	0	0	MCS routing descripter
0	0	1	Console identifier
0	1	0	Userid
1	0	0	Remote number

Figure 8-3. NMRFLAG interpretation. This describes the data in NMROUT depending on the setting of NMRFLAG.

NMRLEVEL Bit Setting	Definition	JES2 Example
X'10'	Non-essential Messages	JCL errors in a job
X'30'	Normal Messages	Invalid password on a line
X'50'	Messages requiring delayed operator action	Job held
X'70'	Essential messages	Device start or drain
X'80'	Messages requiring immediate operator action[1]	Setup messages for printers
X'F0'	Extremely important messages	JES2 abends

Figure 8-4. NMRLEVEL interpretation. This describes some sample uses for the different values of the NMRLEVEL field.

[1] This message will not be deleted except by an operator action. It is also highlighted.

8.8 Connection and Control Records

These records are used to signon and signoff NJE connections. They are also used by the JES2 Network Path Manager to send information about other connections in the network.

8.8.1 Initial Signon Record

This record is sent by the primary NJE node to initiate a signon sequence. It is not compressed or compacted.

INITIAL SIGNON

0	NCCRCB	NCCSRCB I	NCCIDL	NCCINODE
4				
8				NCCIQUAL
C	NCCIEVNT			
10	NCCIREST		NCCIBFSZ	
14	NCCILPAS			
18				
1C	NCCINPAS			
20				
24	NCCIFLG	NCCIFEAT		
28				

NCCRCB 0 1 Binary X'F0'

All Path Manager RCBs are X'F0'.

NCCSRCB 1 1 Character

The particular type of Path Manager record is determined by this field. For initial signon a C'I' is used by the primary.

NCCIDL 2 1 Binary

Length of record from RCB to end of extension (extension currently is void)

NCCINODE 3 8 Character

EBCDIC name of input node left justified (low end is determined by comparing the name of the node)

NCCIQUAL B 1 Binary 0

Member number of the node. JES2 uses 1-7 as member numbers. If only one system is to

communicate with the network this value should be 1 for that system. A value of zero is used for console services to indicate any system

NCCIEVNT C 4 Binary 0

Event sequence number. This field is not used for C'I' and is set to zero.

NCCIREST 10 2 Binary 0

Partial node to node resistance (ignored for pre-defined connections). Resistance values between the 2 connecting systems are added to make up a total resistance

NCCIBFSZ 12 2 Binary

Maximum size transmission record that the sending system can receive. Must be 300 bytes or more.

NCCILPAS 14 8 Character Blanks

Line password.

NCCINPAS 1C 8 Character Blanks

Node password.

NCCIFLG 24 1 Binary 0

Flag byte, zero for initial sign on.

NCCIFEAT 25 4 Binary 0

This defines space for up to 32 features.

Note:

1. Sender does not necessarily know what node will receive the record.
2. Initial signon is only record in transmission.
3. Not compressed nor compacted for transmission.
4. Multiple initial signon records are ignored if the connection is already active.

8.8.2 Response Signon Record

This record is sent in response to an initial signon record.

RESPONSE SIGNON

0	NCCRCB	NCCSRCB J	NCCIDL	NCCINODE
4				
8				NCCIQUAL
C	NCCIEVNT			
10	NCCIREST		NCCIBFSZ	
14	NCCILPAS			
18				
1C	NCCINPAS			
20				
24	NCCIFLG	NCCIFEAT		
28				

NCCRCB 0 1 Binary X'F0'

All Path Manager RCBs are X'F0'.

NCCSRCB 1 1 Character

Response signon SRCB is C'J'.

NCCIDL 2 1 Binary

Length of record from RCB to end of extension (extension currently is void).

NCCINODE 3 8 Character

EBCDIC name of the responding node left justified (low end/high end is determined by comparing the name of the node).

NCCIQUAL B 1 Binary

Member number of the node.

NCCIEVNT C 4 Binary

Normal signon sequence is 0 if high end sends or secondary trunk of a multiple trunk connection. Next higher sequence if low end sends and the trunk is primary. Predefined sequence is X'FFFFFFFF'

NCCIREST 10 2 Binary 0

Partial node to node resistance.

NCCIBFSZ 12 2 Binary

Maximum size transmission record that the receiving system has agreed to.

See B.2.8.8, "Signon Deviations" on page B-21 for JES3 deviations.

NCCILPAS 14 8 Character Blanks

Line password

NCCINPAS 1C 8 Character Blanks

Node password

NCCIFLG 24 1 Bits 0

NCCIFLGM 24 .1 80 bit

On for multi-trunk response

NCCIFEAT 25 4 Binary 0

This defines space for up to 32 new features.

Note:

1. Response signon is only record in transmission.
2. Not compressed nor compacted for transmission.

8.8.3 Reset Signon Record

RESET SIGNON

0	NCCRCB	NCCSRCB K	NCCCDL	NCCCEVNT
4				NCCCREST
8				

NCCRCB 0 1 Binary X'F0'

All Path Manager RCBs are X'F0'.

NCCSRCB 1 1 Character

Reset signon SRCB is C'K'.

NCCCDL 2 1 Binary

Length of record from RCB to end of extension (extension currently is void).

NCCCEVNT 3 4 Binary

0 if high end or secondary trunk of multi-trunk connection. (Reset not transmitted from low end on normal multi-trunk signon sequences.)

NCCCREST 7 2 Binary

Partial node to node resistance.

Note:

1. Reset is used to increment CES and change resistance and is blocked with other control records (first in record).
2. Stepping through the block of records should be done using the length field.
3. Not compressed nor compacted for transmission.
4. A reset signon is scheduled if the trunk being initiated is the primary. It is used in place of a J record if changes are necessary.

8.8.4 Concurrence Signon Record

CONCURRENCE SIGNON

0	NCCRCB	NCCSRCB L	NCCCDL	NCCCEVNT
4				NCCCREST
8				

NCCRCB 0 1 Binary X'F0'

All Path Manager RCBs are X'F0'.

NCCSRCB 1 1 Character

Concurrence signon SRCB is C'L'.

NCCCDL 2 1 Binary

Length of record from RCB to end of extension (extension currently is void).

NCCCEVNT 3 4 Binary

Whatever the reset or response signon indicated.

NCCCREST 7 2 Binary

Total node-to-node resistance.

Note:

1. Low end cannot concur if the trunk is primary.
2. Concurrence is used to acknowledge a Response or Reset Signon record when required and is blocked with other control records.
3. Stepping through the block should be done using the length field.
4. Not compressed nor compacted for transmission.

8.8.5 Add/Subtract Connection Record

ADD/SUBTRACT CONNECTION

0	NCCRCB	NCCSRCB M/N	NCCADL	NCCANODA
4				
8				NCCAQULA
C	NCCANODB			
10				
14	NCCAQULB	NCCAEVNT		
18		NCCAREST		

NCCRCB 0 1 Binary X'F0'

All Path Manager RCBs are X'F0'.

NCCSRCB 1 1 Character

Add connection SRCB is C'M'. Subtract connection SRCB is C'N'.

NCCADL 2 1 Binary

Length of record from RCB to end of extension.

NCCANODA 3 8 Character

Low end node name.

NCCAQULA B 1 Binary

Low end member number (if shared spool).

NCCANODB C 8 Character

High end node name.

NCCAQULB 14 1 Binary

High end member number (if shared spool).

NCCAEVNT 15 4 Binary

Connection Event Sequence.

NCCAREST 19 2 Binary

Resistance (total).

Note:

1. Add and Subtract are used to notify members of a network not directly involved in the connection about the connection status between named nodes/members and are blocked with other control records.

2. Members of MAS have same node name and are distinguished by different member names.

3. If the CES of an add connection is not greater than previously received Add or Subtract Connection, the record should be ignored, (otherwise path manager tables could be down leveled).

4. If the CES of a Subtract Connection is less than a previously received Add or Subtract Connection, the record should be ignored and if the CES is equal to a connection that is already unconnected, it should be ignored, because it has already been received (perhaps by another node on another path).

5. Step through a block of records using the length field.

6. Not compressed nor compacted for transmission.

8.8.6 Signoff Record

SIGNOFF

0	NCCRCB	NCCSRCB B

NCCRCB 0 1 Binary X'F0'

All Path Manager RCBs are X'F0'.

NCCSRCB 1 1 Character C'B'

Signoff SRCB is C'B'.

NOTE

JES2 JES2 only sends a signoff record on BSC and CTCA sessions, not on SNA sessions. JES2 never examines a signoff record.

JES3 JES3 1.3.4 only sends a signoff record on BSC sessions, not on CTCA sessions. JES3 2.1.5 sends a signoff record on both BSC and CTC sessions.

RSCS RSCS sends a signoff record on BSC and CTCA sessions, and uses them.

POWER POWER only sends a signoff record on BSC sessions, not on SNA sessions, but does accept signoff records on an SNA session.

Appendix A. Line Protocols

A.1 Bisynchronous Communications Lines

A.1.1 Initialization

The channel commands used in BSC initialization are given in Figure A-1.

OPCODE	DEFINITION	COMMENTS
X'2F'	Disable	This CCW will cause the adapter to be reset. If this is a switched communications line, the circuit will be terminated.
X'23'	Set Mode	Conditions the adapter. For NJE a byte of X'00' should be used.
X'27'	Enable	Causes a switched line to wait for a connection to be made. On a non-switched line this will immediately terminate with good ending status if the communications equipment is operational. If the line is not operational, a unit check with sense data of timeout or intervention required will occur. It also initializes the adapter to accept read and write commands.
X'01'	Write	Writes data to the adapter.
X'02'	Read	Read data from adapter.

Figure A-1. BSC Channel Commands

BSC connections are established by an exchange of NJE SIGNON records. There is an initiate SIGNON, sent by the primary node followed by a response SIGNON, sent by the secondary node. Before the SIGNON records can be exchanged, the two nodes must agree on which one will be primary and which will be secondary. This is done with the SOH ENQ and DLE ACK exchange. The side which first succesfully sends an SOH ENQ to the other node becomes the primary. The side which first succesfully receives the SOH ENQ becomes the secondary and responds with a DLE ACK. The SOH ENQ exchange followed by the SIGNON exchange is shown in Figure A-2 on page A-2.

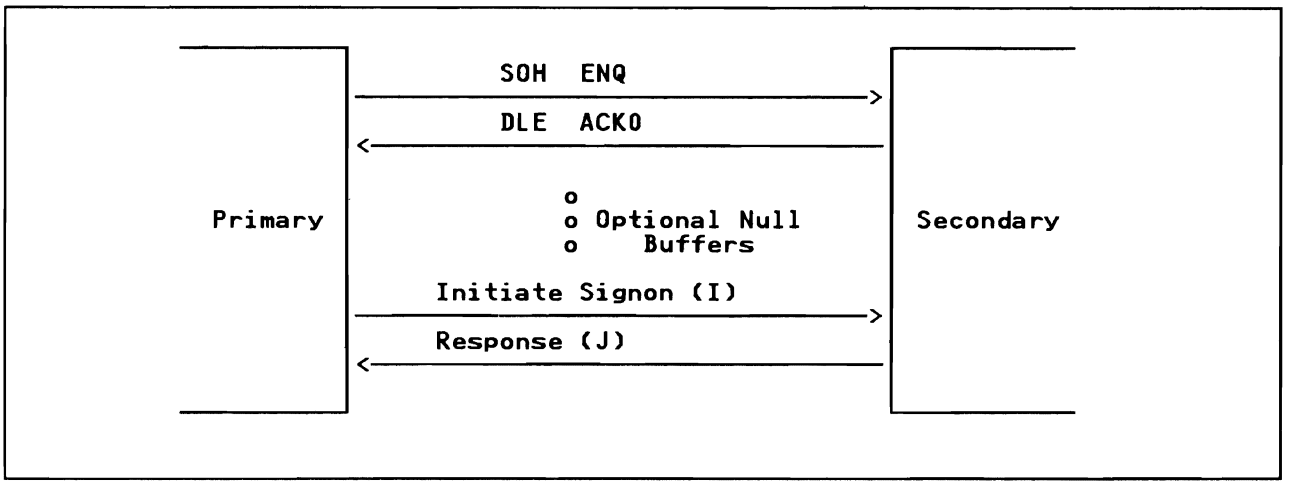

Figure A-2. BSC Initialization

There are two ways that the primary/secondary relationship can be established. The roles may be predefined at each node by the two operators or system programmers such that one end will be secondary and the other primary whenever connection establishment is attempted. In this case the secondary always begins by reading from the line while the primary begins by writing an SOH ENQ to the line.

Note: Only JES2 and RSCS Version 2.1 provide the facility to be pre-defined as a secondary node. JES3, RSCS Version 1 and VSE/POWER can only function as primary (send SOH ENQ) initially.

The other way that the roles may be established is for both nodes to begin as primary, both sending SOH ENQ. The contention for primary is resolved by the collision of SOH ENQ's in one of the two directly connected communication controllers (CC).

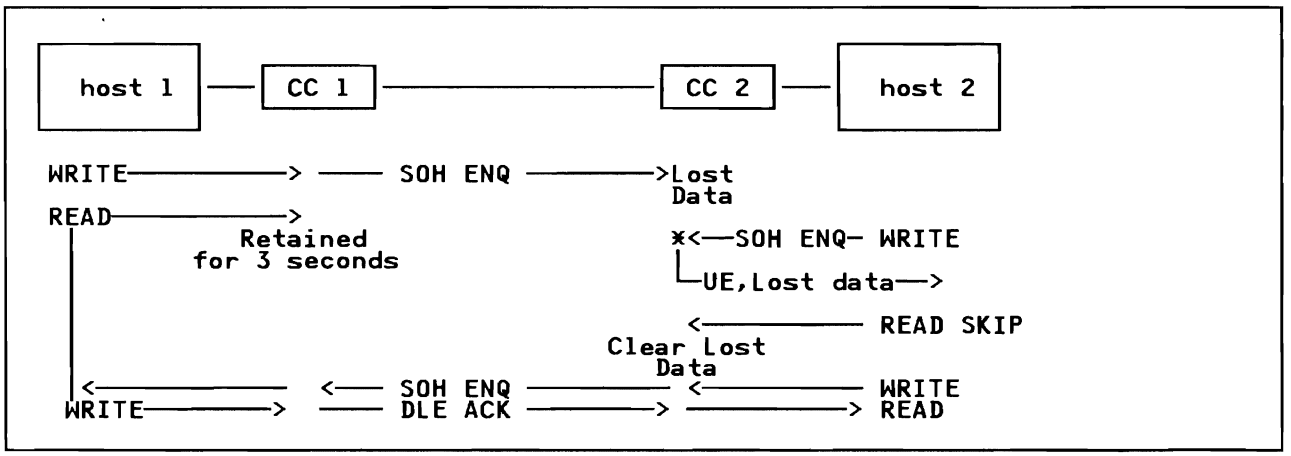

Figure A-3. BSC Signon Contention

If inbound data arrives at a communication controller (CC) and a READ is not up, a lost data condition is set in the controller. If a WRITE is subsequently issued, the CC responds with Unit Exception, Lost Data. The Lost Data condition remains and can only be cleared with a READ command. READ SKIP is issued to clear the Lost Data condition and another WRITE SOH ENQ is issued. At side 1 the READ stays up for 3 seconds and is waiting for the SOH ENQ to come in. Note that the side which first sends an SOH ENQ to the other controller and causes the collision becomes the secondary.

With a satellite connection between the two communication controllers there is no longer a direct connection between them and it is possible for the two SOH ENQ to "pass in the night". Without the collision detection in one of the communication controllers both SOH ENQ's get through and both nodes believe they are secondary. Each waits for an initial SIGNON. Because of this possibility, the primary/secondary roles must be predefined.

A.1.2 Initialization Error Recovery

Initialization error recovery is in effect until either the DLE ACK0 is received properly from the other side, or an SOH ENQ is properly received and a DLE ACK0 is transmitted.
During initialization, errors are responded to in one of three ways:

1. Terminate the line.
2. Retransmit the SOH ENQ.
3. Issue a read with a large CCW count and retransmit the SOH ENQ.

An ending condition of unit exception on the write CCW occurs when data has been transmitted by the other side and a read was not active to receive the data. In this case, a read CCW with a large count and the skip bit should be issued to clear any data that is pending in the adapter. Then the write should be reissued.

An ending condition of unit check with sense data of either command reject or intervention required should cause the line to be terminated. If a command reject occurred, either the adapter was not enabled, or an invalid CCW was executed. In either case a software error has probably occurred and retrying would most likely result in a loop. An intervention required can occur for many reasons. The problems can range from the telephone hanging up on a switched line, to an intermittent clock within the local modem. In any case, retrying the write will usually result in the intervention required happening again. If a unit check with sense data of time out, data check, or lost data is received, the SOH ENQ should be retransmitted. If good status is received, but the data is not SOH ENQ or DLE ACK0, the SOH ENQ should be retransmitted. These conditions are summarized in Figure A-4.

Unrecognized data		Rewrite SOH ENQ
DEVICE STATUS	CCW	RECOVERY ACTION
Unit Exception	Write	Read Skip - Rewrite SOH ENQ
	Read	Rewrite SOH ENQ
DEVICE STATUS	SENSE DATA	RECOVERY ACTION
Unit Check	80 Command reject	Terminate
Unit Check	40 Intervention required	Terminate
Unit Check	20 Bus out check	Terminate
Unit Check	10 Equipment check	Terminate
Unit Check	08 Data check	Rewrite SOH ENQ
Unit Check	04 Overrun	Rewrite SOH ENQ
Unit Check	02 Lost data	Rewrite SOH ENQ
Unit Check	01 Time out	Rewrite SOH ENQ

Figure A-4. BSC Initialization Error Recovery

See B.3.8.3, "Error Handling" on page B-34, "Initialization Error Recovery" on page B-46, and "BSC Initialization Error Recovery" on page B-19 for system specific differences.

A.1.3 Error Recovery

Error recovery after initialization is more complex than error recovery during initialization. Those errors that require the line to be terminated do not change. Other errors require a negative acknowledgment (NAK (X'3D')) to be transmitted. The NAK notifies the other side that either a block was received incorrectly, or no block was received within the three seconds required by the BSC adapter. Reception of a NAK is also considered an error and requires the last transmission that was not a NAK transmission to be resent. The error conditions and recovery actions are shown in Figure A-5.

SOH ENQ		Restart line
NAK		Retransmit last non-NAK transmission
Unrecognized data		Write NAK
DEVICE STATUS	**CCW**	**RECOVERY ACTION**
Unit Exception	Write	Read Skip - Write NAK
	Read	Write NAK
DEVICE STATUS	**SENSE DATA**	**RECOVERY ACTION**
Unit Check	80 Command reject	Terminate
Unit Check	40 Intervention required	Terminate
Unit Check	20 Bus out check	Terminate
Unit Check	10 Equipment check	Terminate
Unit Check	08 Data check	Write NAK
Unit Check	04 Overrun	Write NAK
Unit Check	02 Lost data	Write NAK
Unit Check	01 Time out	Write NAK

Figure A-5. BSC Error Recovery

See "Initialization Error Recovery" on page B-34 and "BSC Error Recovery" on page B-19 for system specific differences.

A.1.4 Termination

Normal termination occurs when a signoff record (see 2.3, "Signoff (All Systems)" on page 2-7) is sent. There is no special BSC character sequence sent for termination. No response is expected to the signoff record. If a read is chained to the write of the signoff, the read will timeout. After the signoff record is written, a disable should be issued to disable the communications adapter.

A.1.5 Normal Sequences

DLE ACK0 (X'1070') is a positive acknowledgment to the previous block. At the line protocol level, a DLE STX (X'1002') is also a positive acknowledgment to the previous block. The DLE ETB that signals the end of the transmission block must be sent by command chaining a separate write CCW for those two bytes of data to the write CCW used to transmit the transmission block. Otherwise, a command reject will occur.

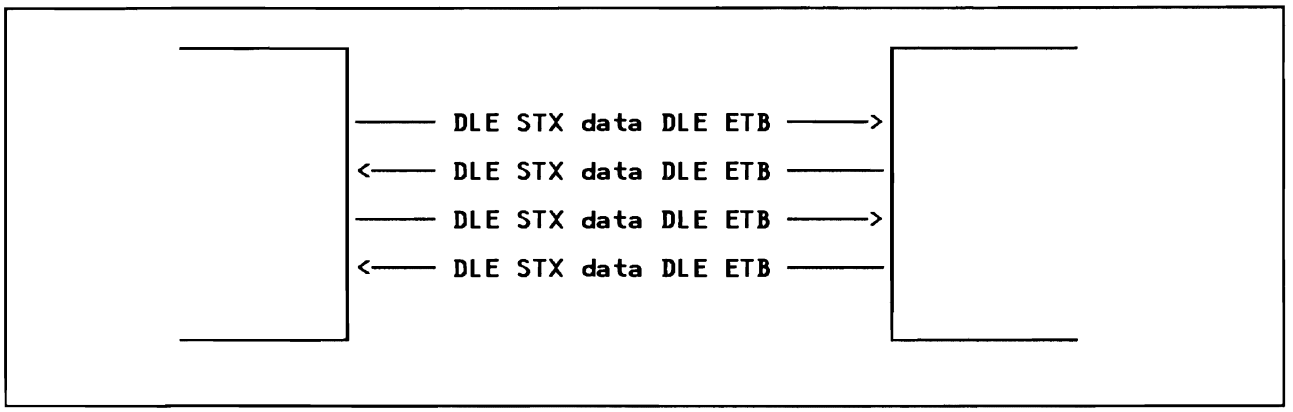

Figure A-6. BSC Normal Communication

A.1.6 V.27 Modem Contention Resolution

Certain modems used in support of CCITT V.27 protocols for BSC communications require special SIGNON contention resolution protocols.

These protocols require that, external to the NJE protocols, there be a way to identify to the program supporting this new protocol, the "mode" of the node that is about to participate in a connection. There are three modes:

- Primary
- Secondary
- Contention

A given node, for V.27 SIGNON, must be identified relative to another specific node as either "Primary" or "Secondary." The other node must have the opposite designation. This information need not be available to the programs supporting the protocols except as a parameter on invocation. Hence, two system operators may establish an arbitrary relationship prior to session initialization, or the product may elect to maintain tables.

A.1.6.1 V.27 SIGNON by the "Primary" Node

This scenario assumes that the node attempting SIGNON has been identified as the "Primary" node.

```
START:

      Allow operator intervention to stop this
      process, otherwise continue until 'DLE ACK0'
      received.

      Issue  WRITE (cc)   SOH ENQ
             READ

      If ANY SERIOUS ERROR (e.g. device not avail)
            -Terminate immediately

      If UNIT EXCEPTION (contention)

            -Issue Message
            -Terminate  (contention not allowed)

      If TIMEOUT

            -Retry immediately the same CCW chain
             up to 5 times
            -If retry works, goto step CKREAD
            -Otherwise set timer for 2 minutes
            -When timer expires, retry the original CCW
             sequence up to 5 times

CKREAD: (Successful I/O completion)

      See if data read is 'DLE ACK0'

      If so, SEND 'I' SIGNON record

       ... normal protocols follow

      If not 'DLE ACK0', go back to START
```

Figure A-7. V.27 Modem Protocols: Primary Node SIGNON

A.1.6.2 V.27 SIGNON by the "Secondary" Node

This scenario assumes that the node attempting SIGNON has been identified as the "Secondary" node.

```
START:

      Allow operator intervention to stop this
      process, otherwise continue until 'SOH ENQ'
      received.

      Issue  PREPARE (cc)
             READ

      If ANY SERIOUS ERROR (e.g. device not avail)
            -Terminate immediately

CKREAD: (Successful I/O completion)

      See if data read is 'SOH ENQ'

      If so, SEND 'DLE ACK0'
             READ

       ... normal protocols follow

      If not 'SOH ENQ', go back to START
```

Figure A-8. V.27 Modem Protocols: Secondary Node SIGNON

A.2 Channel to Channel Adapter

NJE also supports communication by a Channel-to-Channel Adapter (CTC) as a communication line. The support is for 360 mode operations over Block Multiplexor Channels. Currently, POWER is the only NJE product that does not support CTCs because there is no such support in the basic VSE control program.

To properly operate as a NJE device, the UCWs for the CTCA should be plugged nonshared. This is a hardware option on block multiplexor channels. The UCW is used to store the address for performing the channel reconnect. If the UCW for the CTCA is not plugged nonshared, the CTCA will not disconnect while waiting for the control CCW to complete. This will cause the channel to remain busy for a significant length of time during each I/O operation.

A.2.1 Initialization

To begin communication, a sense device command (X'14') is issued by Side "A" to determine the state of Side "B" of the adapter. If a byte of hex zero (X'00') is returned, it indicates that Side "B" has not yet been activated and no command is active. The correct response for Side "A" is to issue a control command (X'07') chained to a read (X'02'). This serves two purposes. First, the control signals an attention to Side "B". Second, the control state remains in Side "A" so that when Side "B" does subsequently start, it will detect that Side "A" is already active. The control CCW causes a channel disconnect on a block multiplexor channel, so the channel is available for use with other devices while waiting for Side "B" of the CTCA to respond.

When Side "B" starts, it issues a sense device command, detects the control state (X'07') in Side "A", and writes a SYN NAK (X'323D'). This allows Side "A" to become the primary. The primary then responds with a SOH ENQ (X'012D'). If Side "B" wishes to be the primary rather than the secondary, it must send a SOH ENQ rather than a SYN NAK. The secondary then responds with a DLE ACK0 (X'1070') as with a BSC line. See B.2.8.1, "CTCA Initialization" on page B-19 for JES3 specific processing. See B.3.8.1, "CTCA Initialization" on page B-33 for RSCS specific processing.

