



Formats and Protocols

[The following text is extremely faint and illegible due to low contrast and scan quality. It appears to be a list of items or a table of contents.]



Network Job Entry

SC23-0070-0

Formats and Protocols

First Edition (November, 1990)

This edition applies to the program products identified in the Preface (which also specifies versions, releases, and program numbers) unless otherwise indicated in new editions or Technical Newsletters. Changes are made periodically to the information herein; before using this publication in connection with the operation of IBM systems, consult the *IBM System/370 Bibliography*, GC20-0001, for the editions that are applicable and current.

References in this publication to IBM products or services do not imply that IBM intends to make these available in all countries in which IBM operates. References to IBM products in this document do not imply that functionally equivalent products may be used. The security certification of the trusted computing base that includes the products discussed herein covers certain IBM products. Please contact the manufacturer of any product you may consider to be functionally equivalent for information on that product's security classification. This statement does not expressly or implicitly waive any intellectual property right IBM may hold in any product mentioned herein.

Publications are not stocked at the address given below. Requests for IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

A form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, Information Development, Department D58, Building 921-2, PO Box 390, Poughkeepsie, N.Y. 12602. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1990. All rights reserved.

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. Introduction to Network Job Entry (NJE)	1-1
Units of Work	1-3
Functions of a Node	1-3
Types of Nodes	1-4
Types of NJE Users	1-6
NJE Addressing	1-6
NJE Protocols	1-7
Formats of NJE Transfer Units	1-8
Presentation Services	1-10
Stream Control	1-14
NJE Connection Control	1-15
Link Protocols	1-16
Chapter 2. Formats of NJE Transfer Units	2-1
NJE Job	2-2
Contents of an NJE Job	2-2
Job Header	2-3
Data Set Header	2-21
Job Trailer	2-54
NJE Nodal Message Records (NMR)	2-58
Identifying the Contents of an NMR	2-58
Format of an NMR that Contains an Unformatted Command	2-59
Format of an NMR that Contains a Formatted Command	2-61
Format of an NMR that Contains a Message	2-63
Networking Connection Control Records (NCCR)	2-67
Initial Signon Record	2-67
Response Signon Record	2-68
Reset Signon Record	2-69
Concurrence Signon Record	2-69
Add/Subtract Connection Record	2-70
Signoff Records	2-70
Chapter 3. Presentation Services	3-1
Non-SNA Buffer Format	3-2
Transmission Buffer Size	3-2
Compressing Data for BSC and CTC Transmissions	3-2
BSC and CTC Buffers	3-2
Block Control Byte (BCB)	3-4
Function Control Sequence (FCS)	3-7
String Control Byte (SCB)	3-8
Record Control Byte (RCB)	3-9
Subrecord Control Byte (SRCB)	3-11
SNA Buffer Format	3-13
Transmission Buffer Size	3-13
Compressing and Compacting Data for SNA Transmissions	3-13
SNA Transmission Buffer	3-13
String Control Byte (SCB)	3-14
Record Identifier (RID)	3-15
Data Records	3-17
Format of Data Records Containing Unspanned Data	3-17
Format of Data Records Containing Spanned Data	3-18
Trailing Blank Truncation	3-19