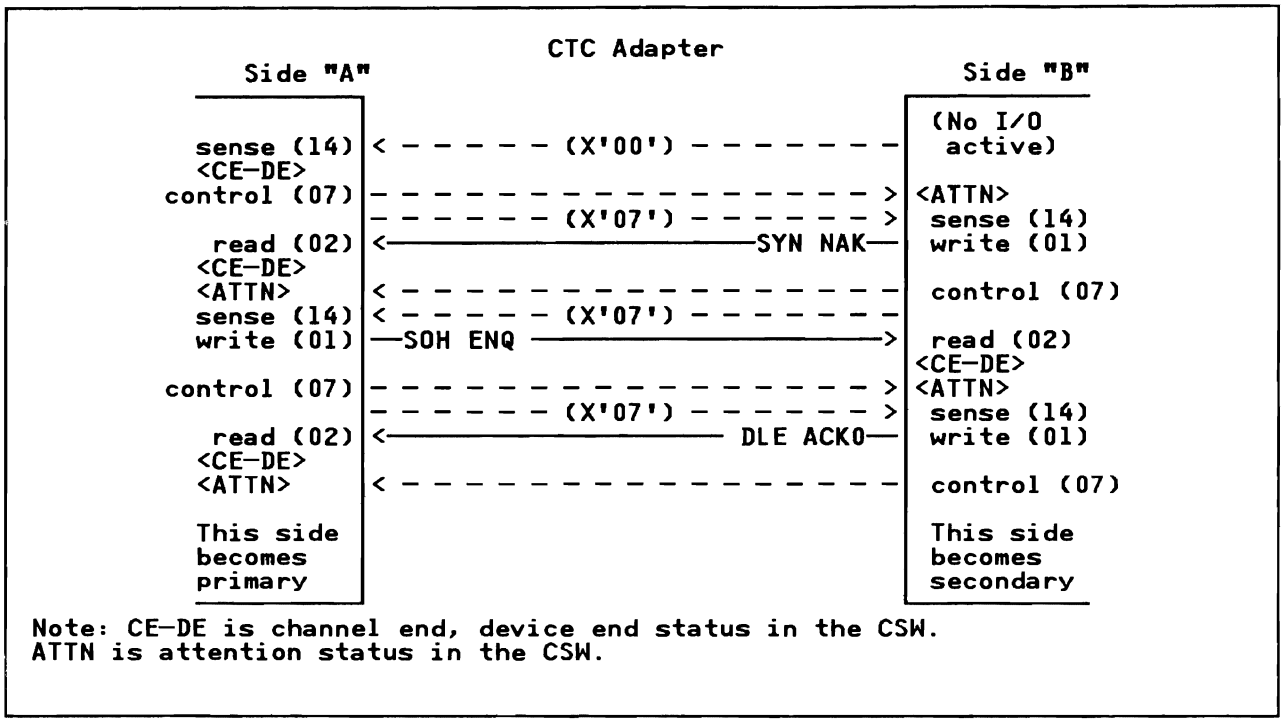

Figure A-9. CTCA Initialization

A.2.2 Error Recovery

No error recovery is used on the CTCA. The CTCA, when running in compatibility mode, cannot generate any status bits other than attention, busy, channel end, or device end. If any channel errors occur it is considered proper to signal connection termination to the higher levels. The session must then be restarted by normal initialization protocol.

Note: Prior to the advent of channels that buffer status, attention and busy could only occur if the two systems had lost synchronization. A nonshared 30xx subchannel may cause an attention or attention and busy status. This is because the subchannel has already buffered the attention from the other sides I/O request and intends to present it. In channels that do not buffer status, the I/O would have been sent to the CTCA and the CTCA would have suppressed the attention status and accepted the command. The only option in this case is to reissue the I/O. On newer, buffered channels, a condition exists where the attention busy status is caused by having another device cause contention. Therefore, the attention busy may occur in normal operation and should not be a termination condition. See B.1.8.2, "CTCA Attention and Busy Status" on page B-10 and B.2.8.3, "CTCA Attention and Busy Status" on page B-19 for system-specific handling of this condition

A.2.3 Termination

Normal termination occurs when a signoff record (see 2.3, "Signoff (All Systems)" on page 2-7) is sent. There is no special character sequence sent for termination. No response is expected after a signoff record is sent. In this case, the standard CTCA channel program should terminate with the write CCW.

See B.2.8.2, "CTCA Termination Deviations" on page B-19 for JES3 deviations.

A.2.4 Normal Sequences

Instead of using CCWs designed for BSC line transmissions the Channel-to-Channel Adapter sequences are as those shown in Figure A-10.

OPCODE	DEFINITION	COMMENTS
X'14'	Sense	This CCW will cause other end of adapter control CCW to fall through to the read CCW.
X'01'	Write	Writes data to the adapter.
X'07'	Control	Causes the channel program to wait until other end of adapter issues sense. A block multiplex channel is freed for other programs.
X'02'	Read	Read data from adapter.

Figure A-10. CTCA Channel Program

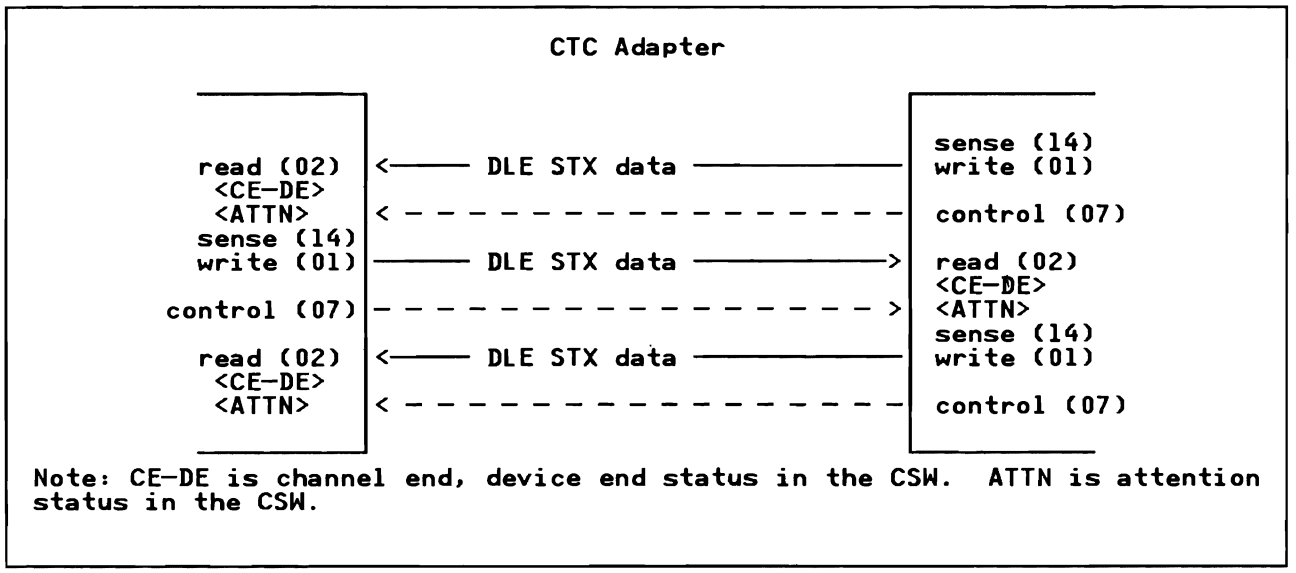

Note: CE-DE is channel end, device end status in the CSW. ATTN is attention status in the CSW.

Figure A-11. CTCA Normal Sequences

A.3 PREPARE Mode Option: Suspend I/O (BSC and CTCA)

This option defines a method by which non-SNA links do not have to transmit to each other every one to two seconds. It is called PREPARE Mode because it uses the PREPARE CCW on a BSC line. For the CTCA, the link is left without any I/O active during periods of inactivity.

Use of PREPARE mode is controlled by a signon concurrence bit (see 2.1.1, "Signon Concurrence Feature" on page 2-2 for the description of signon concurrence). This causes no problems for back level systems, since this has already handled by the signon concurrence protocols. The result is that PREPARE will never be used when communicating with back level systems.

PREPARE could cause idle lines to appear connected when in fact the system at the other end of the line is no longer operational. This problem may be avoided by setting a timer during I/O suspension and attempting to wake up the other side after a fixed interval has passed. If the other side does not respond, the link must be taken down using normal error recovery procedures.

For PREPARE to work properly, both systems using it must send null buffers rather than DLE ACK0 when they have no data to transmit. Signon concurrence bit X'80' controls the use of PREPARE.

This requires no control block changes, but does rely on the Signon Concurrence Feature of NJE. The X'80' bit in the first byte has been assigned to PREPARE.

A.3.1 Normal Sequences

Presently, when no data is being transmitted or received on a given link, a null-buffer is transmitted every one to two seconds. This is shown for BSC in Figure A-12. (The CTCA normal CCW chain is similar except that the read CCW is preceded by a control and the write by a sense.)

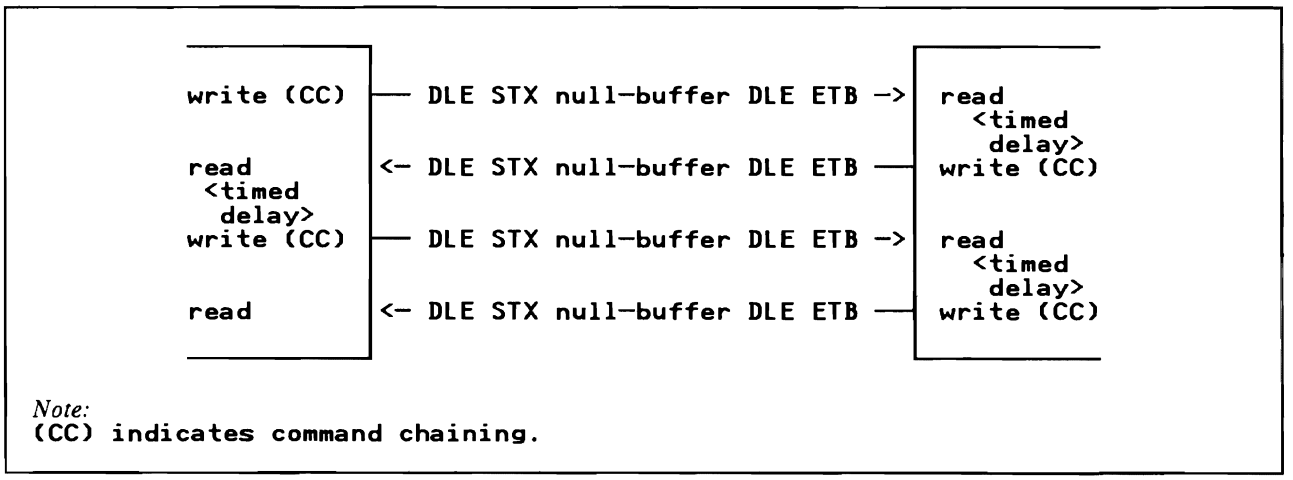

Figure A-12. BSC Normal Communication

A.3.2 Requesting I/O Suspension

When the last non-null-buffer is acknowledged and no files are in the process of being received or sent, one side (called A below) may decide to suspend I/O activity. It does this by writing a null-buffer to the other side (called B) with the terminating sequence 'DLE ETX' instead of the normal 'DLE ETB'. This buffer is sent with the appropriate I/O CCW sequence for CTCA or BSC. Thus, the write is command chained either directly (BSC) or indirectly (through a control for CTCA) to a read. If B wants to accept the I/O suspension, it responds with another null-buffer also ended by 'DLE ETX' instead of 'DLE ETB'. Unlike the normal I/O sequence, this write CCW is not chained directly to a read (BSC) or to a control, read (CTCA). If it does not want to suspend I/O, B sends either a non-null data-buffer or a null-buffer terminated with 'DLE ETB' with its normal I/O sequence. In this case, normal I/O sequences continue to be used by both sides.

As soon as side B sends the acknowledging null-buffer, it can suspend I/O. Side A must also suspend I/O if it receives an acknowledging null-buffer (with DLE ETX) immediately following the buffer it used to request suspension.

Note: That the two null-buffers with DLE ETX must be exchanged without an intervening buffer with DLE ETB in order for PREPARE to be used. If there is an intervening normal buffer, the second DLE ETX buffer will be considered a new request rather than an acknowledgment.

A.3.3 Suspension and Resumption of I/O

How communication is actually suspended and resumed with PREPARE mode differs for BSC and CTCA.

A.3.3.1 BSC I/O Suspension

For BSC suspension, both sides issue a read CCW that is allowed to time out followed by the I/O sequence of a PREPARE CCW (X'06') command chained to a read CCW.

Both systems then wait until either one of them has something to send. When either of the systems has data to transmit to the other system, it must then issue a HDV (or HIO) to terminate the PREPARE. Then before transmitting the actual buffer, it should transmit a DLE ENQ sequence. When a system waiting on a PREPARE receives any data, it must acknowledge the data with a null-buffer before going back to wait on the PREPARE. This allows the system with the pending data to transmit the data and once data has been received, normal processing resumes. This sequence is shown in Figure A-13 on page A-13. If both systems attempt to initiate transmission at the same time, one of the DLE ENQ's may be lost by the HDV, or may be lost when a read-skip is used to recover from a unit exception on the write. In this case, a contention resolution protocol is used (see A.3.4, "BSC Error Protocols" on page A-17) that forces resynchronization.

Note: During resume all null-buffers are terminated with the normal DLE ETB sequence.

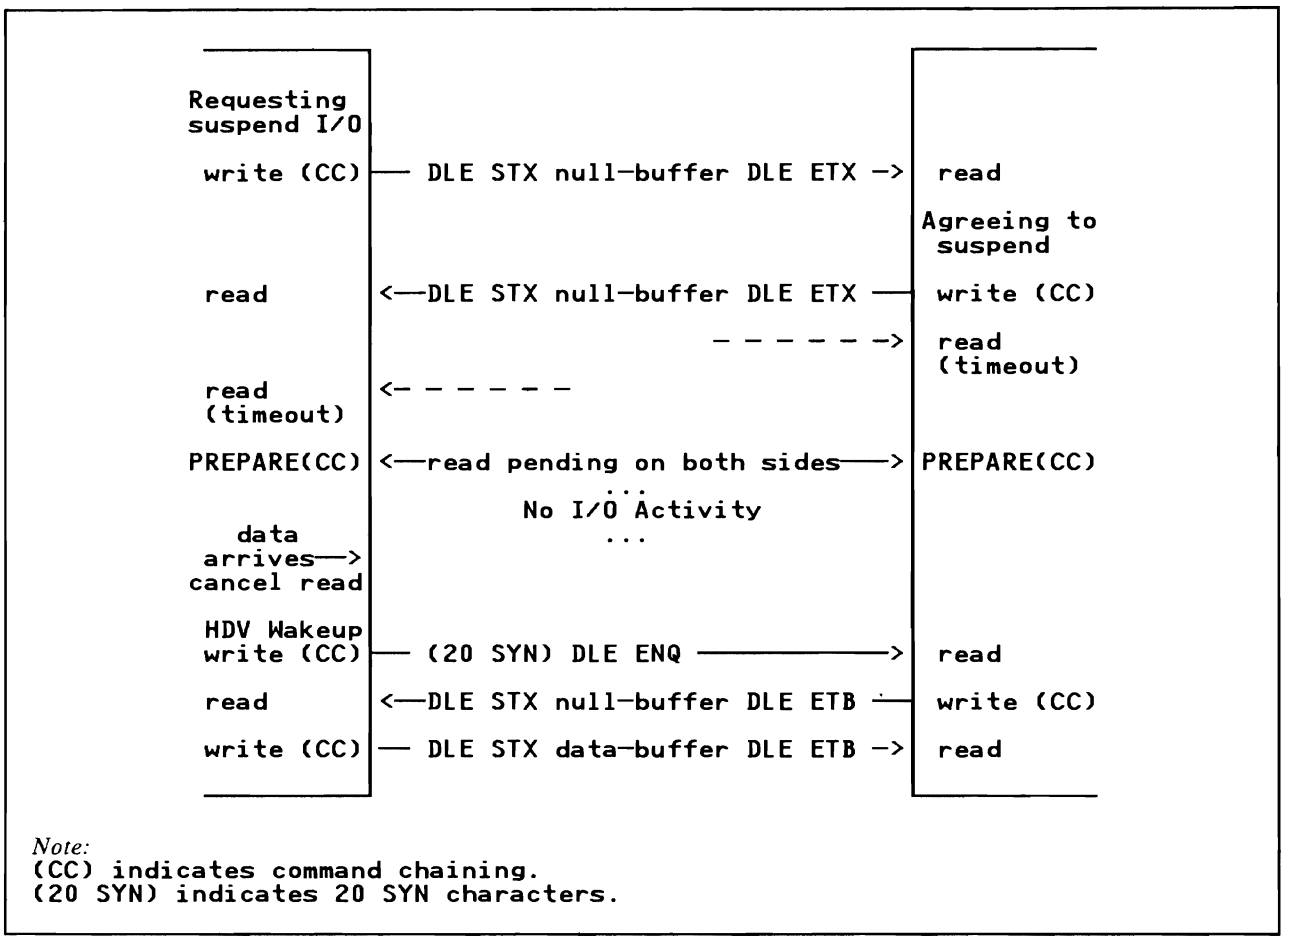

Figure A-13. BSC Communication in PREPARE Mode

A.3.3.2 CTCA I/O Suspension

On a CTCA during I/O suspension, no I/O sequences are issued so that the adapter is left with no I/O in progress during idle periods. A control CCW which signals attention indicates data transmission is to be resumed. The procedure which is illustrated in Figure A-14 on page A-16 is as follows:

- I/O suspension is requested and agreed upon as described in A.3.2, "Requesting I/O Suspension" on page A-12.
- Both sides then wait without any I/O outstanding.
- If either side wants to resume, it (A) must first issue a control CCW which will signal an attention to the other side (B). The control will be command chained to a read.
- When the attention interrupt is detected by B, it will issue its normal CTCA CCW chain (sense, write (null-buffer), control, read).
- The sense causes the control issued by A to complete, and the null-buffer is written by B and read by A.
- Side A which had data to send can send the data by issuing the normal CCW chain as soon as the first "wakeup" read completes.

- If both sides attempt to "wake up" at the same time, one of the control CCWs will get a busy response. When this happens, the normal CCW chain should be issued with the write for a null-buffer. Normal synchronized operation will then commence.

- If one side has gone away during the PREPARE, the "wakeup" control CCW will not complete and a selector channel will be tied up and unavailable for use by other devices. To prevent this, a timer could be set (for less than 30 seconds) at the beginning of I/O resumption. If the timer expired without the I/O completing, HDV would be issued and the link taken down.

The following flows show the procedure graphically. In the diagrams:

S = Sense CCW
C = Control CCW
R = Read CCW
W = Write CCS

Normal Sequence (for CTCA)

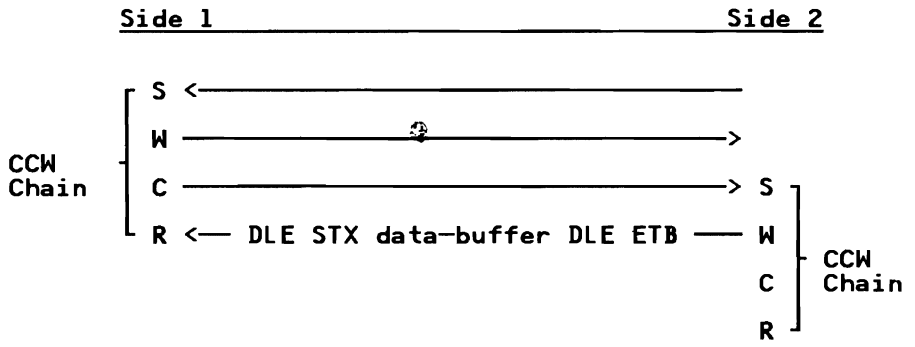

PREPARE Initiation (for CTCA)

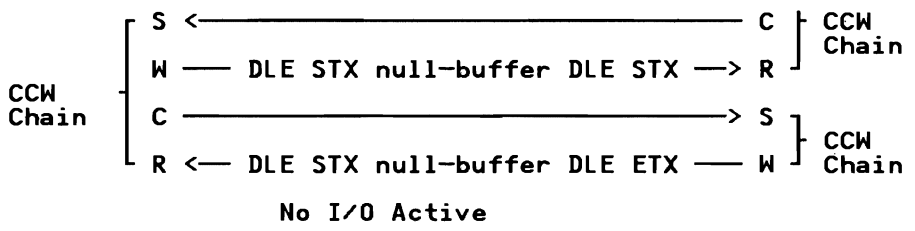

I/O Resumption

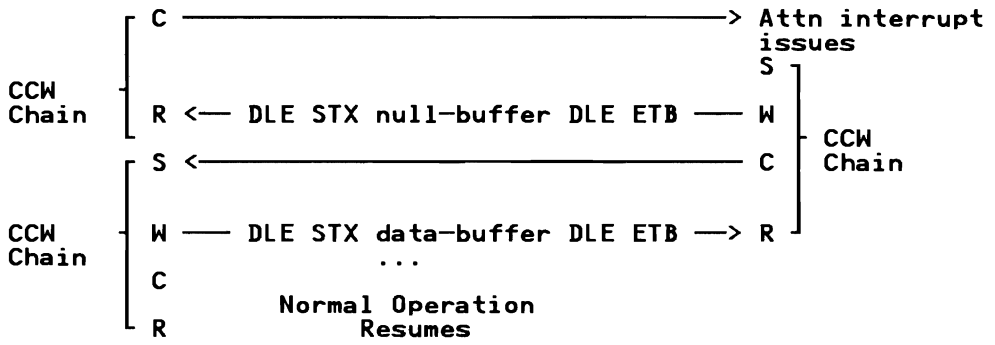

A.3.3.3 Protocols Common to BSC and CTCA

During the I/O suspension period, either side may optionally set a timer for some reasonable interval (i.e. 5 to 15 minutes) and when the time expires, transmit a null-buffer to determine if the other system is still operational. If the other side does not respond at this time, the link is taken down by normal error recovery procedures.

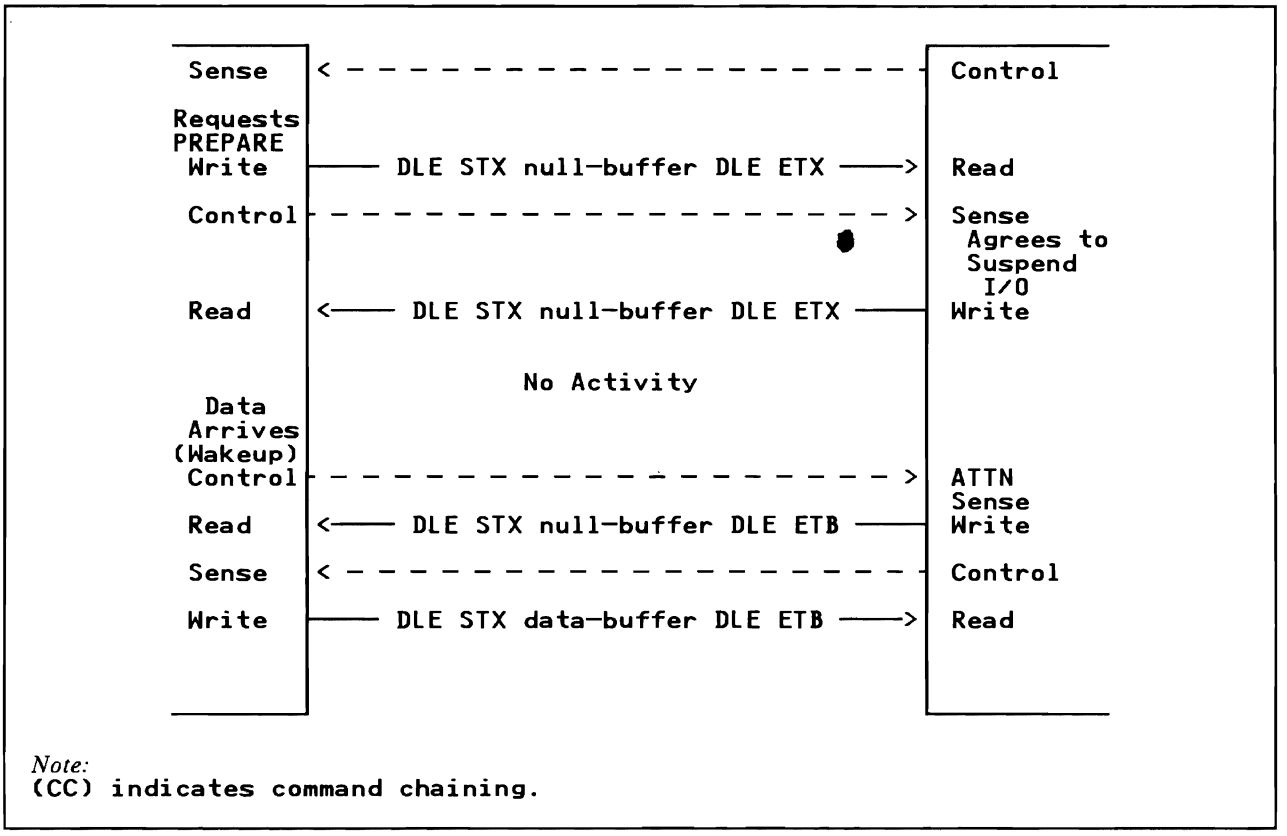

Figure A-14. CTCA Communication in PREPARE Mode

A.3.4 BSC Error Protocols

Symbols in flows below are:

R = Read
P = PREPARE
W = Write
S = Sense
C = Control
TO = Time Out
cc = Command Chain

1. When a DLE ENQ for PREPARE resume is NAKed, the condition is treated as a NAK to a NAK in normal processing. In this case, the last non-NAK is defined to be the last buffer (null or not) for which an acknowledgement has not been received. This buffer will be different depending on which side starts the PREPARE and which receives the NAK. The following example illustrates the different interpretations.

   Note: In the example, an actual buffer is sent in all cases where just the BCB is shown.

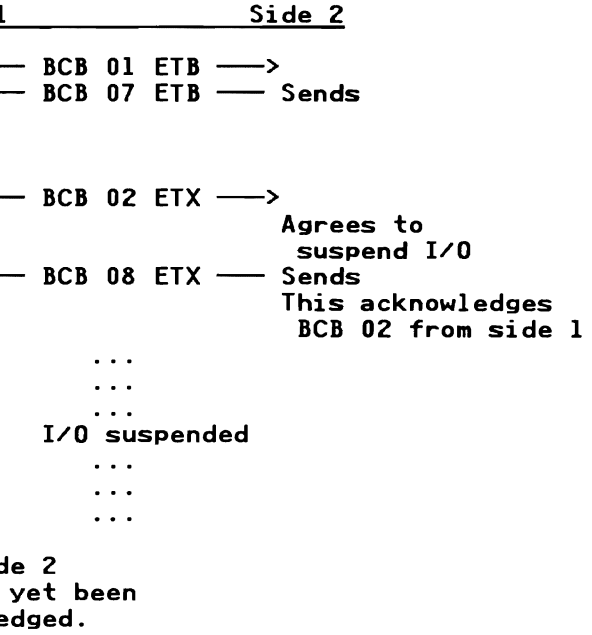

   Now assume:

```
   Side 1                Side 2

a. Side 1 resumes

 Sends ── DLE ENQ ──────>
      <── NAK ──────────
 (last ── BCB 03 ETB ──>
non-NAK)
      <── BCB 09 ETB ──  Other side's BCB 03
                                 acknowledges BCB 08
b. Side 2 resumes

      <── DLE ENQ ──────  Sends
       ─── NAK ─────────>
      <── BCB 08 ETB ── (last non-NAK)
       ── BCB 03 ETB ──>
```

2. In the case of a wakeup contention condition on a BSC line, one DLE ENQ may be lost. This situation is resolved by using the primary/secondary relationship between the nodes (established at SIGNON) to determine which side sends data first.

   If one side receives a Unit Exception when it tries to to write the DLE ENQ, that side uses the primary/secondary relationship in the same way as above after issuing a Read Skip to clear the contention.

   If one side receives DLE ENQ instead of data in response to its DLE ENQ, then that side waits (issues a read CCW) for DLE STX data-buffer DLE ETB to be transmitted if it is the secondary, or sends DLE STX data-buffer DLE ETB if it is the primary.

   The flow for contention is:

```
 Side 1                              Side 2
(Primary)                           (Secondary)

 P(cc)                                P(cc)

    R                                 R

   HDV                                HDV

 W(cc)+++ DLE ENQ ++><-DLE ENQ ──── W(cc)
    R <──────────────++++++++++++++> R

 W(cc)-DLE STX data-buffer DLE ETB-> R

    R <-DLE STX data-buffer DLE ETB-W(cc)

              null-buffer
    W ──────────── or ──────────────> R
              data-buffer
```

3. If one side goes away while I/O is suspended, the other side will detect it at resume time by a timeout as follows:

```
        Side 1                  Side 2

          P(cc)                   P(cc)

          R                       R
                                goes down
    Wake up
     ------>HDV

          W(cc)---- DLE ENQ -->

          R <----Time
                 Out
```

When the read times out in this case, normal time-out processing should be used. This includes sending a NAK.

4. If one side restarts while I/O is suspended, the following scenario will occur:

   Case A

```
        Side 1                  Side 2

           P(cc)                   P(cc)

                                   R

                                 Restarts
          R <---- SOH ENQ ------ W
    Detects
     SOH ENQ
     and goes to
     restart routine

          Normal Signon Processing
```

   If the restart occurs at the same time as a resume, the following can happen:

   Case B

```
          Side 1                  Side 2

             P(cc)
             R                                Restarts
Restarts   HDV <---- SOH ENQ ------ W

             W ------ DLE ENQ ----> R     Note that the DLE ENQ is not
                                          a valid response to SOH ENQ.
             R <---- SOH ENQ ------ W     So signon code tries again.
```

   Detects SOH ENQ and goes to restart routine.

   Note: If both sides try to restart during while I/O is suspended, the signon contention protocol will resolve contention problems.

5. A system which has not agreed to use PREPARE at signon time should never receive DLE ENQ. However, if such a sequence is ever received unexpectedly, the proper procedure is to respond to it with a NAK.

```
                Side 1                Side 2
                 P(cc)
                 R
    Suspends     HDV
    I/O
                 W(cc)- DLE ENQ --> R
                 R <---- NAK ------ W(cc)
                 W(cc)-- null-----> R
                        buffer
                 R
                   No data is lost
```

6. If the resuming side sends DLE ENQ and the sequence is lost (due to a hardware error), a timeout will occur and that side should send a NAK (following normal timeout recovery).

```
                Side 1                Side 2
                 P(cc)                P(cc)
                                      R
                                      HDV      Resumes
                                                 I/O
          Lost       <- DLE ENQ ---- W(cc)
                       Time Out ---> R
                 R <------ NAK ------ W(cc)
                 W(cc)-- null-----> R
                        buffer
```

7. Note that the only valid sequences read by the read which is chained to a PREPARE CCW are DLE ENQ, SOH ENQ, or NAK. All these cases have been covered in the preceding examples. The "resumer" never immediately sends a null or non-null data-buffer, nor is there any way such a buffer can be received until normal sequences are used.

8. The flow below shows what happens when side 1 requests I/O suspension and side 2 does not want to enter this state because it has something to send.

```
          Side 1                          Side 2

          R <-----data-buffer ---------- W(cc)
          W(cc)-- null-buffer DLE ETX--> R
          R <---data-buffer DLE ETB ---- W(cc)
          W(cc)---- null-buffer -------> R
             Normal Sequencing Continues
```

9. Either side may continue to request suspension even if the other side does not agree to it.

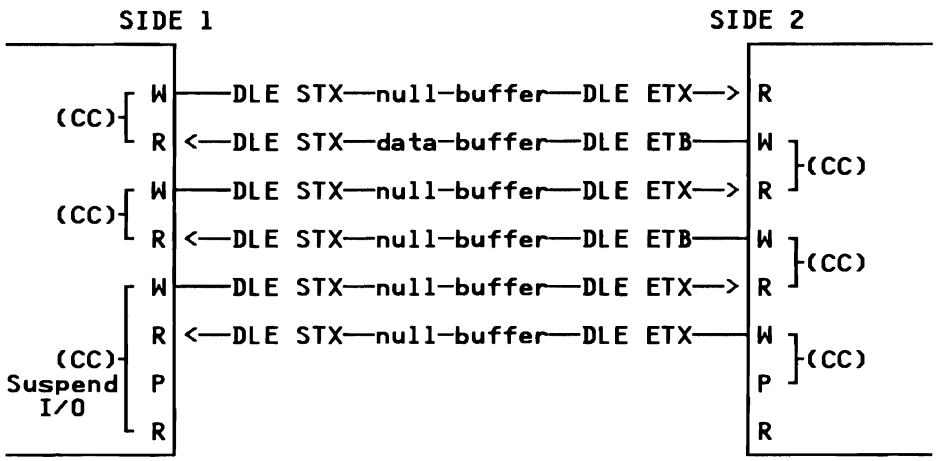

A.3.5 CTCA Error Protocols

Symbols in flows below are:

R = Read
P = PREPARE
W = Write
S = Sense
C = Control
TO = Time Out
cc = Command Chain

1. NAKs are never sent on a CTCA. If a NAK is ever received, the link should be taken down.

2. The hardware prevents any contention during I/O resumption.

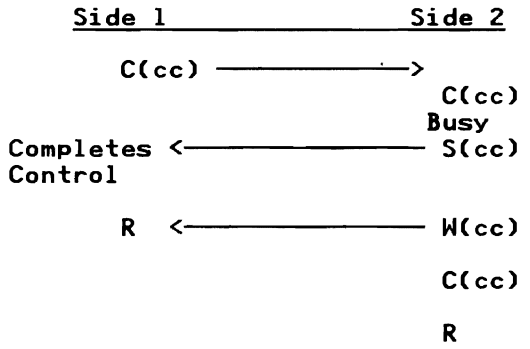

3. If one side goes down while I/O is suspended, the other must detect this during the resume as follows:

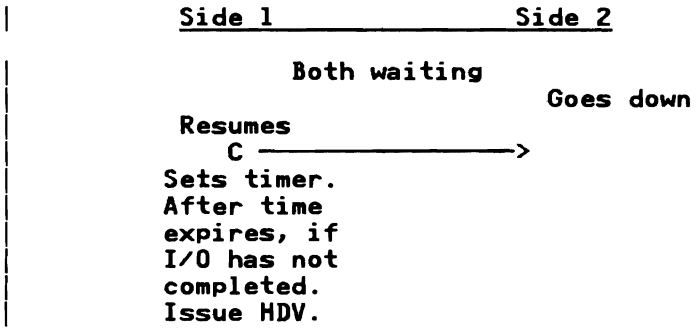
Side 1
Side 2
Both waiting
Goes down
Resumes
C
Sets timer.
After time
expires, if
I/O has not
completed.
Issue HDV.

4. Restart

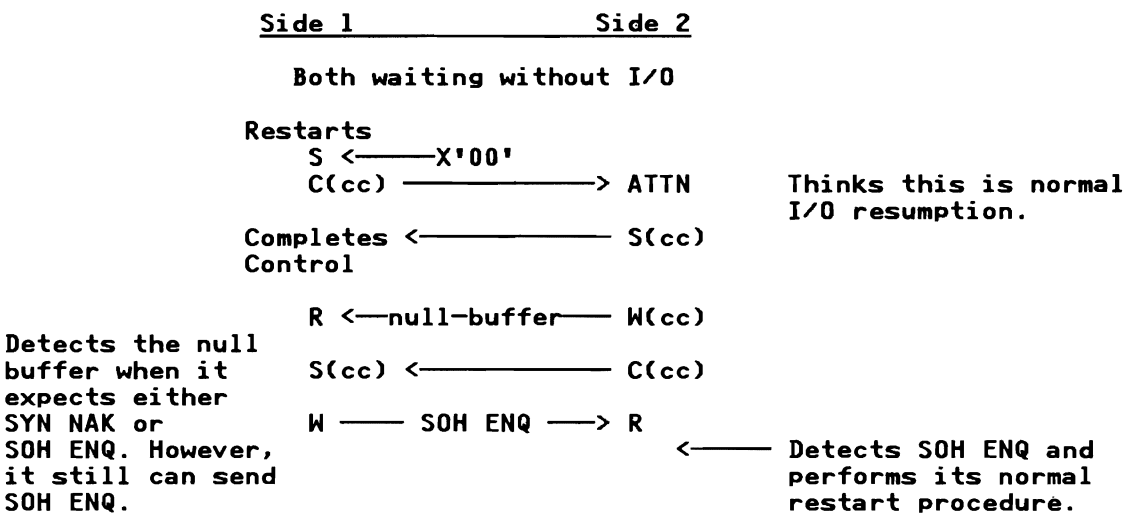

5. If one side does not want to accept a request to suspend I/O, it responds normally to a null-buffer ended by DLE ETX and I/O suspension is not entered by either side.

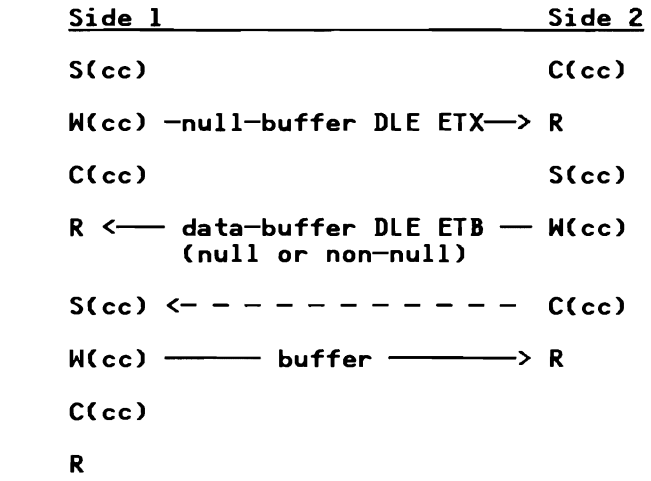

If side 1 still wants to enter I/O suspension after the other side has said no, it can again send a null-buffer ended by DLE ETX. In the above example, the last write (buffer) would be replaced with:

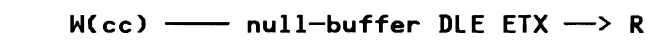

If side 2 responds with null-buffer DLE ETX immediately, suspension will then take place. However, if 2 responds with a normal buffer, normal I/O continues.

6. Side 1 may go on writing null-buffer DLE ETX if it still wants to suspend the I/O. Suspension will not take place until side 2 responds immediately with the null-buffer DLE ETX.

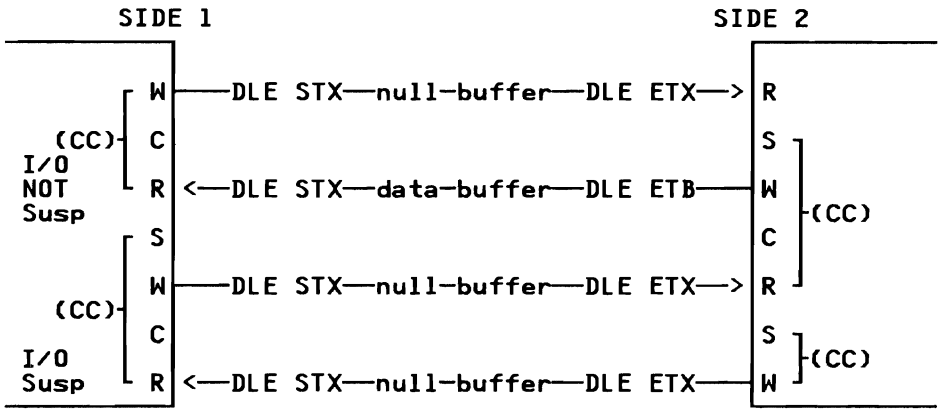

7. All reads should be for a full buffer to prevent possible loss of data in unexpected conditions. This is true even during I/O resumption when the first transmission is expected to be a null-buffer only. Normal BCB sequence checking and error handling apply during I/O resumption.

A.3.6 Wait-a-Bit and PREPARE

1. PREPARE can never be used during any wait-a-bit condition (by definition) because no files are active when a PREPARE is initiated.

2. PREPARE may not be requested with Wait-a-bit all set because that would imply that the side wanting to suspend I/O could not receive data. This would make resumption difficult.

A.4 Systems Network Architecture - LU Type 0

A.4.1 Overview

NJE uses LUTYPE 0 for all its application-to-application communication in an SNA environment. LUTYPE 0 is a full duplex protocol which is not architected in SNA. Bracket protocol is not necessary since no contention problems exist. Both ends of the SNA session are able to send and receive one or more streams.

Compression is mandatory for LUTYPE 0, but compaction is optional. If compaction is desired, each node may specify the compaction table it will use for transmission to the corresponding receiving node. Each table applies for the duration of the session.

NJE (using SNA) allows for store-and-forward transmission directly through the NCP. Either ACF/VTAM, ACF/TCAM or ACF/VTAME (VSE Advanced Functions only), may be used to control path management. Throughout this document wherever ACF/VTAM is mentioned, ACF/VTAME could also be used if the operating system is VSE/Advanced Functions. Currently, only POWER and JES2 use SNA for NJE, although RSCS Networking Version 2 Release 1 is has been announced with this facility.

A.4.2 Session Initialization

In initiating an NJE session, an application must issue an OPNDST OPTCD=(ACQ,SPEC) or SIMLOGON to obtain connection to another application. The application initiating the session automatically becomes the primary Logical Unit (LU) for the life of that session. (This becomes important later in session termination.)

The session parameters are indicated in the BIND area, and this BIND is associated with the OPNDST request. VTAM presents the BIND to the other application in its SCIP exit. This application then becomes the secondary LU. The secondary validates the BIND parameters and may choose to accept the session. In this case, the secondary issues an OPNSEC, which results in a positive response to the BIND.

After a positive response is returned to the primary, VTAM presents the SDT request in the secondary's SCIP exit. VTAM also returns the positive response to the SDT. At this point communication between the two applications may begin.

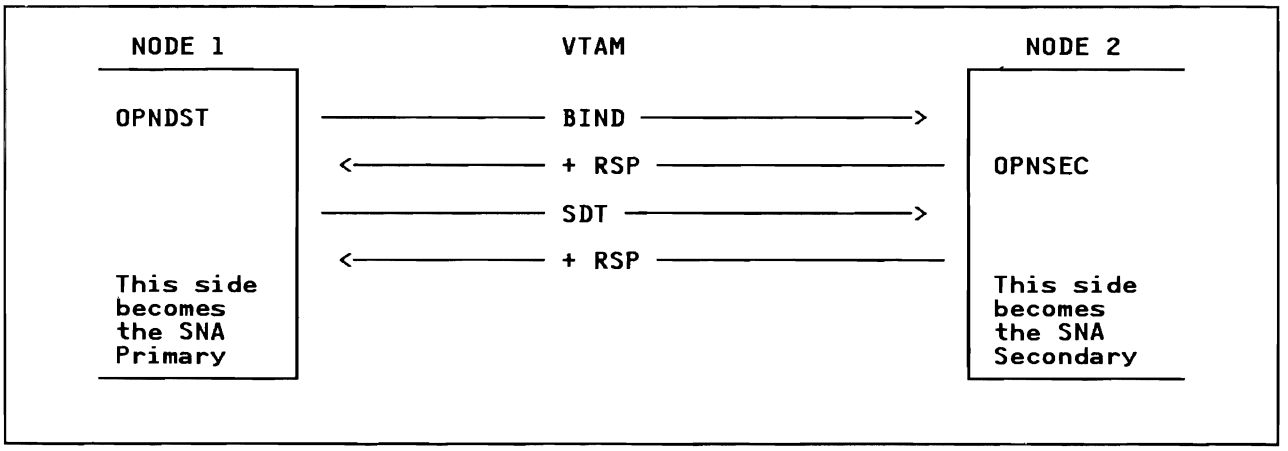

Figure A-15. SNA Session Initialization

The secondary may reject the BIND request. In this case a SESSIONC is issued to return a negative response to the primary.

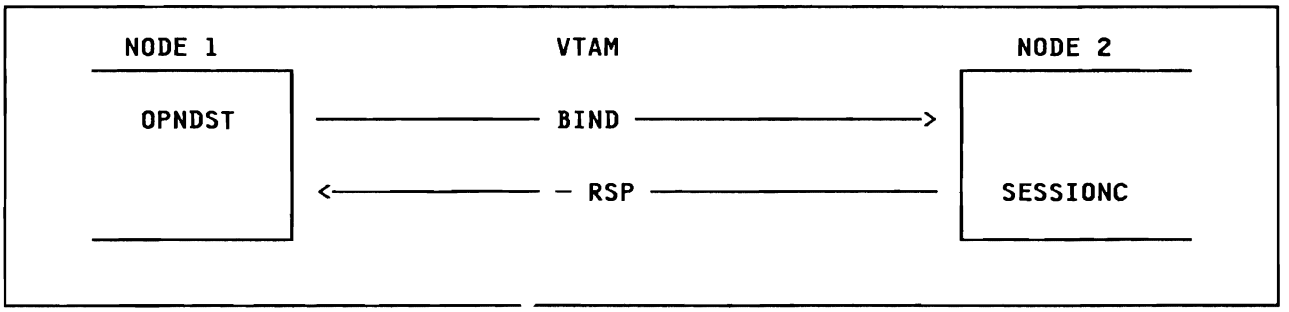

Figure A-16. SNA Session Initialization Error Recovery

See "SESSIONC Usage" on page B-9 for JES2 specific parameters.

A.4.2.1 BIND RU Format

The BIND area is sent from the primary LU to the secondary to activate a SNA session. The BIND contains the parameters which will be in effect for the duration of that session. The BIND image is prefixed in the RU by the RU code for BIND, X'31'. The BIND RU mapping for LUTYPE 0 is as follows:

Byte	Bits	Value	Definition
0	0-7	X'31'	BIND code
1	0-3	B'0000'	LUTYPE 0
	4-7	B'0001'	Cold (non-negotiable)
2	0-7	B'00000011'	FM Profile 3
3	0-7	B'00000011'	TS Profile 3
4			FM usage - primary
	0	B'0'	Single RU chain
	1	B'1'	Delayed request mode
	2-3	B'11'	Definite or exception response
	4-5	B'00'	Reserved
	6	B'1'	Compression allowed
	7	B'0'	No brackets
5			FM usage - secondary
	0	B'0'	Single RU chain
	1	B'1'	Delayed request mode
	2-3	B'11'	Definite or exception response
	4-5	B'00'	Reserved
	6	B'1'	Compression allowed
	7	B'0'	No brackets
6			FM usage - common protocols
	0	B'0'	Reserved
	1	B'1'	FM headers allowed
	2	B'0'	Brackets not used
	3	B'0'	Brackets not used
	4	B'0'	Alternate code not sent
	5-7	B'000'	reserved
7	0-1	B'00'	Full duplex
	2	B'1'	Symmetric responsibility for recovery
	3	B'0'	Reserved (no brackets)
	4-6	B'000'	Reserved
	7	B'0'	Reserved (no brackets)

Byte	Bits	Value	Definition
8	0	B'0'	Pacing for secondary to primary occurs in one stage
		B'1'	Pacing for secondary to primary occurs in two stages
	1	B'0'	Reserved
	2-7	B'bbbbbb'	Secondary Send pacing count
9	0-1	B'00'	Reserved
	2-7	B'bbbbbb'	Secondary Receive pacing count
10	0	B'0'	No maximum RU size for secondary
	1-7	B'0000000'	Secondary maximum RU size
11	0	B'0'	No maximum RU size for primary
	1-7	B'0000000'	Primary maximum RU size
12	0	B'0'	Pacing for primary to secondary occurs in one stages
		B'1'	Pacing for primary to secondary occurs in two stages
	1	B'0'	Reserved
	2-7	B'bbbbbb'	Primary Send pacing count
13	0-1	B'00'	Reserved
	2-7	B'bbbbbb'	Primary Receive pacing count
14			PS profile
	0	B'0'	Basic format
	1-7	B'0000000'	LUTYPE 0
15-25			No protocols specified
26	0-1	B'00'	Reserved-encryption not used by NJE
	2-7	B'000000'	No encryption by VTAM
	2-7	B'bbbbbb'	VTAM encryption
27	0-7	AL1'8'	Length of primary LU name
28-35		CL8' '	Primary LU name
36	0-7	B'00000000'	No user data

See "Bind parameters" on page B-9 for additional JES2 BIND information. See B.4.7.1, "Data Flow Control" on page B-45 for POWER data flow control. See B.3.7.2, "Bind Parameters" on page B-33 for additional RSCS BIND information.

A.4.2.2 Function Management Headers

Function Management Headers (FMH) are used to control the data presentation for the session. FMHs are sent as only-in-chain elements, with a definite response required. FMHs are not compressed/compacted, hence no SCBs are used.

Although several types of FM headers have been defined, only FMH4 and FMH3 are used in LUTYPE 0. These FMHs are exchanged immediately after the session has been initialized and before the initial signon and response signon path manager records are exchanged. FMH4s are exchanged first, indicating whether or not compacted data can be received. Each side then sends either SIGNON, or FMH3 followed by SIGNON, depending on whether the receiver supports compacted data and the sender elects to use compaction.

In other words, both sides must send FMH4 indicating whether compacted data can be received or not, and may send FMH3 indicating whether compacted data is being sent or not, completely independently of each other, except that:

1. If side A indicates compaction not supported as a receiver, then side B must not send an FMH3.

2. Similarly (and independently), if side B indicates compaction not supported as a receiver, then side A must not send FMH3.

Only one FMH may be sent in an RU, hence NJE does not use the SNA FMH concatenation bit. This bits must not be set, and need not be checked. See B.4.7.3, "Functional Management Headers" on page B-45 for POWER specific information. See B.3.7.3, "Function Management Headers" on page B-33 for RSCS specific information.

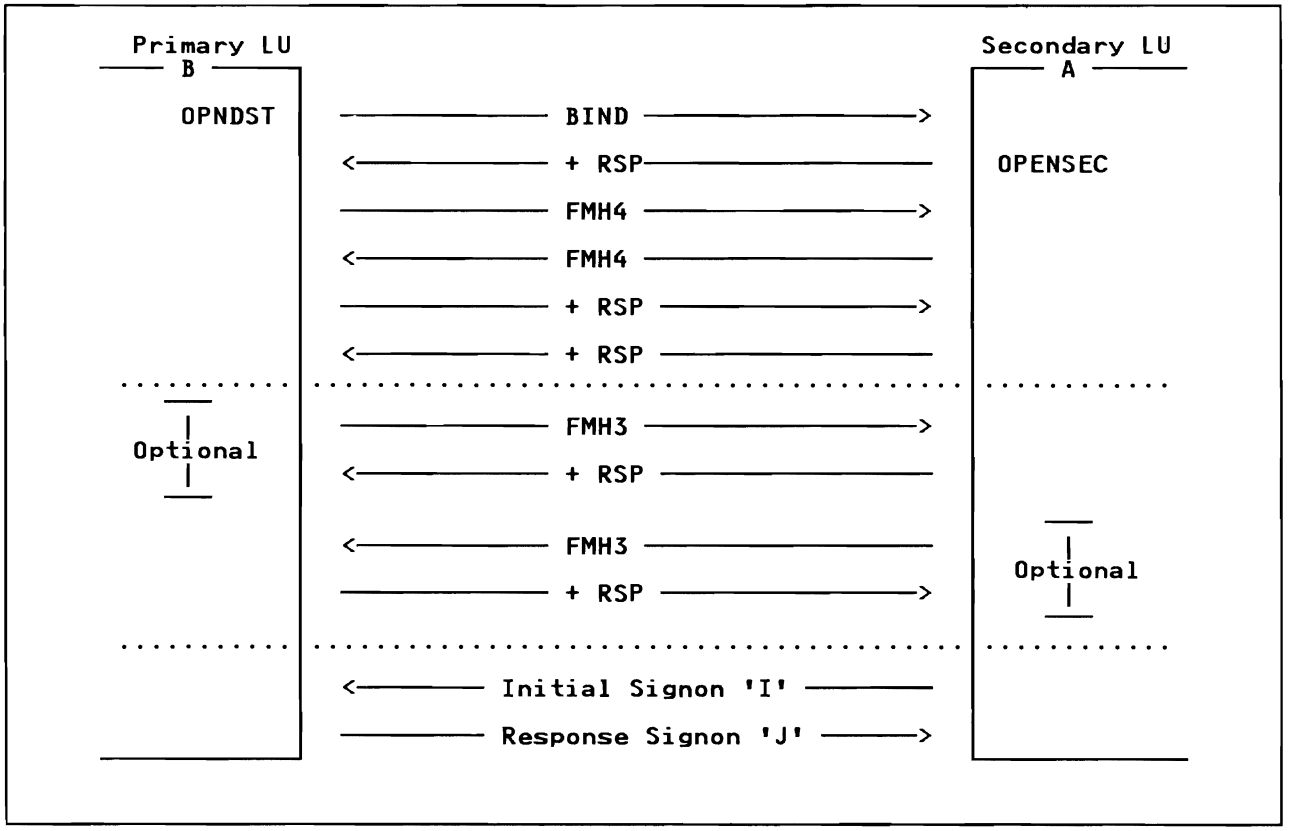

Figure A-17. SNA FMH and SIGNON Flows

The Figure A-17 shows schematically the flows involved in the exchange of FMH4s and FMH3s prior to the SIGNON sequence.

Note that the Primary LU for the BIND flow is not necessarily the NJE Primary Node (the node having the higher NODE name and the sender of SIGNON 'I'). In fact, the figure shows an example where the roles are reversed.

FMH4: FMH4 is a (SNA) non-architected header that is used to exchange network characteristics between the two nodes of a session. Specifically, this header indicates the session RU size. If the RU size is different for the two nodes of this session, the smaller of the two sizes must be used for all data transmission.

Note: The direction of NJE is to eliminate this function of the FMH4. See 8.8, "Connection and Control Records" on page 8-65 and 2.1, "Signon Without Path Manager" on page 2-1 for the preferred method of determining the transmission buffer size via a field in the SIGNON record.

FMH4 indicates whether compaction table and signon records can be received by the node which has sent the header. All NJE products must set the signon accepted bit (B'1'). Compaction may be optionally supported by setting or resetting the compaction supported bit accordingly. NJE products that receive an FMH4 with the compaction supported bit reset (B'0') must not send compacted data and must not send an FMH3. If an FMH3 is received by a sender after having indicated "compaction not supported," then the appropriate error action must be taken as described below.

The optimized fanout flag must be set (B'1'). This means that all NJE implementations must be able to receive data sets preceded by multiple data set headers and perform the necessary fanout. It is not mandatory for all products to create optimized fanout.

RID format 1 is currently the only RID format designed.

The FMH4 format is shown in Figure A-18.