Chapter 4. Stream Control	4-1
Initiating Stream Transmission	4-2
BSC and CTC Protocols for Stream Control	4-3
SNA Protocols for Stream Control	4-3
Abnormal Stream Termination	4-4
Chapter 5. NJE Connection Control	5-1
Signon Concurrence Feature	5-2
External Interface	5-2
Implementation Example	5-2
Protocol Description	5-3
Establishing a Connection Without NPM	5-3
Establishing a Connection or Session With NPM	5-4
Signon Connection Protocol Between Two Nodes With NPM	5-4
Signon Protocols Between A Node With NPM and A Node Without NPM	5-7
Connection Status Information	5-7
Disconnections	5-7
Terminating a Session or Connection	5-8
Chapter 6. Link Protocols	6-1
Bisynchronous Communications Lines	6-2
Initialization	6-2
Initialization Error Recovery	6-4
Error Recovery	6-5
Termination	6-6
Normal Sequences	6-6
V.27 Modem Contention Resolution	6-6
Channel-to-Channel Adapter	6-8
Initialization	6-8
Error Recovery	6-9
Termination	6-10
Normal Sequences	6-10
PREPARE Mode Option: Suspend I/O (BSC and CTC Adapter)	6-11
Normal Sequences	6-11
Requesting I/O Suspension	6-12
Suspension and Resumption of I/O	6-12
BSC Error Protocols	6-16
CTC Adapter Error Protocols	6-21
Wait-a-Bit and PREPARE	6-24
Systems Network Architecture - LU Type 0	6-24
Overview	6-24
Session Initialization	6-25
Session Termination	6-33
Error Recovery	6-35
ACF/VTAM Considerations	6-36
Appendix A. System-Dependent Considerations	A-1
Support of Various NJE Features	A-1
JES2	A-1
Network Control	A-1
Commands and Messages (NMRs)	A-2
SYSIN (Job Input) Transmission	A-3
SYSOUT (Job Output) Transmission	A-5
Stream Support and Control	A-7
SNA Support	A-7
BSC Line and CTC Adapter Support	A-8

Accounting	A-9
Miscellaneous Considerations	A-9
JES3	A-10
Network Control	A-10
Commands and Messages (NMRs)	A-11
SYSIN (Job Input) Transmission	A-12
SYSOUT (Job Output) Transmission	A-14
Stream Support and Control	A-14
SNA Support	A-16
BSC Line and CTC Adapter Support	A-16
Accounting	A-17
Miscellaneous Considerations	A-18
RSCS	A-19
Network Control	A-19
Commands and Messages (NMRs)	A-20
SYSIN (Job Input) Transmission	A-21
SYSOUT (Job Output) Transmission	A-21
Stream Support and Control	A-26
SNA Support	A-27
BSC Line and CTC Adapter Support	A-28
Accounting	A-29
Miscellaneous Considerations	A-29
POWER	A-29
Network Control	A-29
Commands and Messages (NMRs)	A-30
SYSIN (Job Input) Transmission	A-31
SYSOUT (Job Output) Transmission	A-33
Stream Support and Control	A-35
SNA Support	A-36
BSC Line and CTC Adapter Support	A-36
Accounting	A-37
Miscellaneous Considerations	A-37

Appendix B. Abbreviations B-1

Glossary X-1

Index X-7

Figures

1-1.	An NJE Network	1-2
1-2.	Sample NJE Network	1-5
1-3.	NJE Addressing	1-6
1-4.	NJE Protocol Layers	1-7
1-5.	Format of an NJE SYSIN Job	1-8
1-6.	Format of an NJE SYSOUT Job	1-8
1-7.	Format of an NMR	1-10
1-8.	Format of a Transmission Buffer for BSC and CTC Protocols	1-12
1-9.	Format of a Transmission Buffer for SNA Protocol	1-13
1-10.	Stream Control Comparison Between SNA and BSC Protocols	1-14
1-11.	Establishing a BSC or CTC Session	1-16
1-12.	Terminating a BSC or CTC Session	1-17
1-13.	Establishing an SNA Session	1-17
1-14.	Terminating Processing by an SNA Secondary Node	1-18
1-15.	Terminating Processing by an SNA Primary Node	1-18
2-1.	Example of the Bit Notation Used in the Type of Data Column	2-3
2-2.	Job Header Sections used by Each Subsystem	2-3
2-3.	Fields in the Prefix Section of a Job Header	2-4
2-4.	Fields in the General Section of a Job Header	2-4
2-5.	Fields in the JES2 Section of the Job Header	2-16
2-6.	Fields in the POWER Section of the Job Header	2-17
2-7.	Fields in the Job Scheduling Section of the Job Header	2-18
2-8.	Fields in the Security Section of the Job Header	2-19
2-9.	Fields in the Job Header Installation Extension Section	2-20
2-10.	Data Set Header Sections used by Each Subsystem	2-21
2-11.	Fields in the Prefix Section of a Data Set Header	2-22
2-12.	Fields in the General Section of a Data Set Header	2-22
2-13.	Fields in the 3800 Section of a Data Set Header	2-32
2-14.	Fields in the RSCS Section of a Data Set Header	2-36
2-15.	Fields in the POWER Section of a Data Set Header	2-37
2-16.	Fields in the Record Characteristics Change Section of a Data Set Header	2-38
2-17.	Fields in the Output Processing Section of a Data Set Header	2-39
2-18.	OPTU Key Definition Table	2-42
2-19.	Print Server Dumping	2-47
2-20.	File Identifiers used during Output Processing	2-48
2-21.	Functional Subsystem Messages	2-49
2-22.	Trace ids	2-49
2-23.	Fields in the Security Section of the Data Set Header	2-52
2-24.	Fields in the Data Set Header Extension Section	2-53
2-25.	Job Trailer Sections used by Each Subsystem	2-54
2-26.	Fields in the Job Trailer Prefix Section	2-54
2-27.	Fields in the Job Trailer General Section	2-54
2-28.	Fields in the Job Trailer Accounting Section	2-57
2-29.	Fields in the Job Trailer User Section	2-57
2-30.	Identifying the Contents of an NMR	2-58
2-31.	Interpreting the NMROUT Field	2-59
2-32.	Fields in the Unformatted Command NMR	2-59
2-33.	Fields in the Formatted Command NMR	2-61
2-34.	Interpreting Field NMRMSG	2-63
2-35.	Fields in the Messages NMR	2-63
2-36.	Fields in an NJE Initial Signon Record	2-67

2-37.	Fields in an NJE Response Signon Record	2-68
2-38.	Fields in an NJE Reset Signon Record	2-69
2-39.	Fields in a Concurrence Signon Record	2-69
2-40.	Fields in an Add/Subtract Connection Record	2-70
2-41.	Fields in a Signoff Record	2-70
3-1.	Format of a Transmission Buffer for BSC and CTC	3-3
3-2.	BCB Bit Definitions	3-4
3-3.	Results of Not Transmitting Null Records.	3-5
3-4.	Correct Recovery with Null Records	3-5
3-5.	Transmission Block Handling State Table	3-6
3-6.	FCS Stream Bit Definitions	3-7
3-7.	SCB Definition	3-8
3-8.	RCB Definition	3-10
3-9.	SRCB Definition - Part 1	3-11
3-10.	SRCB Definition - Part 2	3-12
3-11.	Format of a Transmission Buffer for SNA	3-14
3-12.	String Control Byte (SCB) Format for SNA	3-14
3-13.	RIDRCB: Byte 1 of the RID	3-15
3-14.	RIDSRCB: Byte 2 of the RID	3-16
3-15.	RIDRLLEN: Byte 3 of the RID	3-17
3-16.	Data Record Without Carriage Control	3-17
3-17.	Data Record With Carriage Control	3-17
3-18.	Spanned Data Record	3-18
3-19.	Spanned Data Record With Carriage Control	3-18
4-1.	Protocols Used to Transmit a Stream	4-2
4-2.	Sending Node Initiates Stream Termination	4-4
4-3.	Receiving Node Initiates Stream Termination	4-4
5-1.	Connection Protocols used by the Non-Path Manager Nodes	5-3
5-2.	Normal Connection (Low Initiates)	5-5
5-3.	Normal Connection (High Initiates)	5-5
5-4.	Secondary Trunk Connection	5-6
5-5.	Trunk Reset Protocol	5-6
5-6.	NPM Connection Protocols	5-7
6-1.	BSC Initialization	6-2
6-2.	BSC Signon Contention	6-3
6-3.	BSC Channel Commands	6-4
6-4.	BSC Initialization Error Recovery	6-5
6-5.	BSC Error Recovery	6-5
6-6.	BSC Normal Communication	6-6
6-7.	V.27 Modem Protocols: Primary Node Signon	6-7
6-8.	V.27 Modem Protocols: Secondary Node Signon	6-7
6-9.	CTC Adapter Initialization	6-9
6-10.	CTC Adapter Program	6-10
6-11.	CTC Adapter Normal Sequences	6-10
6-12.	BSC Normal Communication	6-11
6-13.	BSC Communication in PREPARE Mode	6-13
6-14.	CTC Adapter Communication in PREPARE Mode	6-16
6-15.	BSC Error Protocols	6-17
6-16.	SNA Session Initialization	6-25
6-17.	SNA Session Initialization Error Recovery	6-26
6-18.	SNA FMH and Signon Flows	6-28
6-19.	FMH4 Format	6-29
6-20.	FMH3 Format	6-30
6-21.	Sample Compaction Table	6-32
6-22.	Sample Compacted Data	6-33
6-23.	SNA Session Termination - Initiated by Primary Node	6-34