BYTE	BITS	VALUE	DEFINITION
0	0-7	X'08'	Length of header
1	0-1	B'00'	Reserved
	2-7	B'000100'	Header type 4
2-3		H'300' - H'65,535'	Buffer (RU) size
4-5		H'0'	Reserved
6			Features
	0	B'1'	Optimized fanout accepted
	1	B'1'	Signon accepted
	2	B'0'	Compaction not supported
		B'1'	Compaction supported
	3-7	B'00000'	Reserved
7	0	B'1'	RID format 1
	1-7	B'0000000'	Reserved

Figure A-18. FMH4 Format

See "Compaction" on page B-9 for JES2 specific information.

FMH3: FMH3 is specified by the sender and defines the compaction table that is to be used by its receiving partner on the session. If compaction is supported (as indicated in FMH4) by the receiver, the sender may send FMH3 prior to SIGNON. If the FMH3 is not sent, the sender has elected not to send compacted data (note that the sender may still optionally support receiving compacted data when the data traffic is reversed).

If compaction is not supported by the receiver (as indicated in FMH4), the sender must not send an FMH3. This means that a compaction table may only be specified on a session basis, not on individual SYSIN or SYSOUT streams. However, one compaction table may be used to receive data on one node while a different table may be used to receive data on another node.

The compaction table itself contains the master characters, followed by the non-master characters listed in reverse row - major order. See "Compaction Table Format" on page A-32 for a description of compaction characters.

The format for FMH3 is shown in Figure A-19.

BYTE	BITS	VALUE	DEFINITION
0	0-7	X'24'-X'FF'	Length of header
1	0-1	B'00'	Reserved
	2-7	B'000011'	Header type 3
2	0-7	B'00000010'	Compaction table follows
3	0-7	X'03'-X'10'	Number of master characters
4-x			Compaction table

Figure A-19. FMH3 Format

See "Compaction" on page B-9 for JES2 specific information.

Error Handling Protocols: NJE does not require that all aspects of the FMHs be checked.

Specifically, the FMH4 must be checked for the proper length and a valid RU size (> 299 bytes). All other checking is optional. A bit mask may be used to check the integrity of the required and reserved bit values. As the required bits are only for JES2 compatibility, NJE does not require the receiver to check them; however, the sender must set or reset them as specified.

FMH3s must be checked for the integrity of the compaction table, and valid length. See the section "Compaction Table Format" on page A-32, for a description of the NJE compaction table structure. As with FMH4, checking for other bit settings is optional.

Specific error situations and responses are as follows:

1. Broken FMH4 or FMH3; i.e. length < 300 or bits are set wrong or compaction table format wrong, etc.

   Action: Send -RSP and UNBIND.

2. Receiving an FMH4 or FMH3 when not expected or allowed by NJE

Action: Send - RSP and UNBIND.

3. Multiple FMHs in buffer.

   Action: Product choice. Either ignore the extra data or -RSP and UNBIND.

Compaction Table Format: Compaction is a technique that allows specification of two characters in one byte. Interpretation of compacted data is controlled by a compaction table that is transmitted between two networking nodes.

Note: Compaction is only done on SNA transmissions. JES2 is currently the only system that compacts data. (POWER can accept and decompact data, but does not send it.)

To interpret data that has been compacted, build a 16 by 16 matrix such as the sample matrix shown in Figure A-20 on page A-33. Master characters are placed in the matrix beginning with position F0, F2, F3, etc. In the sample, there are 13 master characters: blank,'ADEGILNORSTU'.

When all of the master characters have been placed in the table, the non-master characters are filled in from left to right and from bottom to top. In the sample, the sequence would be:

```
X'15','.<(+|&',X'1E','$*);',X'0C','-/',',X'6C','_>'
'?\:#@'='"BCFHJKMPQ',X'04','VWXYZ0123456789'
'abcdefghijklmnoprstuvwy'
```

Space in the upper left (mxm) of the matrix (where m is the number of master characters) is left blank. In the sample, m = 13.

The limitations on the number of master characters are derived from recognizing that the maximum length of an FMH3 is 255 bytes, and that there are 4 bytes of fixed overhead. If the number of master characters sent is m, then the number of non-master characters sent (for a 16x16 table) is 256 minus (m x (m + 1)). The smallest value of m for which this works is m = 3, which requires that the FMH3 total length = 251 bytes = 4 overhead plus 3 master plus 244 non-master characters.

The largest value of m that works with this algorithm is m = 15. In this case, only 16 non-master characters are sent, yielding an FMH3 length of 36 bytes = 4 overhead plus 15 master plus 16 non-master characters. Actually, up to 16 master characters may be sent, thus zero non-master characters. In this case, only sequences containing the master characters can be compacted and decompacted, and the 16 by 16 matrix need not be built. Each byte is interpreted as two 4 bit indices into the list of 16 master characters, thus yielding two bytes for each byte of compacted data.

```
The FM header type 3 that would be used to send the table shown
in the figure below is:

BYTE      VALUE          MEANING
----      -----          ----------------------------------------
 0        X'5B'          Length of header: 4+256-(mxm), where m=13
 1        X'03'          Header type 3
 2        X'02'          Compaction table
 3        X'0D'          Number of master characters
 4-X'5A'                 Master characters followed by
                         non-master characters (max 252)
```

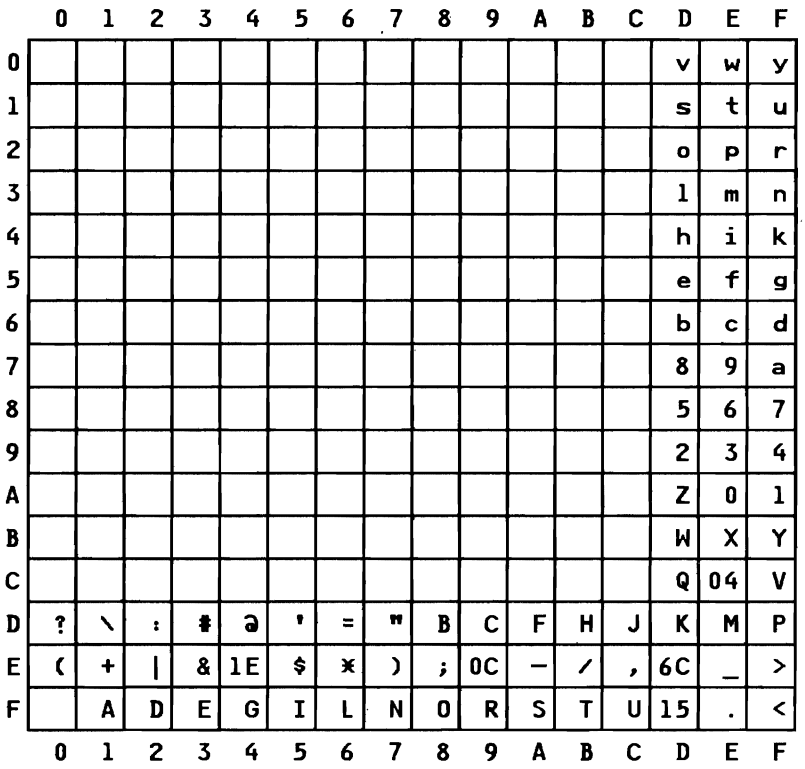

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0														v	w	y
1														s	t	u
2														o	p	r
3														l	m	n
4														h	i	k
5														e	f	g
6														b	c	d
7														8	9	a
8														5	6	7
9														2	3	4
A														Z	0	1
B														W	X	Y
C														Q	04	V
D	?	\	:	#	@	'	=	"	B	C	F	H	J	K	M	P
E	(	+	|	&	1E	$	*	)	;	0C	−	/	,	6C	_	>
F		A	D	E	G	I	L	N	O	R	S	T	U	15	.	<
	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F

Figure A-20. Sample Compaction Table. This figure shows how a sample master and non-master character set are placed in the compaction table.

The corresponding byte sequence for the FMH3 followed by the compaction table is:

```
Meaning      Len type CT  #M  |---- Master Characters ---
Graphic                          A   D   E   G   I   L
Data          5B  03  02  0D  40  C1  C4  C5  C7  C9  D3
Byte Number    0   1   2   3   4   5   6   7   8   9  10

Meaning      --Master (cont'd)----|  |-- Non-Master Characters -->
Graphic       N   O   R   S   T   U       .   <   (
Data         D5  D6  D9  E2  E3  E4  15  4B  4C  4D  ... etc
Byte Number  11  12  13  14  15  16  17  18  19  20
```

Sample Compacted Data: Following is a short data stream in its uncompacted form. The subsequent Figure A-21 on page A-34 shows how the data would appear in its compacted form and how the sample decompaction table (above) would be used:

```
    Sample Data Stream (between the quotes)

'REQ/MODULE/MACRO/SOURCE NAME  -------------' ....
```

```
              Sample Compacted Data ( _ = blank)

Graphic            RE   Q   /    M  OD  UL  E   /    M
Data          55   93  CD  EB   DE  82  C6  F3  EB   DE
Byte Number    0    1   2   3    4   5   6   7   8    9

Graphic       A    C   RO  /    SO  UR  C   E_  NA   M
Data          F1   D9  98  EB   A8  C9  D9  30  71   DE
Byte Number   11   12  13  14   15  16  17  18  19   20

Graphic       E_       -
Data          30   F0  CD  60   0B  0D  15  ...
Byte Number   11   12  13  14   15  16  17

Byte   Value    Meaning
----   -----    ----------------------------------
   1     55     Bits 0 and 1 indicate compacted data
   2            Bits 2-7 ('010101'B) indicate 21 bytes
         93     This value is within the (mxm) portion
                of the matrix. To decode it, use row F.
                Hence, F9 is an 'R' and F3 is an 'E', so
                that 93 when decompacted becomes 'RE'.
   3     CD     This value is not within the (mxm) part
                of the matrix, hence the character 'Q'
                is substituted directly from the table.
   4     EB     /
   5     DE     M
   6     82     OD
 7-20   C6-DE   ULE/MACRO/SOURCE NAM
  21     30     E blank
  22     F0     blank
  23     CD     SCB indicating repeat next 13 characters
  24     60     Dash (-), the character to be repeated
  25     0B     SCB indicating next 11 characters non-
                compressed and non-compacted
  26     0D     Carriage return
  27     15     Line feed
```

Figure A-21. Sample Compacted Data. This figure shows a sample compacted data stream with a byte-by-byte illustration of the results of applying the table to decompact the data.

A.4.3 Session Termination

Two types of termination are described: normal termination with quiesce, and abnormal termination.

Normal termination can be effected by either the application or the VTAM operator. UNBIND flows in these cases, and the LOSTERM exit of the Primary LU is driven.

On the other hand, when a link breaks, UNBIND does not flow from PLU to SLU. Rather, it is simulated to both the PLU and the SLU and the NS exits at both ends are driven.

A.4.3.1 Normal Termination with Quiesce

Normal session termination may be initiated by either the primary or secondary LU. After all session activity has quiesced, the primary LU starts session termination by issuing a CLSDST. From the CLSDST, VTAM presents an UNBIND request to the secondary LU SCIP exit. VTAM returns the positive responses to the UNBIND to the primary LU. The session has then been terminated.

No new sessions may be started until an outstanding CLSDST has completed.

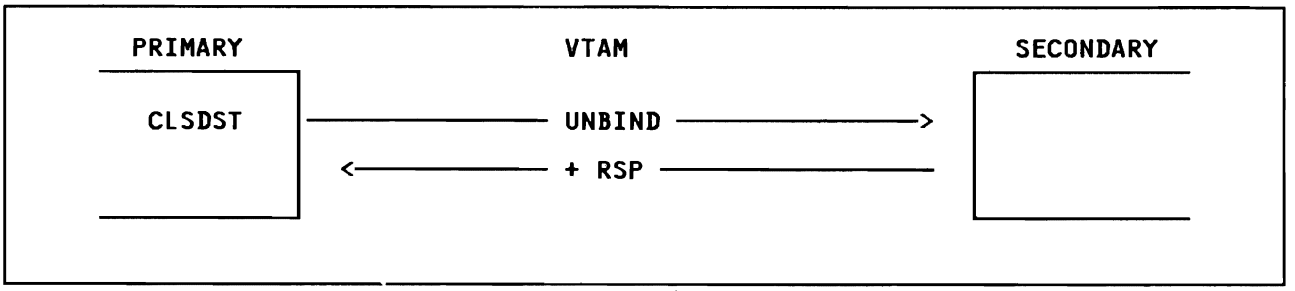

Figure A-22. SNA Session Termination - Primary Initiated

If the secondary LU wishes to terminate the session, it must issue RSHUTD. After all data activity has stopped, the primary will respond with a CLSDST request. This will result in an UNBIND, which will terminate the session. Note that CLSDST may only be issued by the primary LU.

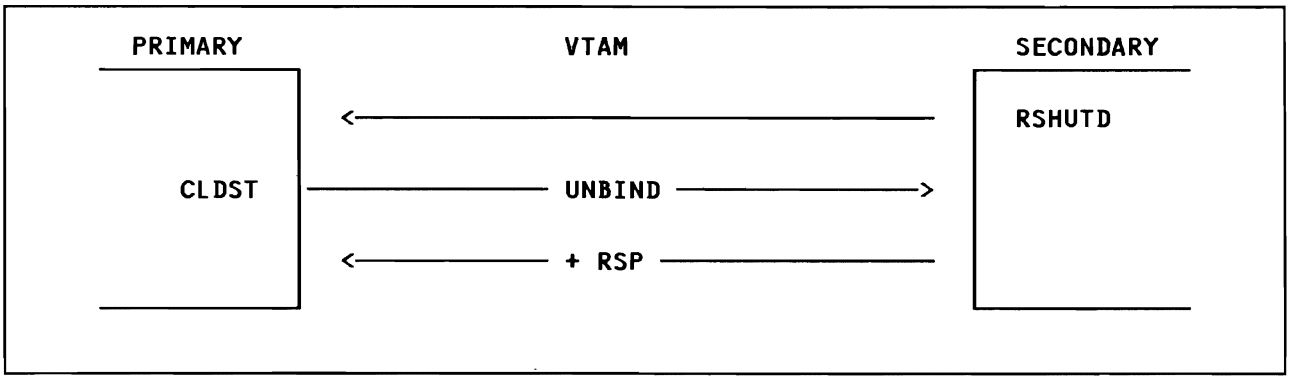

Figure A-23. SNA Session Termination - Secondary Initiated

A.4.3.2 Normal Termination - Immediate

Immediate termination may be initiated by either the primary or secondary LU. It may occur as a result of VTAM errors, a VTAM VARY NET,INACT issued for the application, or a NJE node termination.

If initiated by the primary, CLSDST is requested without waiting for current session activity to quiesce. Otherwise, processing is similar to when the primary initiates normal session termination.

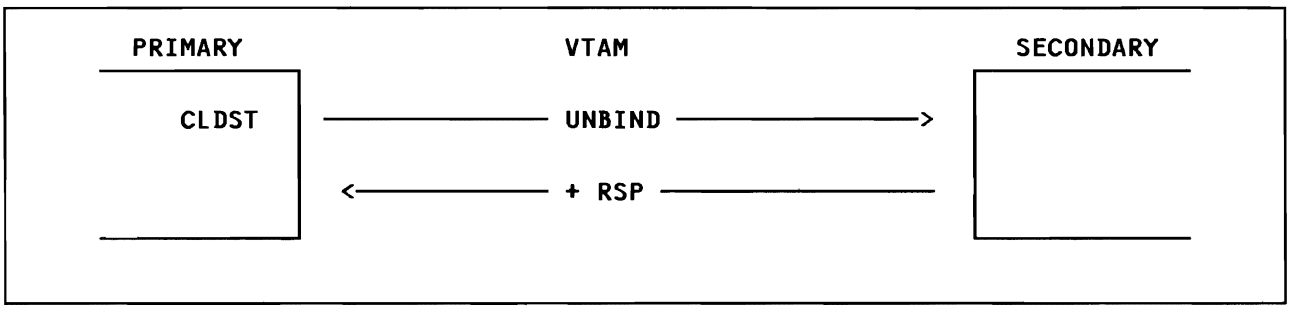

Figure A-24. SNA Immediate Session Termination - Primary Initiated

The secondary LU initiates abnormal termination with a TERMSESS request. Through TERMSESS, VTAM presents the primary LU's LOSTERM exit with a LOGOFF request. The primary LU then responds by aborting current session activity and issuing a CLSDST request. As before, VTAM presents the UNBIND requests to the secondary LU, thereby terminating the session.

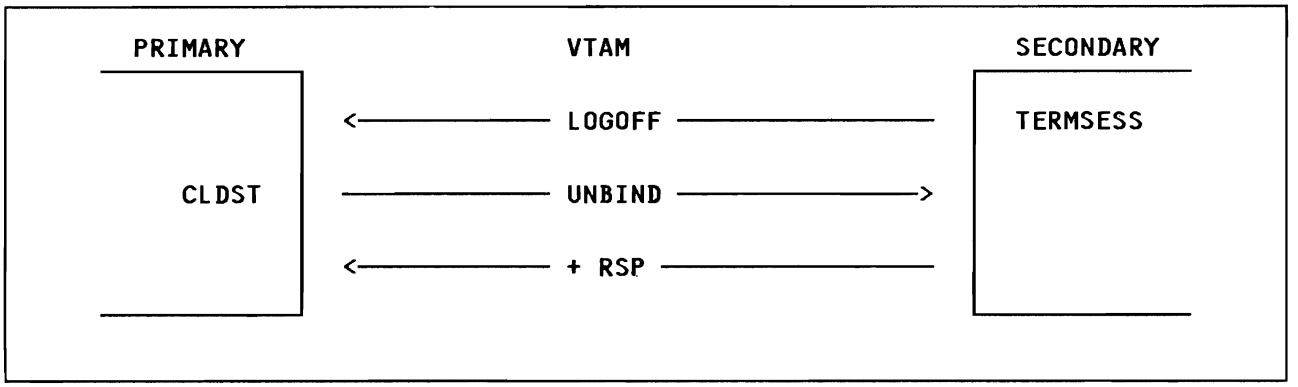

Figure A-25. SNA Immediate Session Termination - Secondary Initiated

See B.1.7.4, "Termination" on page B-10 for JES2 specific information.

A.4.4 Error Recovery

NJE (using SNA) data is sent as only-in-chain elements, with an exception response required. If an exception response is received, all streams will be terminated and the session will be closed using TERMSESS and CLSDST.

If an LUTYPE 0 application receives a data flow command which is not permitted with the LUTYPE 0 architecture, the application will terminate the session.

FMHs are sent only-in-chain, with a definite response required. If an FMH is not acceptable, the receiving application returns a negative response to the sender.

A.4.5 ACF/VTAM Considerations

A.4.5.1 Application Exits

Application exit routines are defined to VTAM during ACB OPEN. The following exit routines could be used by the application for NJE (using SNA).

SCIP: The SCIP exit is scheduled by VTAM to handle BIND, UNBIND and SDT requests.

LOSTERM: The LOSTERM exit is scheduled for normal session termination and for a session failure. The application should either immediately terminate or quiesce the session, based on the completion code in the RPL.

TPEND: This exit is scheduled when a HALT command has been entered or VTAM is abnormally terminating. If a normal HALT is requested, the session should be terminated after all activity has quiesced. For HALT NET,QUICK and VTAM termination, the session should be aborted by the application (using TERMSESS or CLSDST).

NSEXIT: VTAM calls this exit with a CLEANUP RU if the session has been lost.

A.4.5.2 Data Flow Commands

The following data flow commands are defined for VTAM. Only those marked 'YES' may be sent or received by an LUTYPE 0 application.

Normal Flow Commands	Send	Receive
BID	NO	N/A
CANCEL	NO	NO
CHASE	NO	NO
LUS	NO	NO
QC	NO	NO
RTR	N/A	NO

Expedited Flow Commands	Send	Receive
QEC [1]	NO	N/A
RELQ [2]	NO	NO
RSHUTD	YES	YES
SHUTC	N/A	NO
SHUTD	NO	N/A
SIGNAL	N/A	NO

SESSIONC Commands	Send	Receive
CLEAR	NO	YES [3]
RQR [2]	N/A	NO
SDT	YES[3] [4]	YES
STSN [2]	NO	N/A

---

1 Not allowed in FM profile 3

2 Not allowed in TS profile 3

3 Ignored by the application if received

4 Sent automatically by VTAM after OPNDST

Appendix B. System Dependent Considerations and Comments

This section contains system dependent considerations and known deviations from the protocols for each of the NJE products.

Support of Various NJE Features

The following table is provided as a summary chart showing which features are supported by the various products. The features shown are arbitrary and are neither meant to show any "subsets" of NJE, nor which features are optional.

Feature	JES2	JES3	RSCS	POWER
BSC Communications	Yes	Yes	Yes	Yes
SNA Communications	Yes	No	Yes[1]	Yes
CTC Communications	Yes	Yes	Yes	No
Network Path Manager	Yes	No	No	No
Formatted Commands	Yes	(A)	(A)	No
Data Compaction	Yes	No	No	(A)
Output Fan-Out	Yes	(A)	(A)	(A)
Spanned Headers	Yes	Yes	Yes[1]	Yes
Multiple Streams	7	2[2]	Yes[1]	Yes
Signon Concurrence	(A)	(A)	Yes[1]	Yes
Wait-A-Bit All Streams	Yes	(A)	Yes	Yes
Null Buffers vs. DLE ACK0	(A)	(A)	Yes[1]	Yes
Prepare Mode - Suspend I/O	No	No	Yes[1]	No
V.27 Contention Resolution	No	No	Yes[1]	No

Figure B-1. Various NJE Features Supported by the Products

```
Key:
(A) = Accepts, but does not Send.
YES = Supports the Feature
NO  = Does Not Support The Feature
```

---

1 This feature was provided with RSCS Networking Version 2.1.

2 JES3 supports one SYSIN and one SYSOUT stream in parallel.

B.1 JES2

B.1.1 Supported Releases of JES2 Networking

JES2 currently supports the following NJE releases:

1. MVS/SP 1.3.0 and 2.1.0

   - The same level of JES2 was provided for the MVS/SP Version 1, Releases 3.0, 3.1, 3.2 and Version 2 Release 1.0 products. This release adds support for the /*XMIT JECL statement. It also supports up to 1000 nodes and 1000 remotes. In addition, JES2 stores and forwards jobs transparently, without any syntax scanning. This is also the first release to support spool offload which uses an NJE format to store the data offline. Support for this release ends June 30, 1986.

2. MVS/SP 1.3.3 and 2.1.1

   - This release was available in April 1983. SP1.3.3 includes support for the OUTPUT JCL statement. Three new sections are added to the NJE headers: the job scheduling section in the job header, the accounting section in the job trailer and the output processing section (also called the "data stream" section by JES2) in the data set header. Up to 2000 RJE remotes are supported with this release. Support for this release ends June 30, 1987.

3. MVS/SP 1.3.4 and 2.1.2

   - This release was available in April 1984 and provides JES2 support for the 3800-3 Advanced Function Printer (AFP). In this release, JES2 sends stream mode records and also fills in the page-related fields in the headers. Up to 4000 RJE remotes are supported with this release. A PTF is available to allow the SP 1.3.0 and SP 1.3.3 releases to store and forward stream mode data. See APAR OZ75833.

4. MVS/SP 1.3.6 and 2.1.5

   - This release was available December 1985 and provided no new NJE function.

B.1.2 Network Control

B.1.2.1 Network Connection & Control Records

JES2 supports I, J, K, L, M, N, and B (SRCB) records. The initial signon (I) and response signon (J) records must be the only records in their buffers. The other records may be sent in the same buffer.

JES2 allows a user to specify his own function for SRCB types S-Z.

B.1.2.2 Network Addressing, Topology & Routing

Naming Conventions: JES2 supports up to 4000 remotes and 1000 nodes. Remotes and nodes are defined during JES2 initialization. These definitions may only be changed by restarting JES2.

JES2 expects an external remote workstation specification to be either, 'RMTnnnn', 'RMnnnn' or 'Rnnnn'. However, JES2 converts this external specification to an internal binary format and back to the form 'Rnnnn' when sending output to other nodes. Any leading zeros are compressed out.

JES2 allows the installation to specify a 1-8 alphameric name to symbolically define a node. The default node name is of the form 'Nnnnn'. Leading zeros for node numbers are compressed.

All members of a multi-access spool configuration must have the same node name in the network.

Parallel Lines: JES2 supports an unlimited number of parallel BSC lines connecting two nodes.

Dynamic Route Changes: JES2 does not provide a facility for the operator to dynamically change pre-defined connections. To change these connections (which are required for non-JES2 nodes) JES2 must be restarted with different CONNECT statements.

B.1.3 Commands & Messages (NMRs)

JES2 sends an SRCB of X'00' on NMRs.

B.1.3.1 Command Transmission

JES2 builds unformatted Nodal Message Records (NMRs) as a result of the $N operator command.

Formatted (Global) Commands: JES2 builds formatted NMRs for "global" commands. The following global commands are supported by JES2:

Display job - Using $G D operator command

Hold job - Using $G H operator command

Release job - Using $G A operator command

Cancel job - Using $G C operator command

Route job (SYSIN or SYSOUT) - Using $G R operator command

JES2 assumes that destinations for the global route command are syntax checked at the receiving node. As a receiving node, JES2 may reject a global command if only the jobname is specified and JES2 finds more than one job with that jobname. A message is issued to the local console if the global command is rejected.

B.1.3.2 Command Authorization

JES2 supports four levels of command authority, which may be specified on a node basis. The authority levels are as follows:

Network - node has the same authority as local consoles; this includes device, job and system authority

Device - node has the authority to issue device-related commands

Job - node has the authority to issue job-related commands

System - node has the authority to issue certain system commands

B.1.3.3 Message Transmission

Message NMRs are created by JES2 in three instances:

1. In response to a formatted or unformatted command sent across a network link;
2. Via a $DM operator command, which is used to send network messages;
3. For notification messages.

B.1.4 SYSIN (Job) Transmission

B.1.4.1 Store and Forward Transparency

Prior to SP/1.3.0, JES2 would store and forward only MVS jobs with valid JES2 JOB statements. Starting with SP/1.3.0, a JES2 node does not scan any job which is not destined for execution on that node.

As an intermediate node, JES2 will add the JES2 and job scheduling sections to the end of the job header if that header does not already contain those sections. JES2 also adds the accounting section to the job trailer for intermediate node processing.

*/*XMIT:* In SP/1.3.0, JES2 provided a new JECL control statement, /*XMIT. The XMIT statement allows the user to submit non-MVS jobs for network transmission. All JES2 syntax scanning stops after an XMIT statement is encountered within a job. Only the data after the XMIT statement (and before a specified delimiter) is transmitted across the network. Note that a valid MVS JOB statement must precede the XMIT statement.

The jobname placed in the job header of an XMIT job is the name from the preceding MVS JOB statement. If this name is blanks, JES2 transmits a blank jobname in the job header. All relevant job header information (such as priority and accounting information) is taken from the preceding JOB statement.

B.1.4.2 Job Header

The job header and trailer are created when a job is read into the JES2 system. These headers are stored in the same buffer as the Job Control Table (JCT) for the job. Therefore, the total size of both the job header and trailer is limited by the spool buffer size (maximum and most common size is 4008) minus the JCT. (In SP 1.3.6/2.1.5, the maximum spool buffer size is 3992.)

Note: JES2 segments control records in chunks of 253 bytes instead of the preferred 255. This implementation preceded the adoption of 255 as the standard for the NJE protocols.

JES2 Section: The JES2 section of the job header contains two fields, a flag byte and the originator's account number, along with some fields used by the spool offload facility that are zeroed out for NJE transmissions. See 8.4.2, "JES2 Section" on page 8-14 for details.

JES2 appends this section to the job header even at intermediate nodes.

Job Scheduling Section: Beginning with the SP1.3.3 release of JES2, the job scheduling section is sent in the job header. As an intermediate node, JES2 appends this section to the job header. Down-level JES2 nodes require no additional code to store and forward this section.

User Sections: As of JES2 1.3.6/2.1.5, if user sections are added, they must be placed before the spool offload section (even though the spool offload section is not sent on a NJE transmission).

B.1.4.3 Job Trailer

The actual number of lines, cards, pages, and bytes and the execution stop time are set during job termination. These counts are for the total job output, regardless of how many data sets are actually being sent in this transmission. These numbers are used in writing SMF TYPE 26 records for job purge. These fields would be zero for spin data sets, if the job has not completed execution before the data set is transmitted.

Accounting Section: The accounting section is sent in the job trailer beginning with the SP1.3.3 release of JES2. As an intermediate node, JES2 appends this section to the job trailer for SYSIN jobs.

B.1.4.4 Data Set Header

Record Characteristics Change Section: JES2 does not support the receipt of an RCCS prior to the receipt of a JOB statement.

JES2 does not create or receive a segmented record characteristics change section.

Maximum SYSIN Data Record Length: The fix for APAR OZ88264 set the maximum NJE SYSIN record to 252 bytes. Records are written to spool in 256 byte maximum segments. The SYSIN record cannot be spanned, thus the maximum is 252 with the extra bytes being used for record control.

B.1.4.5 Acceptable Job Streams

Multiple Jobs Between a Header and a Trailer: JES2 cannot accept more than one job between a job header and trailer.

Errors in JES2 JECL statements: Prior to the fix for APAR OZ51862, JES2 would issue a receiver cancel if an error was encountered while processing JECL statements in a job. In this fix, JES2 input processing was changed to skip the remaining records in the job and queue the job (with error messages) for output.

With the fix for OZ93366, error messages are returned to the origin node and user.

B.1.4.6 Notification

When a job is first read in by JES2, the TSO/VM userid for notification is stored in the job header if 'NOTIFY' was specified on the JOB statement or a JECL /*NOTIFY statement was included in the job. The originating node name is also stored in the job header. If the 'NOTIFY' node name is not the origin node name, JES2 replaces the origin name in the job header with the 'NOTIFY' name.

If the job header contains a userid and the job is transmitted from its origin node for execution on another node, the job transmitter issues a message to the TSO/VM user indicating that the job was transmitted for execution.

When the job completes execution, a notification message is directed to the TSO/VM userid on the origin node . If the origin node and execution node are the same, this results in an OS 'SEND' command speci-

fying the userid. Otherwise, the notification is sent in a Nodal Message Record to the origin node, and that node issues the 'SEND' command to the userid.

There is also notification when each group of sysout data sets reaches its destination node. See B.1.5.5, "Notification" on page B-8 for details.

B.1.4.7 SYSIN Job Routing

Use of System Qualifier: There is no facility to route a job to a particular system-id in a JES2 multi-access spool complex. Either AFF= must be coded on the /*JOBPARM statement, or a job class structure must be used to control which system executes the job.

Use of Userid: The use of Userid on the /*XMIT or /*XEQ JECL statement is provided for routing SYSIN job streams to a VM userid. As the target execution node, JES2 ignores the NJHGXEQU field.

Operator Rerouting: JES2 operators may change the execution node (but not userid) for jobs on the JES2 queue, through the $R XEQ,... command.

Undefined Node: If JES2 receives a SYSIN job destined for an undefined (not unconnected) node, it will queue the job for local execution (with no error message).

B.1.4.8 Jobid Assignment

The JES2 jobid is a halfword binary number that is assigned when a job first enters the system. This number is unique within a JES2 system. The job header always contains the original (input system's) jobid.

When a job is transmitted from one system to another, the receiving system attempts to assign the original jobid (from the job header) to the job that is being received. If this number is currently in use on the receiving system, a jobid is assigned as if the job were being read in locally, that is, the job counter is incremented by one until an available number is found. A new jobid is assigned even if a part of the original job is on receiving system (as may occur for spin data sets). The newly-assigned jobid is not transmitted in the job header.

B.1.5 SYSOUT (Job Output) Transmission

B.1.5.1 Store and Forward Transparency

All networking levels of JES2 transparently store and forward SYSOUT data sets.

As an intermediate node, JES2 will add the JES2 and job scheduling sections to the end of the job header if that header does not already contain those sections. JES2 also adds the output processing section (also called the "data stream" section by JES2) to the data set header but will not add the accounting section to the job trailer.

JESNEWS: In certain situations, JES2 will attempt to append the JESNEWS data set to the front of the job log data set. This occurs when the job log data set is transmitted from either the execution node or an intermediate node. JESNEWS is appended only if a) it exists on the transmitting system; and b) the job log contains variable length records.

Spanned Records with Carriage Control: Prior to the release of SP 1.3.4, JES2 did not support spanned records with carriage control as either a destination or intermediate node. Support for spanned records with carriage control has been added in SP1.3.4. A PTF (fix for APAR OZ75833) is necessary for down-level nodes to store and forward spanned records with carriage control.

Trailing Blank Truncation: While JES2 truncates trailing blanks on spool, it keeps track of the original record length and restores it upon re-transmission. (JES2/SP 1.3.6/2.1.5 provides a new option to preserve trailing blanks on spool for specified SYSOUT classes.)

B.1.5.2 Job Header

As an intermediate node, JES2 will add the JES2 and job scheduling sections to the end of the job header if that header does not already contain those sections.

Job Copies: APAR OZ67633 changes JES2 processing of multiple job copies. With this APAR applied, JES2 sets and uses the job copies field in the job header. This field is multiplied by the copies field in the data set header to give the total number of copies for a data set. This APAR only applies to JES2 1.3.0 and later releases. (This APAR supersedes APAR OZ57166.)

This APAR also affects processing of the $TO operator command. This command is used to update the characteristics of data sets while they are residing on a JES2 node. The 3800 characteristics is updated only if the existing data set header contained a 3800 section. JES2 does not update any data sets which have multiple 'clone' copies (i.e., created by a $N operator command, or /*JOBPARM copies).

B.1.5.3 Job Trailer

Accounting Section: As an intermediate node, JES2 will not add the accounting section to the end of the job trailer for SYSOUT.

B.1.5.4 Data Set Header

If this is the execution node, the data set header is created when the data set has been selected for transmission. If this is an intermediate node, the original data set header is transmitted. As an intermediate node, JES2 will add the output processing section to the end of the data set header if the header does not already contain that section.

The JES2 SYSOUT receiver spools data set headers as they are received. The data set headers for a particular data set are stored contiguously in a spool buffer. Additional buffers may be used to spool the headers. Data set headers may not span buffers; therefore, the size of the data set header is limited by the spool buffer size (minus the size of the I/O Block (IOB) which precedes the buffer). The spool buffer size is usually 4008 (4096-88 for the IOB),

The first spool buffer of the data set contains a Spool Control Record (SCR) which contains entries for the buffer address and displacement of each data set header. The size of the SCR is limited to the first data set buffer. This means that the number of multiple destinations for a data set (as determined by the number of data set headers) is limited to the number of SCR entries that can fit into the remaining space in a spool buffer after the IOB and the SCR header (which is 3 bytes) have been built. Each SCR entry is 8 bytes long.

3800 Section: On the execution node, JES2 does not create a 3800 section if the data set has only the default 3800 characteristics specified.

Output Processing Section: This section is created on a JES2 SP1.3.3 level and higher. No additional code is required to have down-level JES2 nodes properly store and forward this section.

Page mode records are sent by JES2 SP1.3.4 levels and higher. A PTF is necessary for down-level JES2 nodes to store and forward page mode records. (See the fix to APAR OZ75833.)

JES2 transmits one spool buffer of Scheduler Work Block (SWB) data per data set.

Multiple Data Set Headers: JES2 is able to send and receive multiple data set headers for a data set with more than one destination.

B.1.5.5 Notification

When any of the job's output data sets reaches the destination node, the destination SYSOUT receiver issues a message to the TSO/VM userid specified in the job header. This notification message indicates where the job's data sets were received and is sent to the job's origin node. See B.1.4.6, "Notification" on page B-5 for a description of the end-of-execution message.

B.1.5.6 Job Output Routing

Default Output Routing: Unless otherwise specified, JES2 routes job output back to the origin node and remote workstation (but not back to the Userid in the case of VM) unless specifically routed with a /*ROUTE or other JCL/JECL statement. See "Default Output Routing" on page B-31 for RSCS considerations.

Undefined Node: If JES2 receives a SYSOUT job destined for an undefined (not unconnected) node, it will queue the job for local processing (with no error message).

Operator Rerouting: JES2 operators may change the destination node (but not userid) for output on the JES2 output queue, either by job or by output group.

B.1.5.7 Interactive Data Transmission Facility

A file sent by the Interactive Data Transmission Facility (also called "Netdata") has an external writer name that is identical to the remote/userid field in the data set header. Each transmitted data set is preceded by an internal Netdata header. The data records may be any record length, but TSO/E passes the records to JES2 as fixed-length 80 byte records without carriage control.

When this type of file is received at the destination node, JES2 provides a user exit (Exit 13) to allow the installation to retain or delete the incoming file. This exit may also change the target userid. Based on a return code from the exit, JES2 issues the message "MAIL FROM (node/userid)" to the TSO userid.

If the userid is not valid for this system, the sender is notified with the message "MAIL TO (node/userid) DELETED, INVALID USERID".

B.1.6 Stream Support and Control

B.1.6.1 Multiple Streams

JES2 supports up to 7 job transmitters, 7 SYSOUT transmitters, 7 job receivers and 7 SYSOUT receivers per network line. The total number of job and SYSOUT streams per line cannot exceed 8. This means that the sum of the job receivers and SYSOUT receivers on a line is less than or equal to 8 (similar logic applies for transmitters). The operator has the capability to start and stop each individual stream. A console stream is always defined for an NJE link.

B.1.6.2 Stream Initiation & Suspension

Receiver Cancel: With the fix to APAR OZ79168, JES2 will correctly respond to a "receiver cancel" with an Transmitter Cancel (SCB of X'40') or EOT.

Operator Control (of Streams & Lines): The operator commands that control NJE devices - transmitters (e.g., Ln.JTn) and receivers (e.g., Ln.SRn) start and stop the streams on an individual basis.

B.1.7 SNA Support

B.1.7.1 Initiation

JES2 uses OPNDST OPTCD=(ACQ,SPEC) to connect to another SNA node. The node issuing the start networking command ($SN) becomes the primary LU.

Bind parameters: JES2 creates the BIND from three sources: MODTAB, JES2 initialization parameters and an internal table which forces certain parameters such as TS and FM profile types. JES2 allows only the maximum primary and secondary RU sizes to be variable. JES2 requires the network topology and output dispersal flags to be on in FMH type 4.

SESSIONC Usage: JES2 sends SESSIONC for the following cases:

1. Invalid application name in the JES2 application table
2. No logical SNA line available for this session
3. Invalid parameters in the BIND
4. This node is already in session or another OPNSEC is pending

B.1.7.2 F M Headers

Compaction: JES2 always sets the compaction indicator on even if compaction is not being used for that session. This is because JES2 will always receive a compaction table even though it may not be sending one.

JES2 only sends FMH3 if compaction for the receiving node is indicated during JES2 initialization (via APPL or Nnn statements) and the receiving node has indicated compaction is accepted via the flag bit in FMH4.

B.1.7.3 RU Composition

RU Size Determination: The NJE RU size is determined by the &TPBFSIZ JES2 initialization parameter. It is set in three fields during NJE session establishment: in the BIND, the FMH type 4 and the initial and response signon records. However, JES2 only checks the RU size when receiving the FMH type 4. If the received size from the other node is different from this node's buffer size, the smaller of the two RU sizes is used.

RU Multiplexing: JES2 sends only 1 type of record in an RU, but is able to receive multiplexed records within one RU.

B.1.7.4 Termination

A JES2 NJE session using SNA is normally terminated by a $PLNE command. Normal session termination may be initiated by either the primary or secondary LU.

The purpose of abnormal termination is to clear the session as quickly as possible. A session is abnormally terminated when the line is restarted ($ELNE).

TERMSESS is sent by JES2 for the following reasons:

1. CLEANUP RU received in NS exit

2. Logic error - no OPNDST AUTH=(ACQ)

B.1.8 BSC & CTC Support

B.1.8.1 CTC Initialization

JES2 issues the SYN NAK as specified in A.2.1, "Initialization" on page A-8.

B.1.8.2 CTCA Attention and Busy Status

JES2 performs one retry in this condition which resets the attention and busy. If the attention busy occurs during the retry, it is considered a hardware error and the line is drained. If it is a temporary condition, the one retry will always clear the condition.

B.1.8.3 Error Recovery

JES2 follows the actions in Figure A-5 on page A-4, and terminates after 10 errors.

B.1.8.4 Use of Null Buffers

JES2 uses DLE ACK0 instead of null buffers for positive acknowledgement when there is no data to send.

B.1.9 Accounting

B.1.9.1 Accounting Records

JES2 uses SMF Type 26 (job purge) records to record all successful SYSIN job transmissions, and SMF Type 57 records to record all successful sysout transmissions. Since multiple sysout data sets may be transmitted within a job header and trailer, this record may represent multiple sysout data sets.

None of these records contains the node name of the local node. This could be a problem when combining records from multiple sites.

Type 26 Records: Execution node name (and other fields related to the execution node) are not recorded in the type 26 record when the jobs executes on the origin node.

Type 57 Records: The following standard job-header information is missing from the type 57 record cut by JES2:

- Jobname
- Time and Date on reader at original node
- User Identification from common exit parameter area

NJE Network Management Records: JES2 records the following information reflecting network events.

- 55 Network Sign-on
- 56 Network Integrity (invalid password)
- 58 Network Sign-off

Other Records: The following information is missing from the type 6 record cut by JES2 and other records (e.g., type 4, 5)

- Original Job Number
- Original Node Name

B.1.9.2 Network Account Number

JES2 uses the Network Account Number in the job header as follows:

1. If specified by the user on the /*NETACCT JECL statement

2. If not explicitly specified in the JECL, then it may be converted from a local account number to a network account number through local-to-network account table set up by the NETACCT JES2 initialization statements,

3. Otherwise, it defaults to the local account number (if fix for OZ60994 is applied).

B.1.10 Miscellaneous Considerations

B.1.10.1 User Exits

Exit 13 is invoked from the sysout transmitter when the data set header has been read and processed for a file sent by the TSO/E Interactive Data Transmission Facility. The exit can be used to screen incoming files or to notify the recipient that an incoming file has arrived.

There are no other exits specifically for NJE, but the exits in JES2 input processing could be used to screen jobs coming into a JES2 node for execution.

- Exit 2 (Job Statement)
- Exit 3 (Job Statement Accounting Parameters)
- Exit 4 (JCL and JECL)
- Exit 20 (End of Input Processing)

The above exits are always taken at the execution node. They are also taken at the submitting JES2 node with the following exception: Exit 4 (and JES2 input processing) does not scan any JCL or JECL after the /*XMIT statement.

B.1.10.2 Spool Offload Considerations

JES2 uses the NJE interface for its Spool Offload feature. This allows jobs and data sets to be transferred in NJE format, using existing header protocols.

B.2 JES3

B.2.1 Supported Releases of JES3 Networking

The JES3 Networking support package was first introduced as a PRPQ to JES3/Release 3 (SU26). The package first became a part of the base JES3 package for JES3 System Product Release 3.0. Current releases of JES3 are SP 1.3.4 (370) and SP 2.1.5 (XA). This chapter addresses only the current release.

B.2.1.1 MVS/SP 1.3.4

This MVS/370 release contains JES3 support for the 3800-3 Advanced Function Printer (AFP). In support of AFP in an NJE environment, this release sends page mode data records, fills in the page-related fields in the headers and supports spanned headers.

B.2.1.2 MVS/SP 2.1.5

This is the latest release of JES3 and is only supported on MVS/XA.

B.2.2 Network Control

B.2.2.1 Network Connection & Control Records

Path Manager Records: JES3 does not include a path manager. It supports only SRCB record types I, J, and B.

B.2.2.2 Network Addressing, Topology and Routing

Naming Conventions (Remote vs. Userid): To the end user and operator in JES3, there are only two levels of qualification for the specification of destinations and origins. These are node and either remote id or userid. For example, to specify the destination of a data set, a user might use the JES3 FORMAT statement in a job as follows:

```
//*FORMAT PR,DDNAME=,DEST=NODEX.SECOND
```

The destination for SYSOUT data sets from the job is node NODEX and 'secondary destination' SECOND. The secondary destination SECOND could be a VM userid or a remote workstation id. The networking code does not know what this secondary destination is, and usually considers it to be a remote id. If the secondary destination must be placed in a field which, by definition, could contain either a userid or a remote id, and a flag set to indicate which it is, the flag which is set will indicate a remote id. An example of this is the destination field in the NMR (NMROUT). When a secondary destination value is placed in this field, it is flagged as being a remote, even if the destination is actually a userid.

Use of the System Qualifier in a JES3 Complex: A JES3 complex can consist one to eight processors. One of the processors is called the Global, and is responsible for complex-wide data set integrity, job scheduling for the complex, processing of SYSOUT, and other functions. The networking code also runs only on the Global. The other processors are called Locals and are mainly responsible for running jobs under MVS. The Locals communicate with the Global via CTC adapters.

TSO users may be attached to any system in the complex and may submit network jobs and receive status and notify messages. The mechanism used to indicate which system in the complex the user belongs to in

the NJE headers is the system qualifier. When a TSO user submits a job, a value is placed in the job header qualifier field NJHGORGQ. Other nodes may place this value in the NMRTOQUL field in the NMR for status or notify messages.

The value which is used as the qualifier is a number from 1 to 8, which corresponds to the relative position of the Main Processor Control table (MPC) in the MPC chain for the appropriate system (as determined by the sequence of MAINPROC statements in the initialization deck). For TSO notify messages, a value of 0 is invalid and the message will go to the default NJE message class and not the TSO user. For jobs which are not submitted from TSO, NJHGORGQ will be 0.

Parallel Links: Zero to three lines may be defined for each node in the network. One or two logical senders will be generated for each line, unless zero lines have been defined, in which case no logical senders will be generated. Please see B.2.6, "Stream Support and Control" on page B-17 for a description of logical senders.

Dynamic Route Changes: A JES3 operator may modify the routes with a command. (There is no dynamic path manager.)

B.2.3 Commands & Messages (NMRs)

B.2.3.1 JES3 Console Service

JES3 supports both MCS consoles (for MVS) and JES3 consoles in a JES3 complex. All JES3 consoles are usually attached to the Global JES3 system, with the addition of one MCS console per JES3 Local for IPLing MVS. This is in support of the 'single system image' that JES3 projects for a given complex. JES3 consoles are addressed using a halfword console number, which is looked up in a table called the CST (console status table). Messages destined for JES3 consoles specify a message class, which in turn is mapped to one or more consoles. Messages destined for MCS consoles specify MCS routing and descriptor codes. (JES3 maps MCS routing codes to JES3 message classes.) JES3 Networking only supports JES3 consoles as networking consoles. Networking messages, if not sent to a specific destination or console, are sent to the default NJE message class specified at initialization.

B.2.3.2 JES3 Use of the NMR

A field exists in the NMR which, for messages, designates the destination. This 8 byte field (NMROUT) is redefined several times to indicate a userid, a remote, or a console destination. The console destination usually has meaning only for command responses for a command that originated at a specific console. For JES3, the console destination always designates a JES3 console, not an MCS console. Therefore, only two bytes are required to contain the JES3 console number. This console number is placed in the field NMRROUT, which is the second two bytes of the 8 byte destination field.

B.2.3.3 Commands

Formatted (Global) Commands: JES3 supports most global (formatted) commands for input. The ROUTE command is ignored. These commands never originate at a JES3 node. The formatted commands are translated into equivalent JES3 commands if one exists.

NMR Command Length Restriction: A command sent to JES3 is placed into a buffer and inserted into the system with a JES3 INTERCOM macro. The INTERCOM buffer can be a maximum of 80 bytes. In addition, certain keywords are added to the commands, with the result that commands greater than 59 bytes long are rejected.

To prevent needless command rejection, JES3 removes any trailing blanks from the command to reduce the length prior to checking for the 59 byte limitation.

B.2.3.4 Command Authorization

Only the following commands may be sent to a JES3 node for execution; all other commands are invalid:

- *I J={job name or job-number}

  display job status by job name or job number

- *I B

  display statistics for number of jobs waiting to be processed

- *I Q[,N={xx|ALL}]

  display status of all jobs

- *F J=jobno,{H,R,C,CP,CO}

  hold, release, cancel, or cancel with print a job

The above inquiry commands will only display information about jobs submitted by the node issuing the command. The modify command will only modify jobs submitted by the node and userid/remote id issuing the command.

The installation can provide a user exit (IATUX35) which will accept other commands, or place further restrictions on the commands listed above. This user exit replaces JES3 standard validation/authorization of the command. (In addition, exit IATUX18 may be used for additional authorization checking.)

B.2.4 SYSIN (Job) Transmission

B.2.4.1 Multi-leaving Header Record Expansion

NJE control records (headers) are not padded with blanks when received by JES3. If a received header has had trailing blanks truncated by the sending system, the blanks are not restored and unpredictable results could occur.

B.2.4.2 Data Set Header (RCCS)

JES3 does not support receipt of the Record Characteristics Change section that accompanies SYSIN data greater than 80 bytes. If received, this section is ignored. It will not be stored-and-forwarded.

B.2.4.3 Acceptable Job Streams

Requirement For Two Job Statements: To transmit a job from a JES3 system for execution at another node, the following JCL stream must be submitted:

```
//abc    JOB  xxxxx          first job statement
//*NETACCT  ....             (optional)
//*ROUTE XEQ  nodename
//xyz    JOB  yyyyy          second job statement
  .
  .
  .
//
```

The first job statement and the NETACCT and ROUTE XEQ statements are stripped off at the submitting node; what is transmitted is the second job statement and whatever follows up to the next job statement or end-of-file. This places the following restrictions on the user who wants to submit jobs from JES3:

- Only one job can be submitted for execution elsewhere.
- Jobs which are submitted must begin with a statement that looks like a MVS JOB statement.
- The user must be familiar with the requirements of the execution node for job statements (accounting syntax, etc.).
- Users submitting jobs from TSO must remember to code NJB in place of JOB on the second job statement, or the second job statement will signal the beginning of a new job.

B.2.4.4 Notification

JES3 sends notification messages at the following times during SYSIN job transmission and execution:

- When a job has finished transmission at an intermediate node (this does not include the submitting node)
- When the job arrives at the execution node

See B.2.5.3, "Notification" on page B-17 for notification during sysout processing.

B.2.4.5 SYSIN Job Routing

Use of System Qualifier: See "Use of the System Qualifier in a JES3 Complex" on page B-13 for details.

Use of Userid: See "Naming Conventions (Remote vs. Userid)" on page B-13.

Operator Re-routing: JES3 does not permit operator re-routing of SYSIN jobs.

B.2.4.6 Jobid Assignment

JES3 tries to assign the original job number to SYSIN and SYSOUT jobs received over the network. Otherwise, it assigns the next available number.

B.2.5 SYSOUT (Job Output) Transmission

B.2.5.1 Store and Forward Transparency

Spanned Record Support: JES3 supports spanned records.

B.2.5.2 Data Set Header

Multiple Data Set Headers: JES3 always splits jobs received with multiple data set headers into separate jobs. Multiple data set headers for one data set will never exist in a job which is sent by JES3.

B.2.5.3 Notification

JES3 sends notification messages for job output as follows:

- When output is queued for transmission at the execution node.
- When output has finished transmission at an intermediate node.
- When output arrives at the print node.

See B.2.4.4, "Notification" on page B-16 for notification during SYSIN processing.

B.2.5.4 Job Output Routing

Use of System Qualifier: See "Use of the System Qualifier in a JES3 Complex" on page B-13 for details.

Use of Userid: See "Naming Conventions (Remote vs. Userid)" on page B-13.

Default Output Routing: As of SP 1.3.4, JES3 will set the default output node to the submitting node.

Operator Re-routing: SYSOUT data sets on the transmission or output queues can be re-routed to another userid/node.

B.2.6 Stream Support and Control

Senders and Receivers: In JES3, the transmitters and receivers are not discussed in the same terms as elsewhere in this document. One to six logical senders per node are generated at initialization time which are analogous to transmitters. The number generated depends on other options such as the number of lines for each node and whether or not multi-streaming (discussed below) is to be used. Each sender is either capable of transmitting both jobs and SYSOUT (normal mode of operation), or only one of the two (multi-streaming). Only one receiving function exists, and it is not generally referred to as a receiver in JES3 publications. The receiver processes each transmission buffer as it is received, so it does not care what it is receiving (job or SYSOUT).

B.2.6.1 Multiple Streams

JES3 can support up to one stream for job transmission and one stream for SYSOUT transmission. In each case, the stream supported is stream 1 (RCB of X'98' or X'99'). However, with multi-streaming, one job stream and one SYSOUT stream can be active on a line simultaneously. To accomplish this, two senders are generated for each line - one to send jobs and one to send SYSOUT. The line manager alternates the two streams for transmission.

Operator Control of Lines: The JES3 operator can control individual transmitters (senders). Limited transmission control is also available in that the operator can place a remote node in hold status; all subsequent jobs scheduled for that node will be held. The operator may later release the node from hold status, then release the individual jobs.

Receiver Initiated Processing: JES3 uses an RCB of X'D0' to indicate to a remote node that this node's receiving function has been turned on. The scenario for this function is as follows:

1. The operator at NODE1 issues the command to stop data reception on the specified line: *S LINEX,NORCV.
2. NODE2 sends a request permission to initiate stream sequence (RCB X'90').
3. NODE1 responds with negative permission (RCB X'B0') due to the NORCV (no receive) in effect.
4. NODE2 varies the sender for the job offline and places the job in specialized reschedule status (the job will wait for the associated device - sender, in this case - to become available).
5. If multi-streaming is in effect, another job of the opposite type (job vs. SYSOUT) could ask for permission to start transmission, with the same result. This could cause more than one sender to be varied offline.
6. Some time later, the operator at NODE1 allows data reception to take place by specifying the command: *S LINEX,RCV.
7. NODE1 now sends a receiver initiated sequence (RCB X'D0') to NODE2.
8. NODE2 now varies all the senders to NODE1 online. Jobs in specialized reschedule status which are waiting for the senders to become available will now be available for scheduling.
9. NODE2 can now reissue the request permission to initiate stream. If all else is well, the job will be transmitted.

JES3 always uses an SRCB of X'D7' to accompany the RCB of X'D0'. This is in itself meaningless, and is only present because an SRCB should always be associated with an RCB. Because JES3 only has the concept of a single receiving function which handles both jobs and SYSOUT, "receiver on" can not be associated with either. Hence, it is impossible to place the RCB of the stream to be initiated in the SRCB that goes with the X'D0' RCB.

B.2.7 SNA Support

JES3 does not support SNA as a communication medium.

B.2.8 BSC & CTC Support

B.2.8.1 CTCA Initialization

JES3 checks the data returned from the sense CCW and if it is a X'00' a control-read sequence is issued, if it is a X'01' a read-sense-write-control-read is issued, and otherwise a write-control-read is issued.

B.2.8.2 CTCA Termination Deviations

In SP 1.3.4, JES3 does not send a SIGNOFF record to terminate a CTCA connection. Instead, a NAK is transmitted and the connection is terminated. In 2.1.5, JES3 sends a signoff to terminate a CTCA.

B.2.8.3 CTCA Attention and Busy Status

Some systems (RSCS, for example) place a TIC command between the standard CONTROL and READ commands for the CTCA. This can cause the remote node's channel program to terminate with ATTENTION plus BUSY status. If JES3 receives an ATTENTION plus BUSY status, the following occurs:

If the last I/O issued was a SENSE command, then it is reissued. If the last channel program issued terminated on the first CCW, and it was a SENSE command, then the channel program is retried. Otherwise, a SENSE command is issued. If the SENSE indicates that either a CONTROL or READ is pending from the remote node, the last channel program is restarted from the WRITE command. If none of the previous conditions is met, then the connection is terminated.

B.2.8.4 CTCA Zero Byte Operation

It is possible for a READ or WRITE operation on a CTCA to terminate with good ending status, but with a CSW residual byte count equal to the original byte count (i.e. a zero-byte operation).

This can occur if a system reset is issued on a remote node. When JES3 detects this condition, the connection is terminated immediately. In JES3 1.3.4, a NAK is not sent; in 2.1.5, a signoff is not sent.

B.2.8.5 Error Recovery

BSC Initialization Error Recovery: JES3 differs from Figure A-4 on page A-3 in the following ways:

1. If data other than SOH ENQ, DLE ACK0, or NAK is received the line will be canceled.
2. If command reject occurs on a write, the line will be terminated. Otherwise the operation will be retried.
3. If intervention required occurs on a read, retry will be performed.

JES3 will retry on Bus Out Check and Equipment Check before terminating.

BSC Error Recovery: JES3 follows the actions in Figure A-5 on page A-4 with the following exceptions:

1. A retry is attempted for unit exception on other than a read or write.
2. On command reject for a read, a NAK will be sent. On other than a read, retry is attempted.
3. On a unit check other than a command reject or intervention required, if the CCW is a read, a NAK is transmitted; otherwise a retry is performed.

JES3 terminates after 20 consecutive errors.

B.2.8.6 BCB Handling

JES3 has the following deviations from the correct BCB recovery procedures in Figure 6-4 on page 6-3:

1. When in state S1 or S3 and a NAK is received (E3), if the line is a CTCA it will be canceled. Also, a NAK counter is maintained (similar to the retry limit), such that after 10 consecutive NAKs are received, the line will be canceled.
2. When in state S2 and a duplicate BCB is received, that data block is ignored, but the last block is not resent.
3. When in state S2 and a missing BCB is detected, the line is canceled.
4. During signon, a duplicate BCB causes the line to be canceled.

B.2.8.7 Null Buffers

JES3 does not send null buffers as a response. DLE ACK0 is always sent as an acknowledgement.

B.2.8.8 Signon Deviations

JES3 does not negotiate the buffer size in the SIGNON records. If a mismatch occurs, the line is terminated.

B.2.9 Accounting

Use of Job Header & Trailer Fields: There are several fields in the NJE job header which relate to the recording of SMF information. For example, NJHGETIM, which is the estimated job execution time, and NJHGELIN, the estimated output print lines. Also, the job trailer is composed entirely of accounting type information. JES3 does not use most of the job trailer fields, and uses only some of the job header fields in recording information for SMF. Some of the accounting information which is recorded by JES3 is culled from other sources, such as internal control blocks.

B.2.9.1 Accounting Records (SMF)

SMF (System Management Facilities) is a function of MVS that allows the collection and recording of various types of system and job-related information. This information is recorded in the form of a number of different records, which are numbered. Installations can process the SMF records with any number of application programs to analyze the data, produce reports, etc.

Two SMF records are recorded by JES3 which contain network related information. These are SMF type 26 and SMF type 57 records. The type 26 record is produced at job termination time and contains various job summary information. The type 57 record is produced for each transmitted job (or SYSOUT) after successful transmission and contains summary and resource usage information related to the network processing of the job.

There are no records produced in JES3 which record information of a network management nature.

B.2.9.2 Network Account Number

Users submitting jobs from JES3 can specify certain accounting information which they wish to accompany their job. This is accomplished by using the //*NETACCT control statement in the job's JCL stream. Refer to "Requirement For Two Job Statements" on page B-16 for information on the placement of this statement in the JCL stream. The information which can be supplied on this statement is as follows:

- PNAME - programmer's name (1-20 characters)
- ACCT - network account number (1-8 characters)
- USERID - userid for origin or notify (1-8 characters)
- DEPT - user's department (1-8 characters)
- BLDG - user's building (1-8 characters)
- ROOM - user's location (1-8 characters)

This information is placed in the NJE job header which is built for the job.

B.2.10 Miscellaneous Considerations

B.2.10.1 Use of Utility Jobs in JES3

JES3 uses what are termed "utility jobs" to process network traffic. When a job executes which produces output destined for another node, or a job is submitted for execution on another node, or a job or SYSOUT data is received for either local processing or store-and-forward, a utility job is created. A "job" in the JES3 sense is a set of well-defined, schedulable elements which define work to be done on behalf of some entity, be it a "real" job, or, in the networking case, a collection of network data. Utility jobs are created to contain specific "scheduler elements" to process the network data. This results, however, in a network job having more than one job number during its life on a JES3 node. For example, when a job is received for local execution, a utility job is created and assigned a job number. When the job actually executes, it is assigned another job number. If the job's output is to be transmitted, another utility job is created to do this and assigned another job number.

B.2.10.2 User Exits

Several user exits exist in JES3 networking code. Some of the exits allow the installation to modify fields in the job or data set headers for received data, and some allow modification of the headers before transmit. A brief description of each of the user exits follows.

1. IATUX35 - this exit allows the installation to perform special command validation/authorization for commands which execute locally.

2. IATUX36 - this exit allows the user the pull accounting information from the first segment of the job header for received jobs and SYSOUT.

3. IATUX37 - this exit allows the modification of the first segment of the data set header for received SYSOUT data sets.

4. IATUX38 - the user can perform special processing of SYSOUT class for received data sets with this exit.

5. IATUX39 - this exit provides access to the data set header constructed by this node for the SYSOUT from locally executed jobs. The exit is called prior to the transmission of the SYSOUT.

6. IATUX40 - this exit allows modification of the job header which is built for locally submitted jobs.

7. IATUX43 - this exit allows access to the entire job header for SYSOUT data sets from locally executed jobs which are to be transmitted. The exit is called at intermediate nodes.

B.3 VM/370 RSCS

Terminology: In this section, the term 'file' is used to refer to a unit of data between a job header and job trailer. In previous sections this same unit is referred to as a 'job'. This terminology is not intended to confuse the reader but is used because the section is written from a VM/370 orientation. Because VM/370 has no job awareness, the term 'job' has no meaning in this environment but 'file' does.

B.3.1 Supported Releases of RSCS Networking

RSCS Networking (Version 1) Release 3 and RSCS Networking Version 2 Release 1 are the current NJE products for VM/370. Version 2 provides native SNA support for NJE transmissions, and corrects a number of deviations noted in this chapter.

Note: Any references to RSCS in this section which use Release n without a Version m are to Version 1.

B.3.2 Network Control

B.3.2.1 Signon

RSCS Release 3 will only act as the 'primary' for sign-on. RSCS sends the signon record with reset BCB specified. In Version 2, Release 1 the above restriction is eliminated.

B.3.2.2 Network Path Management

RSCS does not support a path manager. Any connection between RSCS and JES2 must be predefined by a JES2 CONNECT statement. RSCS will not forward path manager topology (NPM) records.

B.3.2.3 Network Connection & Control Records

RSCS supports only (SRCB) type I, J and B records.

B.3.2.4 Signoff

RSCS will not automatically restart the line after a signoff is received.

B.3.2.5 Network Addressing, Topology and Routing

Remotes and Userids: VM/370 makes extensive use of 'userid's (which it uses to identify virtual machines). Therefore, RSCS always knows or wants to know the userids of originators and final receivers of network data. It also handles data for remote work stations. However, a 'remote' as defined for JES2 and other NJE systems is not used in exactly the same way in a VM/370 environment. The VM/370 'userid' is not the same as a JES2 'remote'. RSCS uses the 'remote' name as a link (or line) name. It does not replace a userid in a network address and RSCS uses it more like a node name than a userid. RSCS is able to handle the same 'remote' ids for work stations that are used by other NJE systems with no restrictions on names. However, a remote and userid can not have the same name on a single node. Also because the networking control blocks (headers and NMRs), do not contain fields for both types of identifiers, RSCS often must use fields defined for 'remotes' for what it thinks of as 'userids'.

Parallel Links: RSCS does not support multiple links to an adjacent node.

Dynamic Route Changes: The RSCS operator may dynamically change links (direct connections to adjacent nodes) and routes (indirect paths to other nodes) through the use of operator commands.

B.3.3 Commands & Messages (NMRs)

B.3.3.1 General NMR Processing

| RSCS Release 3 does not store and forward all fields in the NMR record. Only TO and FROM nodes and TO and FROM userids are handled as well as the message or command text. When the NMR record is used to contain console ids, logical routing information and node qualifiers, this information is lost if the NMR record is sent through an RSCS node. Fields which are not supported are NMRTOQUL, NMRFMQUL, NMRUCM, NMRDESC and most other redefinitions of NMROUT except for userid or remote id.

| In Version 2, Release 1 the above restriction is eliminated.

RSCS returns command responses to the node and user that issued the command as determined from the incoming command NMR. It does not route command responses based on qualifiers and MCS console ids in the command NMR. The console command response is treated like any other message originating from the node upon which RSCS is running.

B.3.3.2 Commands

Formatted Commands: RSCS Networking does not support formatted commands which originate at its node. It will pass formatted commands (received from another NJE node) on correctly. It also will translate | formatted commands into an equivalent RSCS command as follows according to the setting of the NMRFOP field:

1. NMRFOPD (display job) - Query File spoolid
2. NMRFOPC (cancel job) - PURge spoolid
3. NMRFOPA (release job) - CHange spoolid NOHold
4. NMRFOPH (hold job) - CHange spoold HOld
5. NMRFOPR (route job) - TRANsfer spoolid TO locid userid

| With Release 3, if a formatted command passes through a VM/370 to VM/370 connection, the command will not be understood. The reason for this is that the connection between two RSCS nodes must use the VMB or VMC line drivers which do not implement the full NJE protocols. (Two NJI line drivers cannot talk to each other because each insists on being NJE primary.) An example of a configuration which will not handle the commands properly is shown below.

However, in the following configuration, the destination node will translate the command to an equivalent RSCS command. In this case, the RSCS nodes use the NJI line drivers which use the full NJE protocols.

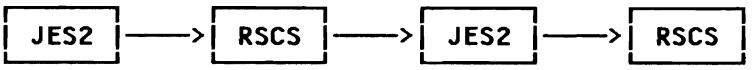

(In Version 2.1 the above restriction is eliminated, because of the common NJE line drivers.)

Fields NMRLEVEL and NMRPRIO are always set to 7 when an NMR is forwarded regardless of what they contained when RSCS received the record.

With Release 3, certain NMRTYPE bits are not stored and forwarded under various conditions, but this was fixed in Version 2.1.

NMRFLAGW and NMRFLAGT fields are both always turned on in messages when the message contains a destination userid.

B.3.3.3 Command Authorization

Remote system operators may issue commands to RSCS to affect the link between RSCS and their system (without authorization). (The AUTH statement was not intended to be used to authorize system operators at remote systems.)

B.3.4 SYSIN (Job) Transmission

B.3.4.1 Store and Forward Transparency

Segmented Headers: RSCS supports segmented job headers, dataset headers (for both SYSIN and SYSOUT) and segmented job trailers up to 4092 bytes.

Trailing Blank Truncation: For any type of file (SYSIN or print or punch SYSOUT) the following modification occurs:

Any records with trailing blanks have the blanks truncated when they are forwarded by RSCS.

In Version 2, Release 1, blanks are only truncated when SYSOUT files are sent with RECFM=V or U (as specified in NDHGRCFM). Files with fixed length records are store and forwarded without loss of trailing blanks.

B.3.4.2 Data Set Header (Record Change Characteristics Section)

RSCS Release 3 does nothing with input dataset headers. It does not look at them at all but merely puts them in the same file with the rest of the SYSIN job they are contained in. All SYSIN data which is store and forwarded through RSCS must be 80 bytes or less. RSCS will reject with receiver cancel any SYSIN job which contains records longer than 80 bytes.

In Version 2, Release 1 the above restriction is eliminated.

B.3.4.3 Notification

RSCS sends a notification message when the file is successfully transmitted on a link. (DMTxxx147I)

B.3.4.4 Spool File Id Assignment

CP assigns a new spool fileid when it receives a file. No attempt made to re-assign the same number it had on the origin node (in NJHGJID).

B.3.5 SYSOUT (Job Output) Transmission

B.3.5.1 Store and Forward Transparency

Segmented Headers: RSCS supports segmented headers.

Print File Transparency: RSCS modifies store and forward data in the following ways for SYSOUT print files only:

1. When records without carriage control of any kind are received, RSCS will forward them with a machine carriage control character of X'09' added. The NDHGRCFM and NDHGLREC fields in the dataset header are changed to reflect this change.

   In Version 2.1 the above restriction is eliminated as long as NDHGRCFM indicates no carriage control. When this field indicates carriage control, RSCS will leave a machine carriage control of X'09' on each record that was received without carriage control.

2. When records with ASA carriage control are store and forwarded, RSCS converts the ASA carriage control to a machine carriage control. The NDHGRCFM header field is changed accordingly.

   Note: Records sent with any machine carriage control character are store and forwarded without any change.

3. Any record (with or without carriage control of any type) sent as a spanned record is store and forwarded without change. In Release 3, this is not true for spanned records in a VM/SP virtual 3800 file or a normal print file sent with a 3800 section and OPTCD=J specified. (See "3800 Files" on page B-29.))

   In Version 2, Release 1 this item is correct as written.

4. RSCS truncates any characters in print SYSOUT records which exceed the number allowed for the printer whose type is used to store the file in CP Spool. The method in which this printer type is determined is described in detail in section "RSCS/CP Spool Interface Considerations" on page B-28. An example is that if RSCS receives a file which it stores in CP spool as a virtual 3211 file, no records longer than 150 characters in the file will be forwarded without truncation by RSCS.

   In Version 2.1 the above restriction is eliminated for store and forwarded files. At the destination node, records are truncated as described above.

Punch File Transparency: RSCS handles incoming SYSOUT punch files as follows:

1. In Release 3, no punch files are received if the NDHGRCFM and NDHGLREC fields of the data set header indicate that the records are greater than 80 characters long (not including carriage control). Such files, including spanned records, are rejected with Receiver Cancel.

   (This problem is eliminated in Version 2.1.)

2. If NDHGRCFM indicates that a punch file contains ASA or machine carriage control, RSCS will replace the carriage control in each record with a machine punch op code. For store and forwarded files, the data set header field NDHGRCFM will be changed accordingly to indicate machine carriage control.

3. If NDHGRCFM indicates no carriage control, the file is forwarded without change (as 80 byte records, no carriage control).

Trailing Blank Truncation: For any type of file (SYSIN or print or punch SYSOUT) the following modification occurs:

Any records with trailing blanks have the blanks truncated when they are forwarded by RSCS.

In Version 2, Release 1, blanks are only truncated when SYSOUT files are sent with RECFM = V or U (as specified in NDHGRCFM). Files with fixed length records are store and forwarded without loss of trailing blanks.

RSCS/CP Spool Interface Considerations: RSCS does not do its own spooling, but relies on the the CP component of VM/370 to create and manage spool files. These files may be stored on spool as either virtual print or virtual punch files. Virtual punch files can contain up to 80 bytes of data and, within the spool, they also contain a punch opcode of X'41'. Virtual print files which RSCS uses (note CP itself supports more virtual print types) are 1403, 3211 and 3800. Each print record is stored with a machine carriage control (or machine opcode) and is limited in its length by the maximum number of records a real device of that type can handle, based on the data set header.

For any records that conflict, RSCS Release 3 will either modify them so that they do not conflict or reject the file with a receiver cancel (RCB of X'B0'). RSCS Version 2.1 will modify any records that conflict (truncate them) so they no longer conflict. This is done only at the destination node.

Figure B-2 is a table of virtual device types which are defined for different incoming NJE file types:

NJE data	CP device type
SYSIN	Punch - limited to 80 characters
SYSOUT[3]	
If NDHGF2PU is on	Punch - limited to 80 characters
If NDHGF2PR is on	Print -
	1403 if NDHGLREC is 133 or less (132 if NDHGRCFM indicates no carriage control) and file does not meet the 3800 criteria
	3211 if NDHGLREC is greater than limits specified for 1403 and file does not meet the 3800 criteria
	3800 if header contains a 3800 subsection with NDHAF1J flag on (3800 file processing is discussed in section "3800 Files" below in greater detail.)

Figure B-2. RSCS Virtual Device Types.

Note: The table is accurate only for RSCS Release 3 and above running on a VM/SP Level 2 and above system and assumes all data headers contain the print/punch flag.

---

[3] When a dataset header contains an RSCS subsection, RSCS uses the originating device type in this section to determine the VM/370 device type. Therefore, the SYSOUT part of this table only applies to files which do not originate on a VM/370 System.

| Once RSCS has determined the device type it is going to use for a file at header processing time, the type will not be changed even if the data records themselves have different characteristics. Once the file is stored in CP Spool, it will not be changed if it is forwarded to another non-VM/370 NJE system. Hence, if RSCS receives a file with individual data records of 150 but the NDHGLREC field says 132, the file will be stored as a 1403 file and all characters after the 132nd character in each record will be truncated for both store and forward and for printing on the RSCS node if an end node. It is extremely important to note, that one result of this processing is that RSCS can not store and forward print records greater than 150 unless they are sent as spanned records. When data for the 3800 printer (which can have data records up to 204 characters long) is stored and forwarded through RSCS, only the first 150 characters will be forwarded intact unless the user turns on OPTCD=J (which causes NDHAF1J flag in the dataset header to be turned on).

| Version 2.1 defines files in the same way as shown in Figure B-2 on page B-28. However, it is able to store | records which are longer than the normal length specified for the device type, without truncating data.

3800 Files

Version 1, Release 3: RSCS handles 3800 print files specially. There are essentially two types of 3800 files: regular print files which are meant to be printed on a 3800 printer (indicated by the fact that their data set header contains a 3800 section) and virtual 3800 files which are currently only understood by a VM/SP system and thus are only originated by RSCS. Virtual 3800 files are also sent with a 3800 section in their dataset headers. For purposes of this discussion, virtual 3800 files may be distinguished from regular 3800 files by two characteristics: First they contain actual 3800 CCW op codes as well as machine carriage control characters and secondly they are the *only* type of file with a 3800 subsection which are originated by RSCS. It also only applies when RSCS is running on a VM/SP2 or above level system.

When virtual 3800 files are created by an RSCS node, they are sent with both an RSCS subsection and a 3800 subsection in the dataset header. RSCS sets the virtual device type in the RSCS subsection to be virtual 3800 so another RSCS system can process the file correctly. Thus RSCS can always recognize a virtual 3800 file it has created. However, when a non-VM/SP system originates a 3800 file, there is no RSCS section and no VM/370 device type field. RSCS does not generally define any of these regular 3800 files as a virtual 3800 file unless the file was created with OPTCD=J. This case is discussed in detail in the next paragraph. If any user or any other NJE subsystem wants to be certain that all 3800 data is preserved (including lines up to 204 characters long) when store and forwarded through RSCS, the file must be created with OPTCD=J specified (which causes NDHAF1J to be turned on).

RSCS creates a virtual 3800 file out of any 3800 print file which has the OPTCD=J flag on in the 3800 section of the data set header. This occurs for store and forward files and for files for which RSCS is acting as an end node. RSCS then removes the TRC byte from each record (unless it is spanned or contains no data) and inserts select CCWs which cause the proper character arrangement table to be selected. If such a file is printed on a real 3800 printer on the receiving node or another VM/SP node, the printout should be as the originator intended. If the same file is sent back to another non-VM/SP NJE system, the TRC bytes are re-inserted and the select CCWs removed.

In addition, 3800 files have ASA carriage control characters changed to machine and machine X'09' CCTLs inserted in records without carriage control in the same way as other type print files (see B.3.5.1, "Store and Forward Transparency" on page B-27). Any spanned record in a file that RSCS is storing as a virtual 3800 file, is always unspanned before it is stored. Hence for virtual 3800 files, these records also may have their type of carriage control modified by RSCS.

The above processing does not cause any data to be lost but it does imply that RSCS modifies data on store and forward. This is a deviation from the protocol outlined in this document. However, RSCS does not intend to change this processing nor does it consider changing a desirable alternative. The current method of handling best insures the data is printed as its originator intended.

Version 2, Release 1: 3800 support is the same as it is on Version 1, Release 3 except that

- RSCS is always running on a level of CP which has virtual 3800 support and so the 3800 files are always handled the same way (as described here).

- Virtual 3800 records sent without carriage control are not modified on forwarding. They are forwarded without cc. However, records sent with ASA carriage control are changed to machine carriage control. NDHDRCFM is changed to reflect this.

- Spanned 3800 records are not unspanned when they pass thru an RSCS node. This is only done when the files are printed on the node RSCS is running on.

- Data is not lost in any record which are store and forwarded.

Support of SRCB X'B0' RSCS Release 3 will reject with a X'B0' RCB any file which it finds containing this SRCB (page mode records).*:* In Version 2.1 the above restriction is eliminated. RSCS will store and forward records with the x'B0' SRCB. It continues to throw these records away when it tries to print a file which contains them.

Spanned Records: RSCS does not support spanned records in any releases below Release 3. Starting with Release 3, it will store and forward print file spanned records up to 32k bytes. It will originate spanned records for virtual 3800 files but will send these records with an SRCB indicating no carriage control (although the first byte in the record will actually be a machine opcode).

Punch files containing spanned records will be rejected because they are longer than 80 bytes.

| *Note:* In Version 2.1 the above restrictions are eliminated and 3800 spanned records are sent with machine
| carriage control indicated in the SRCB.

Release 3 will handle spanned records with machine carriage control and no carriage control when acting as an end node provided the maximum length field in the dataset header (NDHGLREC) is filled in correctly. ASA carriage control is not supported in Release 3 for spanned records if RSCS is the end node.

| In Version 2, Release 1 the above restriction is eliminated.

B.3.5.2 Data Set Header

Use of CP TAG Command: RSCS has no direct JCL equivalent which allows users to specify parameters for jobs or datasets. However, RSCS allows a file originator a limited method of specifying certain fields in the dataset header RSCS sends with a file. The originator can use this method to either override default settings of fields used by RSCS or to specifically indicate values he wants included for fields RSCS does not set. The CP TAG command is used to specify these parameters and it is described in detail in RSCS Publications. The TAG command can only be used for fields in the dataset header. There is currently no method to override job header fields set by RSCS.

Multiple Data Set Headers: The protocol allows multiple data set headers to appear within a job. These headers may indicate different routing for a single dataset or many datasets may be grouped in the job because they are all output from its execution. In the second case, each dataset is preceded by a dataset header. RSCS handles these cases as follows:

1. When a single dataset is received by RSCS with multiple destinations, the dataset is stored and forwarded as separate files (one for each destination).

In Version 2.1 RSCS will only make separate files if the datasets involved are going out on different links.

2. When any dataset within a job has different characteristics from any other in the same job, each dataset is stored and forwarded as a separate file.

3. Spin datasets (flag NDHGF1SP on) are always stored as separate files.

Note: When datasets are stored and forwarded as separate files, each file has the original job header and job trailer.

Use of HOLD in Dataset Header: RSCS Release 3 does not respond to flag NDHGF1HD is the dataset header. This flag, if on, says that the dataset is to be held at its final destination not at the receiving node. Because RSCS does not process headers on all links (see section B.3.10.1, "VM-Specific Line Drivers" on page B-36), the final node is not always aware that this flag is on and it can not be checked and used with consistency in all cases. Therefore, the general principle is never to put the received file in HOLD status.

In Version 2.1, the above restriction is eliminated.

B.3.5.3 Notification

RSCS sends a notification message when the file is successfully transmitted on a link. (DMTxxx147I)

B.3.5.4 Special Processing

Print File Processing: If RSCS is the destination, it will try to print records with CCWs it does not understand and will space after them.

Punch File Processing: Punch files originated by RSCS Release 3 are sent as fixed 80 byte records without carriage control. In Version 2, Release 1 punch files may be sent either with as 81 byte records with machine carriage control, or as fixed 80 byte records without carriage control. This is specified by the file originator on the TAG command (PUNCC option).

B.3.5.5 Job Output Routing

Default Output Routing: When no specific routing is given (such as by JES2 OUTPUT or ROUTE statements), output returned to RSCS from a job executed on a non-VM/370 NJE system goes to the system printer or punch and not to the reader of the user who originated the job.

Fan - Out (Optimized) Support: RSCS will not send multiple data set headers with a file. If it receives multiple data set headers from another node, Release 3 will make multiple copies of the file. Version 2.1 will only make multiple copies in certain situations. See "Multiple Data Set Headers" on page B-30.

Operator Rerouting: The TRANSFER command can be used by the RSCS operator (and by the file originator in Version 2.1) to re-route files to other nodes. In Release 3, this command only changes the tag block; it does not change the network header fields. Therefore, if the file gets sent to another node over an NJE link, the updated routing information will be lost. This restriction is removed in Version 2.1.

B.3.6 Stream Support and Control

B.3.6.1 Multiple Streams

RSCS supports only one stream of each type. For transmission, only a single stream (SYSIN or SYSOUT) can be active at a given time. For reception, a single SYSIN and a single SYSOUT stream may be active simultaneously.

With Version 2, Release 1 the above restriction is eliminated,

RSCS Networking supports only RCBs of X'98' for input and X'99' for output. RSCS rejects any other streams another system tries to send it.

In Version 2, Release 1 the above restriction is eliminated.

B.3.6.2 Stream Rejects

RSCS will reject with an RCB of X'B0' a request to open any stream other than the first one of each type.

In Version 2, Release 1 the above restriction is eliminated.

B.3.6.3 Abnormal Termination Protocol

Release 3: RSCS fully supports the abnormal termination protocol outlined in this document. Note that Release 3 will send both aborts and receiver cancels for a variety of error conditions which may not have been detected in prior releases.

When RSCS receives either permission denied (when it asks to initiate a stream) or receiver cancel (when a stream is started), it puts only the file involved in HOLD status. The operator must use an RSCS command to free the file if he later wants to send it. No other files are affected and RSCS continues to use or attempt to use the stream for transmission.

When RSCS initiates abnormal termination as either a receiver or a sender only the protocol is used to inform the other side of this termination. Operator messages are sent only to the RSCS operator and not to the other system operator.

| *Version 2.1:* When RSCS receives permission denied (when it asks to initiate a stream) it stops using that
| stream until either the RSCS operator resets that link, or until a receiver ready is received (X'D0') for that
| stream. Receiver cancel is still considered to be a file, rather than stream, control record.

| Whenever RSCS Version 2.1 initiates abnormal termination as either a receiver or a sender an operator
| message is sent to the remote system operator. Additionally, the file originator is also informed.

B.3.6.4 Stream Initiation & Suspension

Stream Suspension: Because RSCS Release 3 does not support multi-streaming, it does not react to individual stream suspension. It does not check the low order 12 bits of the FCS. It will suspend all streams when the second bit (wait-a-bit) is on. When wait-bit is on, transmission remains suspended until an FCS with wait-a-bit off is received or the other side sends a DLE ACK0. RSCS does send wait-a-bit for a stream it is receiving when it only has one buffer left to hold the next input.

| In Version 2, Release 1 the above restriction is eliminated.

B.3.7 SNA Support

| RSCS Version 1 does not support SNA sessions for NJE. Version 2 does provide native SNA support.

| B.3.7.1 Data Flow Control

| The only Data Flow Control (DFC) function supported by RSCS Version 2.1 is RSHUTD. If any other
| function is received, RSCS will terminate the session with a CLSDST or a TERMSESS.

| B.3.7.2 Bind Parameters

| RSCS creates the bind from the bind image VTAM returns in the CINITRU which is passed to the RSCS
| LOGON exit after a SIMLOGON has been issued.

| VTAM gets this bind image from the Logon Mode Tables associated with the remote system's VTAM.
| Through use of RSCS START command parameters, it is possible to specify which Logon Mode Table
| entry is returned to RSCS, but the Logon Mode Table which is searched must be specified when defining
| the remote system to it's local VTAM.

| B.3.7.3 Function Management Headers

| *Compaction:* RSCS does not support compaction. If an FMH3 is received RSCS will terminate the
| session.

| B.3.7.4 RU Composition

| *RU Size Determination:* The transmission block size is defined in RSCS START parameters for each link.
| This value is used in both the signon record, and FMH type 4.

| On a received FMH4, the transmission block size is checked to verify it is 300 or greater. Actual trans-
| mission block size is the smaller of the two values specified in the signon records.

B.3.8 BSC & CTC Support

B.3.8.1 CTCA Initialization

| RSCS Release 3 does not follow the procedure outlined in A.2.1, "Initialization" on page A-8. During this process, it always sends an SOH ENQ and will accept either another SOH ENQ or a DLE ACK0 as a valid response. It then sends the I signon record. If the other side replies with a non-valid response, RSCS will resend SOH ENQ. If the other side sends an I signon record, RSCS will ignore it and still send the I itself.

| In Version 2.1, the above restriction is eliminated.

B.3.8.2 Wait-a-bit Processing

RSCS Release 3 does not reply with a NAK when it gets data when it has set wait-a-bit. Any data sent in such a condition is lost. When RSCS receives wait-a-bit from the other side, it responds in its normal time sequence (i.e. no special delay) with a null buffer. There is no difference as to response time as to whether the wait-a-bit was received with data or a null buffer. When wait-a-bit is received, null buffers are sent until a DLE ACK0 sequence or a data buffer with wait-a-bit off is received.

| Version 2.1 will accept and process received buffers which have wait-a-bit set. It responds with a null buffer
| without any response time difference (i.e. no special delay).

B.3.8.3 Error Handling

Initialization Error Recovery: RSCS does not follow the procedure outlined in Figure A-4 on page A-3. It re-tries any type of error 15 times and only terminates after this many errors or if there is no I/O device defined for the line. All retries involve rewriting the SOH ENQ.

Normal (non-initialization) Error Recovery: RSCS Release 3 follows the actions in Figure A-5 on page A-4 except on command reject, bus out check, equipment check, data check, overrun and lost data, RSCS does not specifically diagnose the cause of the error. If this happens on a write, the data is resent. If it happens on a read, a NAK is sent.

| Version 2.1 follows the actions in Figure A-5 on page A-4 except for Data Check, Overrun, or Lost Data. If
| these errors happen on a write, the data is resent. If they happen on a read, a NAK is sent.

B.3.8.4 Use of Null Buffers

RSCS Release 3 uses null buffers only when wait-a-bit is set. It uses DLE ACK0 for positive acknowledgement when there is no data to send.

| In Version 2, Release 1 the above restriction is eliminated.

B.3.9 Accounting

B.3.9.1 Accounting Records

The VM/370 Network Accounting Card is recorded for all files transmitted and received by RSCS.

RSCS Network Accounting Card

0	RSCS LOCAL USERID		
4			
8	ACNTUSER		
C			
10	ACNTDATE		
	. . .		
18			
1C	ACNTOID		ACNTID
20	ACNTILOC		
24			
28	ACNTDEST		
2C			
30	ACNTCLAS	ACNTINDV	RESERVED
34	ACNTRECS		
38	ACNTTOVM		
3C			
40	RESERVED		
44			
48	ACNTSYS		
4C		ACNTCODE	RECID

The fields and bits are described below:

SYMBOL disp len type default range

RSCS 0 8 Character

Local network USERID fixed by CP

ACNTUSER 8 8 Character

Originating location USERID

ACNTDATE 16 12 Character

Data and time of record (MMDDYYHHMMSS)

ACNTOID 28 2 Binary

Origin spool file ID

ACNTID 30 2 Binary

Local spool file ID

ACNTILOC 32 8 Character

Originating location ID

ACNTDEST 40 8 Character

Destination location ID

ACNTCLAS 48 1 Character

Class

ACNTINDV 49 1 Character

Origin device type ('8N' = PUN/'4N' = PRT)

ACNTRECS 52 4 Character

Number of records in file

ACNTTOVM 56 8 Character

Destination location USERID

ACNTSYS 72 5 Character

System ID (Serial + Model)

ACNTCODE 77 1 Character

Transmission Code (01 = SEND/02 = RECV)

RECID 78 2 Character

Record identifier ('C0') fixed by CP

B.3.10 Miscellaneous Considerations

B.3.10.1 VM-Specific Line Drivers

RSCS Release 3 was designed so that there is different code used to communicate with non-VM/370 NJE systems from that used to communicate with other RSCS systems. The result of this structure is that certain information about jobs and NMRs which is carried in the non-VM/370 NJE code is lost when sent thru an RSCS node. The headers and trailers described in this document are *not* sent on a link which goes from one RSCS system to another when a file originates on one of the RSCS nodes. Headers and trailers are not processed on an end node which receives a file on an RSCS-RSCS link. (See diagram in section "Formatted Commands" on page B-25 for an example.)

| In Version 2, Release 1 the above restriction is eliminated. RSCS uses the same code to communicate with
| itself in this release as it does to communicate with non-VM systems.

B.3.10.2 User Exits

There are no user exits in RSCS to examine headers, trailers or transmitted files.

B.4 POWER

B.4.1 Supported Releases of POWER Networking

POWER currently supports POWER Version 2 Release 1 as the only networking system. The same networking support is also available in SSX Version 1 Release 2, as this includes POWER as an integral part of the system.

| Version 2 Release 2 (V 2.2) is the most current level. It contained some NJE improvements in message and command routing, increased transmission buffer sizes, and RAS enhancements.

B.4.2 Network Control

B.4.2.1 Network Connection & Control Records

POWER does not have a network path manager and supports only SRCB types I, J, and B.

| Any connection between POWER and JES2 must be predefined by a JES2 CONNECT statement.

B.4.2.2 Network Addressing, Topology and Routing

Naming Conventions: POWER supports up to 200 remotes (RJE stations) in Version 2.1, 250 in V 2.2, and both support an unlimited number of nodes. Remotes are defined during the POWER generation and are loaded during POWER initialization. Nodes are also pre-assembled and are loaded during the initialization phase.

Remote workstations are referenced by Rnnn, where nnn must be three numeric characters.

All members of a shared spool configuration must have the same node name in the network. They are distinguished by the node name qualifier.

Parallel Links: Multi-trunk is NOT supported.

Dynamic Route Changes: A new network definition table (NDT) can be loaded by the PLOAD operator command. This can be used to alter dynamically the network routing and topology.

Network Routing: POWER does not support the Path Manager of JES2. All routing information must be predefined by the user with the aid of the PNODE macro. A route table is loaded into storage when POWER is initialized. This table may be dynamically replaced while POWER is active by means of the PLOAD PNET operator command.

If a Job or OUTPUT is received from the network and the destination is not found in the Network Definition Table, then the Job/OUTPUT is put into 'HOLD' status in the transmission queue and POWER tries to send a message to the originating node and userid.

Alternate routing is supported by POWER. If the prime route is not active, but the secondary route is active, then the secondary route will be chosen. Both routes will NOT be used together, i.e. as soon as the prime route is available all new transmissions will be sent over the prime route and the secondary route will cease to be used as soon as the current transmission finishes.

B.4.2.3 Transmission Recovery Techniques

Jobs and OUTPUT data are handled on a store and forward basis. Responsibility for a unit of work does not pass from a transmitting node until it receives positive acknowledgment to the end of file which follows the Job Trailer Record. If transmission is broken off before that point, the receiver discards any data it may have received so far and the transmitter requeues its work for retransmission from the beginning.

Transmission may be discontinued voluntarily or involuntarily: Voluntarily by either the transmitter or receiver, and involuntarily in the case of a line disconnection. In all cases the receiving system discards the job or OUTPUT data that it has partially received. In the case of voluntary termination of transmission, the transmitting system will always requeue the job or OUTPUT data for retransmission.

If a line break occurs before the transmitter has sent end of file, the transmitter will requeue its work for transmission. If a line break occurs after the end of file is sent but before acknowledgment is received, the work is requeued in 'HOLD' status to avoid immediate retransmission. A message will be issued to inform the operator; he should ascertain whether the receiving system successfully received the entire job before releasing it or canceling it. If the job were to be automatically retransmitted in this case, it might result in duplicate jobs on the receiving system.

B.4.3 Commands & Messages (NMRs)

B.4.3.1 Commands

POWER builds an unformatted Nodal Message record (NMR) as a result of the PXMIT operator command. The syntax of the command is not checked by the transmitting system. The destination of the command is checked by the transmitting system and the command is rejected if the destination is not known.

B.4.3.2 Store and Forward Transparency

All commands and messages are only transmitted to their final destination if there is a path available. If the prime route is not available then any alternate route may be chosen. If no route is available then commands are thrown away but a message is sent to the originator informing them of the inability to deliver the command. Messages are thrown away if the route is not available without informing the originator.

B.4.3.3 Formatted (Global) Commands

Formatted commands are not supported by POWER

B.4.3.4 Command Authorization

POWER supports three types of command authorization, which must be specified on a node basis in the node definition table. The three levels are:

- Network - node has the authority to control a lot of functions that the local operator can control. Transmitters/ receivers can be started remotely and transmission flushed. No control is possible for local I/O devices.

- JOB - node can only manipulate jobs or output which either originated with or are destined for that node.

- NOJOB - node is not allowed to manipulate anything within the network. Users are still allowed to do displays of queues on other nodes.

B.4.3.5 Message transmission

Messages are created by POWER in the following instances:

- In response to a unformatted command
- Via a PBRDCST command issued either by the operator or by any interactive user having authority.
- For notification at end of execution of a job or on receipt of jobs or output.

B.4.4 SYSIN (Job) Transmission

B.4.4.1 Store and Forward Transparency

All sections in the job header are forwarded without change. As an intermediate node, no sections are added to the job header. The operator at the S&F node can change the execution node of the job and this will be updated in the job header.

B.4.4.2 Job Header

The job header is created as soon as the job enters the POWER system. The header is used to store job specific data, e.g. execution destination, default print and punch routes, etc.. Segmented job headers are supported up to the maximum allowed by the architecture.

POWER Section: The POWER section contains a field specifying the disposition of the job as specified in the * $$ JOB statement. This is required because POWER allows additional dispositions to JES2/JES3. The section also contains user information from the * $$ JOB statement.

B.4.4.3 Job Trailer

The job trailer is built as soon as the job enters the POWER system.. Since POWER doesn't use a lot of the counts that JES2/JES3 use for scheduling purposes, the majority of the fields in this record are not used and are set to defaults.

B.4.4.4 Data Set Header

Record Characteristics Change Section: This section is put into the job stream as soon as POWER determines that the stream may contain records greater than 80 bytes. This means that when a SYSIN stream is read from a diskette with a header that says 128 bytes, this section is sent directly after the Job header. There is also a POWER section appended to this section to allow the transmission of the diskette address used when reading the file.

B.4.4.5 Acceptable Job Streams

POWER cannot accept multiple POWER jobs between a job header and job trailer. A POWER job can consist of multiple job steps.

B.4.4.6 Notification

When a job is first read in by POWER, the node and userid for notification are stored in the Job Header Record if 'NTFY=' was specified on the JOB card. This means that if a node, other than the originating node, is to be notified, then the originating nodeid is destroyed by the notify node name.

If the job header contains a notify userid and the job is transmitted from its origin node for execution on another node, the job transmitter of all intermediate nodes will issue a message to the user indicating that the job has been transmitted.

When the job completes execution, the execution processor on the execution node will issue a notification message to the user specified by the notify userid field and to the node specified by the origin node field.

B.4.4.7 SYSIN Job Routing

Use of System Qualifier: If the job is entered on a POWER system and will be executed on a POWER system, then there is the capability to ensure that a certain system in a shared-spool complex executes the job. The system qualifier is specified in the * $$ JOB statement and is contained in the POWER section.

Use of Userid: The userid field on the XDEST parameter can be used to route a job to a VM virtual machine for execution. If POWER is the target execution node, then the field NJHGXEQU is ignored.

Operator Rerouting: POWER operators may change the execution node and userid for any job on the POWER queue by use of the PALTER command.

B.4.4.8 Assignment of JOBID (Job Name and Job Number)

In POWER the JOBID is a combination of Job Name and Job Number. The POWER job number is a halfword binary number assigned by POWER when the job first enters the system. This number may not be unique within the POWER system. POWER uses the job name as the primary identifier for all system control, and the job number as a secondary delimiter.

When a job is transmitted from one system to another the receiving system will assign the original jobname (from the Job Header) to the job that is being received. A new job number is always assigned to received jobs or OUTPUT. The new and original job numbers are displayed in the messages which are issued to acknowledge transmission of jobs or OUTPUT to the next node. The Job Header always contains the original (input system's) JOBID. The JOBID assignment procedure outlined above will be followed by both job receivers and OUTPUT receivers.

If the OUTPUT has been segmented, then the job name will remain the same as the JOB and each segment will retain the original job number.

B.4.5 SYSOUT (Job Output) Transmission

B.4.5.1 NJE Unit of Transmission

Within NJE a unit of transmission is defined as being everything between a Job Header and a Job Trailer record. It may consist of several Data Sets of different characteristics. Within POWER spooling it is necessary to differentiate between PRINT and PUNCH output, because there are separate queues for each type. This means that if output is received containing mixed OUTPUT types, i.e. PRINT and PUNCH data it will be split.

B.4.5.2 Store and Forward Transparency

All SYSOUT data sets are forwarded through POWER transparently. A maximum length of spanned record as defined in the architecture is supported. No sections are added to the record by POWER although the user may decide to add his own section.

Spanned Records with Carriage Control: POWER supports spanned records with carriage control both as intermediate node and as final destination.

Trailing Blank Truncation. Trailing blanks are truncated from all data both on spool and on transmission. The original record length is however retained and on receipt of a file, the blanks are appended.

Print File Transparency: POWER Version 2 Release 1 converted all ASA carriage control to machine carriage control. Future releases of POWER will not convert any data before it goes to the physical device.

B.4.5.3 Job Header

The Job Copy count field of the job header is not used by POWER. This means that a user should use other means of specifying that he wishes to print multiple copies of his output. The default print and punch destination and userid fields in the job header are not used for SYSOUT as their information has already been copied to the data set headers.

B.4.5.4 Job Trailer

The majority of the fields in the job trailer are not used by POWER and are set to default values.

B.4.5.5 Data Set Header

The Data Set Header contains the information necessary to handle correctly OUTPUT data sets on the receiving system. It contains three types of information:

- Identification (data set number)
- Routing Control (destination node name and remote)
- Data Set Characteristics (output class, copy count, FCB name)

Data Set Headers exist by POWER as soon as any job has executed on a POWER node or OUTPUT is received from another node. They are built from information contained in the Job Header and in the POWER queue record, which in turn is built from JECL statements, if present. The Data Set Header is always present on POWER spool, even if the data set is not intended for transmission. The Data Set Header record is flagged as an internal record so that is will be ignored by local print/punch processors and external writers on the destination node. Certain information is retrieved from the Data Set Header record by the print/punch processor when the data set is finally printed/punched.

POWER does not mix OUTPUT data types in one transmission stream. This means that if there is Print and Punch output to be sent, then this will be transmitted as two distinct transmissions, i.e. with two Job Headers/Job Trailers.

3800 section: POWER will create a 3800 section for all output for which a 3800 device had been specified at execution time. The values in the section will be taken from the values specified in the * $$ LST statement or from the SETPRT defaults.

Note: In VSE there are two sets of defaults for the 3800, the hardware defaults and the software defaults. Unfortunately a value of X'00' in VSE means use hardware defaults and X'40' means use software defaults whereas in JES2/JES3 X'00' means use software defaults and hardware defaults are unknown. This has meant that when POWER is the end destination, it must determine whether the file originated at a POWER node and when that is not the case, the defaults are converted.

Multiple Data Set Headers: POWER does not create any multiple destination data sets. The receipt of multiple headers with no OUTPUT data in between indicates to the OUTPUT receiver that the data set has multiple destinations, and it will build multiple queue entries accordingly. When the data set is transmitted again it will be transmitted as multiple copies of headers and data.

POWER always splits jobs when they are received, and never generates multiple dataset headers.

B.4.5.6 Notification

When a job's system output (or any part of it) reaches its ultimate destination, the OUTPUT receiver on the destination node will examine the userid in the job header. If 'NTFY=' was specified on the JOB card, the OUTPUT receiver will issue a message to the specified user indicating that the job's OUTPUT was received.

B.4.5.7 Job Output Routing

Default Output Routing: Unless otherwise specified on control statements, the default print and punch nodes are set the the origin node. The information is taken from the job header record.

Operator Rerouting: The operator can reroute any SYSOUT that is currently resident in his queue. Both the node and the userid can be changed.

Fan-Out (Optimized) Support: POWER does not support fan-out. Any input stream containing multiple data set headers will be broken down into streams consisting of single datasets using the criteria specified below:

POWER will keep all data sets belonging to a transmission together provided certain important criteria do not change. The criteria are as follows:-

- Target Node and Userid
- Type of output, list or punch
- FCB name
- Priority
- Output Class
- Copy Count
- Form name
- 3800 Characteristics (Copy Group, Burst, Copy Index)

If any of these characteristics change within one unit of transmission, then a new spool entry is built together with Job header and Data set header. This is necessary because the operator has the possibility to change the destination of any Job or OUTPUT which is in the spool file and make its destination the local node. This means that all entries on spool are capable of being processed locally.

B.4.6 Stream Support and Control

POWER supports up to 7 job transmitters, 7 SYSOUT transmitters, 7 job receivers and 7 SYSOUT receivers per network line/session. The total number of job and SYSOUT streams concurrently active on any line/session cannot exceed 8. The operator has the capability to start and stop each individual stream. A console stream is always defined for every line/session and cannot be stopped by the operator.

B.4.6.1 Stream Initiation & Suspension

Operator Control (of Streams & Lines): The operator can start and stop any stream on any line/session. If the operator has network authority he can also manipulate the transmitters and receivers on another connected node. The PACT and PDRAIN commands are used for this purpose.

NJE Tasks: When the line is initialized, all receivers are set in active status, but only JOB transmitter 1 and OUTPUT transmitter 1 are set in active status. If the VSE operator wishes to activate more transmitters he may do so with the aid of the PACT command. The operator must enter:-

- PACT TRn,JOB|OUT

  where n is a number from 1-7 and corresponds to transmitter 1 to 7. The specification of JOB or OUT is necessary so that we internally may set the correct RCB and FCS bits and inhibit the starting of invalid combinations of JOB and SYSOUT transmitters, e.g. JOB transmitter 5 and SYSOUT transmitter 4.

Under POWER receiver tasks live only for the time span of one transmission, i.e. for everything between Job Header and Job Trailer. After the transmission is complete, i.e. the End-of-File record has been received and acknowledged, the receiver task detaches. The task will be created again when another RIF for the same RCB is received. The detaching of tasks saves resources and can be critical in small systems.

Transmitter tasks, on the other hand, live for as long as there is anything in the transmit queue for the Node they are serving. As soon as there is nothing eligible in the queue, the task detaches and will be created again only when JOB or OUTPUT is placed in the transmission queue for this node.

B.4.7 SNA Support

B.4.7.1 Data Flow Control

The only Data Flow Control (DFC) function supported by POWER is RSHUTD. If any other function is received, POWER will terminate the session with a CLSDST or a TERMSESS.

B.4.7.2 Bind Parameters

POWER does not negotiate any of the bind parameters. They are "hard-coded" and most of the parameters received by the other node are ignored.

B.4.7.3 Functional Management Headers

The session will be terminated if the FMH received by POWER is not acceptable.

Compaction: POWER does not support compaction for NJE transmissions. The fix for APAR DY33141 allows POWER to accept and handle an FMH3, and decompacted inbound data. POWER will still not be able to compact outbound data.

B.4.7.4 RU Composition

RU Size Determination: The transmission block size is is defined in the network definition table. This value is used in the SIGNON record that is exchanged with the communicating node. The value that is received from the partner node is compared with the defined value and the smaller of the two values is taken. For SNA the maximum size allowed by POWER is 32,000.

RU Multiplexing: POWER sends only 1 type of record in an RU and is not able to receive RU's containing mixed record types.

B.4.8 BSC & CTC Support

B.4.8.1 BSC

Buffer Size: For BSC, the maximum transmission buffer size allowed by POWER Version 2.2 is 1800.

B.4.8.2 Error Recovery

Initialization Error Recovery: POWER will retry on Bus Out Check and Equipment Check before terminating.

B.4.8.3 CTC

The channel-to-channel adapter is not supported by DOS/VSE.

B.4.9 Accounting

See page 4-7 in the "VSE/POWER Networking User's Guide", SC33-6140.

B.4.10 Miscellaneous Considerations

B.4.10.1 User Exits

See page 4-7 in the "VSE/POWER Networking User's Guide", SC33-6140.

Appendix C. NJE Bibliography

The following list of publications is a subset of the various product documentation libraries including some *IBM Systems Journal* reprints and Systems Center Technical Bulletins. For a complete list of product documentation, see the General Information Manuals listed under the appropriate product heading, or *IBM System/370 Bibliography*, GC20-0001.

C.1.1 NJE - General

- *Job Networking (IBM Systems Journal V17 N3, 1978)*, G321-5071
- *Job Networking Facilities*, GG22-9042
- *NJE Concepts and Protocols Overview*, GG66-0224

C.1.2 Systems Network Architecture (SNA)

- *Concepts and Products*, GC30-3072
- *Technical Overview*, GC30-3073
- *Sessions Between Logical Units*, GC20-1868
- *Reference Summary*, GA27-3136
- *Format and Protocol Reference Manual: Architectural Logic* , SC30-3112
- *IBM Synchronous Data Link Control General Information* , GA27-3093

C.1.2.1 ACF/VTAM

- *General Information*, GC27-0608
- *Planning and Installation Reference*, SC27-0610
- *Programming*, SC27-0611
- *Operation*, SC27-0612
- *Messages and Codes*, SC27-0614
- *Diagnosis Guide*, SC27-0615
- *Diagnosis Reference*, LY38-3053 (OS/VS), LY38-3058 (VSE)
- *Data Areas*, LY38-3054 (OS/VS), LY38-3059 (VSE)

C.1.3 Binary Synchronous Communications (BSC)

- *General Information*, GA27-3004

C.1.4 Channel-to-Channel Adapter (CTC)

- *S/370 Special Feature Description: CTCA*, GA22-6983

C.1.5 JES2

- *Network Job Entry (IBM Systems Journal V17 N3, 1978)*, G321-5072
- *Running JES2/NJE on an ACF/VTAM Network*, G320-5855
- *How JES2 Uses SNA for RJE and NJE*, GG22-9378

C.1.5.1 MVS/SP JES2 - Version 1 (5740-XYS)

- *General Information*, GC28-1025
- *JES2 Initialization and Tuning*, SC23-0046
- *JES2 Commands*, SC23-0048
- *JES2 Command Language Reference Summary*, SX23-0008
- *System Messages Volume 1*, GC28-1374
- *System Messages Volume 2*, GC28-1375
- *JES2 Messages*, GC28-1354
- *MVS JCL User's Guide*, GC28-1349
- *MVS JCL Reference*, GC28-1350
- *JES2 Modifications and Macros*, LC23-0067
- *JES2 Logic*, LY24-6006

C.1.5.2 MVS/SP JES2 - Version 2 (5740-XC6)

- *General Information*, GC28-1118
- *JES2 Initialization and Tuning*, SC23-0065
- *JES2 Commands*, SC23-0064
- *JES2 Command Syntax*, SX23-0010
- *System Messages Volume 1*, GC28-1376
- *System Messages Volume 2*, GC28-1377
- *JES2 Messages*, GC28-1353
- *MVS JCL User's Guide*, GC28-1351
- *MVS JCL Reference*, GC28-1352
- *JES2 Modifications and Macros*, LC23-0069
- *JES2 Logic*, LY24-6008

C.1.6 JES3

C.1.6.1 MVS/SP JES3 - Version 1 (5740-XYN)

- *General Information*, GC28-1025
- *JES3 Initialization and Tuning*, SC23-0041
- *JES3 Operator's Library*, SC23-0045
- *JES3 Operator's Reference Summary*, SX23-0007
- *MVS JCL User's Guide*, GC28-1349
- *MVS JCL Reference*, GC28-1350
- *JES3 Messages*, GC23-0044
- *System Messages Volume 1*, GC28-1374
- *System Messages Volume 2*, GC28-1375
- *JES3 Modifications and Macros*, LC28-1371
- *JES3 Diagnosis*, LC28-1369
- *JES3 Logic - Volume 1*, LY28-1507
- *JES3 Logic - Volume 2*, LY28-1509
- *JES3 Logic - Volume 3*, LY28-1511
- *JES3 Logic - Volume 4*, LY28-1513
- *JES3 Logic - Volume 5*, LY28-1515
- *JES3 Logic - Volume 6*, LY28-1517
- *JES3 Logic - Volume 7*, LY28-1519
- *JES3 Logic - Volume 8*, LY28-1521
- *JES3 Logic - Volume 9*, LY28-1523
- *JES3 Logic - Volume 10*, LY28-1525
- *JES3 Logic - Volume 11*, LY28-1527

C.1.6.2 MVS/SP JES3 - Version 2 (5665-291)

- *General Information*, GC28-1118
- *JES3 Initialization and Tuning*, SC23-0059
- *JES3 Commands*, SC23-0063
- *JES3 Command Syntax (Reference Card)*, SX23-0012
- *MVS JCL User's Guide*, GC28-1351
- *MVS JCL Reference*, GC28-1352
- *JES3 Messages*, GC23-0062
- *System Messages Volume 1*, GC28-1376
- *System Messages Volume 2*, GC28-1377
- *JES3 Modifications and Macros*, LC28-1372
- *JES3 Diagnosis*, LC28-1370
- *JES3 Logic - Volume 1*, LY28-1529
- *JES3 Logic - Volume 2*, LY28-1531
- *JES3 Logic - Volume 3*, LY28-1533
- *JES3 Logic - Volume 4*, LY28-1535
- *JES3 Logic - Volume 5*, LY28-1537
- *JES3 Logic - Volume 6*, LY28-1539
- *JES3 Logic - Volume 7*, LY28-1541
- *JES3 Logic - Volume 8*, LY28-1543
- *JES3 Logic - Volume 9*, LY28-1545
- *JES3 Logic - Volume 10*, LY28-1547
- *JES3 Logic - Volume 11*, LY28-1549

C.1.7 VM/370

- *Evolution of Virtual Machine Subsystem (IBM Systems Journal, V18, N1, 1979)*, G321-5089

C.1.7.1 RSCS Networking Version 1 (5748-XP1)

- *General Information*, GH24-5004
- *Program Reference & Operations Guide*, SH24-5005
- *Program Logic Manual*, LY24-5203
- *Operators Reference Summary Card*, SX24-5119

C.1.7.2 RSCS Networking Version 2 (5664-188)

- *General Information*, GH24-5055
- *Operation and Use*, SH24-5058
- *Planning and Installation*, SH24-5057
- *Diagnosis Reference*, LY24-5228

C.1.8 DOS/VSE

C.1.8.1 VSE/POWER Version 2 (5666-273)

- *General Information*, GH12-5131
- *Installation & Operations Guide*, GH12-5329
- *Networking User's Guide*, SC33-6140
- *Installation and Operations Guide*, SH12-5329
- *Reference Summary*, SH12-5435
- *Messages*, SH12-5520
- *Diagnosis Reference Manual*, LY12-5027
- *Networking Design Guide*, GG24-1570

Glossary and Abbreviations

This glossary defines NJE terms and other data processing terms used in this publication. For definitions of terms not included in this glossary, see *IBM Vocabulary for Data Processing, Telecommunications, and Office Systems*, GC20-1699.

A

ACF/VTAM. Advanced Communications Facility/Virtual Terminal Access Method

ACK (or ACK0,ACK1). In BSC, an affirmative acknowledgement, indicating that the previous block was accepted without error, and the receiver is ready to accept the next block of transmission.

adjacent. In NJE, two nodes are said to be "adjacent" if they are connected directly by a single link or session.

AFP. Advanced Function Printing

ASA. American Standards Association

ASP. Attached Support Processor, also called Asymmetric Multiprocessing Support - The predecessor to JES3.

B

BCB. See "block control byte".

BCB sequence error. This refers to a multi-leaving record beginning with an RCB of X'E0'. This record indicates that an error was detected in the BCB of the previous block. At this time, data has been lost, and the only correct response is to terminate the line and restart it.

BIND. In SNA, a request to activate a session between two logical units.

block control byte (BCB). A byte used to maintain the integrity of data during a multi-leaving transmission.

BSC. Binary Synchronous Communication, or Bisynchronous Communication - Communication using binary synchronous line discipline in which transmission of binary-coded data between stations is synchronized by timing signal generated at the sending and receiving stations.

C

CCTL. Carriage Control. The first character of an output record which indicates the placement of that record in the print or punch medium.

CCW. Channel Command Word

CES. See "connection event sequence".

channel to channel adapter (CTC or CTCA). A feature on S/370 channels that allows two processors to communicate directly to one another. It is described in *IBM S/370 Special Feature Description: Channel-to-Channel Adapter*, GA22-6983. Also see the IBM 3088.

compaction. A method of reducing the length of records for transmission by representing certain 8 bit characters with only 4 bits.

compression. A method of reducing the length of records for transmission by removing blanks and duplicate characters.

connection event sequence (CES). A number in each path manager record, based on the current TOD clock value (in GMT), to prevent the use of redundant path manager records by the Network Path Manager.

control record. All records with an RCB with the low four bits of zero or an SRCB with the two high order bits of one.

CPDS. Composed Page Data Stream. An architected set of control codes used to control page printers running in all-points-addressable (i.e., page) mode.

CTC. See "channel to channel adapter".

D

data record. In NJE, this refers to all records with an RCB in the range 98-F8 and the range 99-F9 which do not have an SRCB with the two high order bits set to one.

data set header (NDH or DSH). An NJE control record that generally precedes a unit of SYSOUT data. It may also appear in the middle of SYSIN data to indicate a change in the format of the SYSIN data.

DBCS. Double-Byte Character Set. A method of representing information, in which each character to be printed is represented by two bytes. This contrasts with EBCDIC, in which each character to be printed is represented by one byte.

decompaction. The process of restoring a compacted record to its original form.

decompression. The process of restoring a compressed record to its original form.

DLE. Data Link Escape - a BSC control character used to provide supplementary line control characters.

DOS/VSE. Disk Operating System/Virtual Storage Extended

E

end of block (EOB). A multi-leaving record control block (X'00') indicating the termination of a transmission block.

end of file (EOF). This refers to a null record that is transmitted at the end of a job, following the job trailer, to indicate that nothing remains to be transmitted. It is acknowledged by a Stream Complete.

end of record (EOR). A multi-leaving string control block (X'00') indicating the termination of a logical record.

end of transmission block (ETB). A BSC control character indicating the end of a block of characters started with a SOH or STX.

end of transmission (EOT). A BSC control character indicating the end of a message transmission.

ENQ. Enquiry - a BSC control character used to bid for a line, or obtain a repeat transmission of the response to a message.

EOB. See "end of block".

EOF. See "end of file".

EOR. See "end of record".

F

fan-out. The ability in NJE to send a SYSOUT data set to multiple destinations without sending multiple copies down the same link.

FCS. See "function control sequence".

file. A collection of records represented by a job or a data set.

FMH. See "function management header".

function control sequence (FCS). Control bytes used to manage streams in multi-leaving transmissions.

function management header (FMH). In SNA, specialized control format to select a destination and control the way data is sent or presented at the destination.

H

HASP. Houston Automated Spooling Priority - The predecessor to JES2.

J

JCL. See "job control language".

JECL. See "job entry control language".

JES2. Job Entry Subsystem 2. An MVS system facility for spooling, job queuing, and managing I/O.

JES3. Job Entry Subsystem 3. An MVS system facility for spooling, job queuing, and managing I/O.

JES3 global. In a JES3 complex, the processor responsible for managing the complex, i.e. Controlling all input and output, insuring complex-wide data integrity, etc. The global also handles all networking functions.

JES3 local. A processor in a JES3 complex whose primary function is to run MVS jobs. Locals access the global for various services.

job. A job is a unit of work within the network. It consists of all data beginning with a job header control record and ending with a job trailer control record.

job control language (JCL). In OS/VS, an esoteric command language used to specify batch work.

job entry control language (JECL). Specialized control language, interspersed in JCL, read by the job entry subsystem (JES2, JES3, or VSE/POWER).

job header (NJH). The NJE control record that provides general information relating to the job as a whole.

job network. A collection of peer-coupled systems connected by communication links, using NJE protocols.

job trailer (NJT). The NJE control record the terminates the job and generally provides accounting information.

L

link. A connection, or ability to communicate between two adjacent nodes

logical unit (LU). In SNA, a port through which an end user or application accesses the network in order to communicate with another end user or application.

LRECL. Logical Record Length

LU_0. Logical Unit Type 0. In SNA, a port through which two applications communicate using implementation-defined protocols. NJE uses such protocols for SNA transmissions.

M

multi-leaving (ML). Fully synchronized two-directional transmission of a variable number of data streams between terminals and/or computers using BSC facilities.

MVS. Multiple Virtual Storage - an alternative name for OS/VS2.

MVS/SP. Multiple Virtual Storage/System Product

N

NAK. In BSC, a Negative Acknowledgement, indicating that the previous block was received in error, and the receiver is ready to accept a retransmission of the erroneous block.

NCC. See "network connection control".

NDH. Network Data Set Header - see Data Set Header

network connection control (NCC). An NJE control record used to establish or break a connection. Also used by the network path manager to send information about other NJE connections.

network job entry (NJE). (1) A facility for transmitting jobs (JCL and in-stream data sets), sysout data sets, (job-oriented) operator commands and operator messages, and job accounting information from one computing system to another. (2) A facility that provides access to batch computing facilities from other host systems. It enables users to transfer work and data throughout a distributed network of batch computing facilities. ("NJE" is not a part of "Systems Network Architecture (SNA)", but is an application layer which uses SNA, BSC and CTC transmission facilities.) (3) The JES2 program product implementation of the NJE Protocol.

network job interface (NJI). The original HASP, RSCS, ASP, or JES3 Programming RPQ implementation of the NJE Protocol.

network path manager (NPM). A facility to manage the topology of a network through the sending, receiving and processing of network connection and control (NCC) records.

NJE. See "network job entry".

NJH. Network Job Header - see "job header".

NJI. See "network job interface".

NJT. Network Job Trailer - see "job trailer".

NMR. See "nodal message record".

nodal message record (NMR). A record for transmitting commands and messages to other locations.

node. A computer participating in an NJE network.

nodeid. Node identifier. The name by which a node is known to all other nodes in a network.

notification. The sending of a message to an interactive user of an event associated with the processing of a job.

NPM. See "network path manager".

null buffer. A transmission buffer containing only an RCB of X'00'.

O

OPTB. Output Parameter Text Block - an optional section in the data set header.

P

path management record. In NJE, this refers to a record beginning with an RCB of X'F0'. It is a non-compressed record and contains network connectivity information.

permission granted. In NJE, this refers to a record with an RCB of X'A0'. It is used following a Request Permission to Initiate Stream when the receiving system is willing to accept the new stream.

permission rejected. In NJE, this refers to a record with an RCB of X'B0'. It is used following a Request Permission to Initiate Stream when the receiving system is not willing to accept the new stream. This same code also goes under the name of Receiver Cancel, and is used whenever the receiver wishes to cancel the job that is being received.

POWER. "Priority Output Writers, Execution Processors and Input Readers" - a spooling subsystem for VSE systems

R

RCB. See "record control byte".

receiver cancel. In NJE, this refers to a record with an RCB of X'B0'. This is used after a Permission Granted has been transmitted and the receiver wishes to terminate the job on a stream.

record. In NJE, this is defined as all bytes beginning with an RCB up to but not including the next RCB. This term encompasses Control Records, data records, and nodal message records.

record control byte (RCB). The byte that defines the stream for each record within a transmission buffer.

record identifier (RID). The Record Identifier used in SNA transmissions. It is a three byte field made up of the RCB, SRCB, and a byte containing the length of the data record.

request permission to initiate stream. In NJE, this refers to a record with an RCB of X'90'. It is used prior to transmitting a stream of data. The other system will either respond with a Permission Granted, or a Permission Rejected.

request/response unit (RU). In SNA, an element of the Basic Link Unit containing data and data stream controls.

RID. See "record identifier".

RJE. See "remote job entry".

remote job entry (RJE). Submitting a job, and receiving output, through an I/O device that is connected to a computer via communications equipment.

RSCS. Remote Spooling Communications Subsystem - A program product for VM/SP, it is a special-purpose subsystem that supports the reception and transmission of messages, files, commands, and jobs over a computer network.

RU. Request/Response Unit

S

SAF. Store and Forward

SCB. See "string control byte".

scheduler work block (SWB). A structured data block containing the data produced by subsystem JCL verification and internal processing.

SCIP. Set Control Interval Processing

SEGL. Segment Length

SNA. (1) (Noun) See "systems network architecture". (2) (Adjective) Adhering to the Systems Network Architecture definitions. "SNA" is used to describe links, I/O devices, telecommunication controllers, etc.

SOH. See "start of heading".

spool. Simultaneous Peripheral Operation On-Line (1) (Noun) An area of auxiliary storage defined to temporarily hold data during its transfer between peripheral

equipment and the processor. (2) (Verb) To use auxiliary storage as a buffer storage to reduce processing delays when transferring data between peripheral equipment and the processing storage of a computer.

spoolid. Spool file identifier. A number between 1 and 9900 that is assigned to a spooled file by VM spooling facility.

SRCB. See "sub-record control byte".

start of heading (SOH). In BSC, a character preceding a block of heading characters.

stream. A logical flow of information.

stream complete. In NJE, this refers to a record beginning with an RCB of X'C0'. This is sent by the receiver after the transmitter sends an EOF. It is at this point that the transmitter may purge the job from its queue.

string control byte (SCB). A byte within the data stream that is used in compression algorithms.

sub record control byte (SRCB). Defines individual types of records within an RCB.

SWB. See "scheduler work block".

SYSIN. SYSIN refers to a type of job which is intended to be processed as an execution type job by the operating system at the receiving node. After processing, SYSOUT is usually produced which is usually returned to the origin node. It is possible for a job to execute and produce SYSIN to be executed at the same or another node.

SYSOUT. SYSOUT refers to the output from some program. When received by the networking component at the destination, it is not inserted into the execution job queue. It may be printed or punched immediately on a locally connected output device, or placed into a state from which a user or operator of the system may specify its further processing.

systems network architecture. A formal definition of the format, protocol, and sequencing of information that is sent through a network.

T

transmission block. A collection of one or more records to be transmitted over the network as a unit. Transmission blocks are that portion of each transmission that is independent of the access method.

transmission buffer. An area in storage for building and receiving transmission blocks.

transmitter abort. In NJE, this refers to a record with an SCB of X'40'. It is used when a transmitter wishes to abort the job being transmitted.

TSO. Time Sharing Option

TSO/E. Time Sharing Option/Extensions

TSS. Time Sharing System

U

unit control word (UCW). The word in a block multiplexor channel that is used to save the address for channel reconnect. A separate word must be defined for each CTCA for proper channel sharing.

userid. User identifier. (1) (VM) the name by which a virtual machine and its user are known to others. (2) (MVS) the name by which a time sharing user is known to others.

V

VM. Virtual Machine

VM/370. Virtual Machine/370

VSE. "Virtual System Extended". Another name for DOS/VSE.

VSE/POWER. "Virtual System Extended/Priority Output Writers, Execution Processors and Input Readers". A spooling system for VSE.

VTAM. Virtual Telecommunications Access Method. A program product that controls communication and the flow of data in a computer network. It provides single-domain, multiple-domain, and multiple-network capability. VTAM runs under OS/VS1, MVS, VSE and VM/SP.

W

workstation. An I/O device from which jobs can be submitted to a host for processing, and/or to which output can be returned.

Index

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

Z

Numerics

READER'S COMMENT FORM

Title:

Network Job Entry Formats and Protocols
for System/370 Program Products
Washington Systems Center
Technical Bulletin GG22-9373-02

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Please state your occupation: _____

Comments:

Please mail to:

J. M. Hutchinson
Washington Systems Center
IBM Corporation
18100 Frederick Pike
Gaithersburg, MD 20879

Reader's Comment Form

Cut or Fold Along Line

Fold and tape | Please Do Not Staple | Fold and tape

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE:

J. M. Hutchinson
Washington Systems Center
IBM Corporation
18100 Frederick Pike
Gaithersburg, MD 20879

Fold and tape | Please Do Not Staple | Fold and tape

IBM®

READER'S COMMENT FORM

Title:

Network Job Entry Formats and Protocols
for System/370 Program Products
Washington Systems Center
Technical Bulletin GG22-9373-02

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Please state your occupation: _____

Comments:

Please mail to:

J. M. Hutchinson
Washington Systems Center
IBM Corporation
18100 Frederick Pike
Gaithersburg, MD 20879

Cut or Fold Along Line

Reader's Comment Form

Fold and tape | Please Do Not Staple | Fold and tape

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE:

J. M. Hutchinson
Washington Systems Center
IBM Corporation
18100 Frederick Pike
Gaithersburg, MD 20879

Fold and tape | Please Do Not Staple | Fold and tape

IBM®

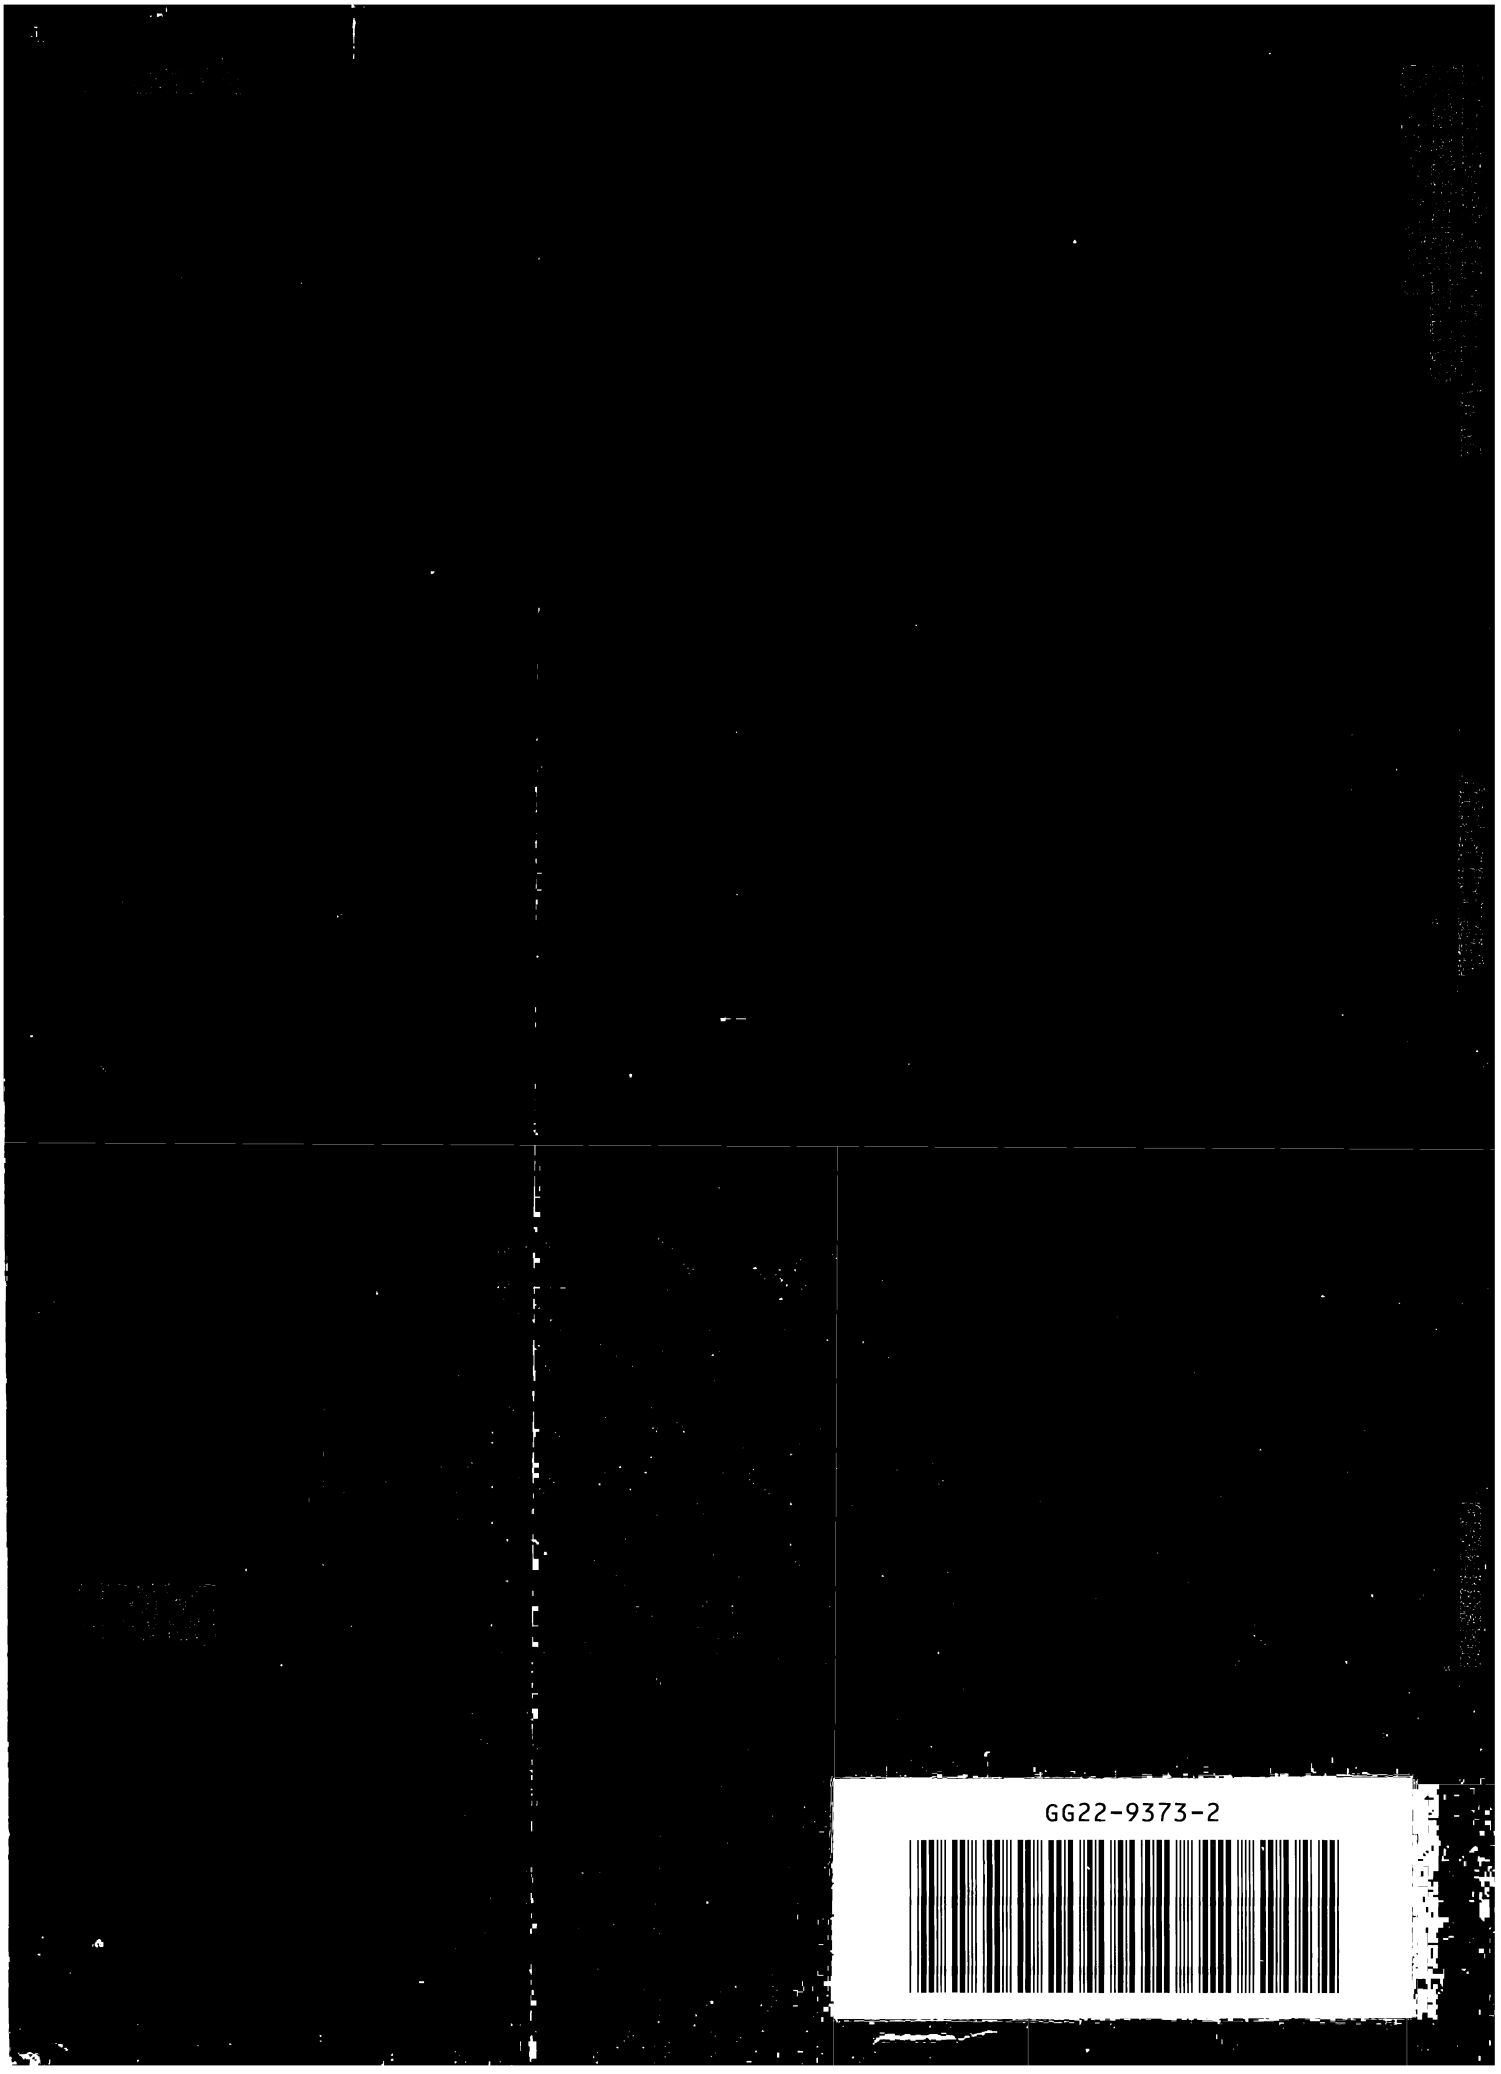
GG22-9373-2

GG22-9373-02

Network Job Entry Formats and Protocols for System/370 Program Products

Printed in U.S.A.

GG22-9373-02

IBM