6-24.	SNA Session Termination - Initiated by Secondary Node	6-34
6-25.	SNA Immediate Session Termination - Initiated by Primary Node	6-35
6-26.	SNA Immediate Session Termination - Initiated by Secondary Node	6-35
A-1.	Various NJE Features Supported by Networking Facilities	A-1
A-2.	RSCS Virtual Device Types	A-23

Special Notices

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates.

Any reference to an IBM licensed program or other IBM product in this publication is not intended to state or imply that only IBM's program or other product may be used. Any functionally equivalent program which does not infringe any of IBM's intellectual property rights may be used instead of the IBM product. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Commercial Relations, IBM Corporation, Purchase, NY 10577.

This *book* is intended to help the customer understand the Network Job Entry (NJE) formats and protocols used by System/370 Program Products that support networking and describes the similarities and differences between them. It also helps the customer diagnose NJE problems and customize products for specific environments or applications. This *book* documents no programming interfaces for use by customers in writing programs that request or receive the services of *MVS/System Product Version 3*.

Preface

Network Job Entry Formats and Protocols explains the Network Job Entry (NJE) formats and protocols used by System/370 Program Products that support networking and describes the similarities and differences between them. It consolidates information about how a network job entry subsystem represents an NJE transfer unit, control format, and communications information. This document allows users and systems programmers to make full use of job processing facilities available with the program products that support NJE. This document is also useful for diagnosing NJE problems and customizing products for specific environments or applications.

Product-specific comments refer to the latest shipped release unless otherwise stated. Refer to "Program Products" for the System/370 Program Products that support NJE.

This document does not replace, but supports, the existing product manuals listed in "Product-Related Publications" on page xii. For most product-related material, the product manuals should be consulted; however, Appendix A, "System-Dependent Considerations" contains detailed descriptions of product-specific implementations of, and deviations from, the networking formats and protocols described herein. It is recommended that you read it carefully.

Organization of this Manual

Network Job Entry Formats and Protocols is divided into the following chapters:

Chapter 1, "Introduction to Network Job Entry (NJE)" introduces the concepts needed to understand the material in the remaining chapters. This chapter defines a network and describes the 5 layers of networking.

Chapter 2, "Formats of NJE Transfer Units" describes the format of the control blocks used to transmit data across the network.

Chapter 3, "Presentation Services" discusses how the networking protocols span, compress, and compact data to be transmitted.

Chapter 4, "Stream Control" discusses the multileaving functions of SNA, BSC, and CTC.

Chapter 5, "NJE Connection Control" discusses establishing or terminating a connection between two nodes.

Chapter 6, "Link Protocols" discusses SNA, BSC, and CTC protocol sequences for initialization, normal processing, and termination.

Appendix A, "System-Dependent Considerations" describes system-dependent considerations and deviations from the protocols described in this book.

Appendix B, "Abbreviations" defines the abbreviations and terms used in this manual.

Program Products

The following lists the program products that use the protocols as they are documented in this manual.

- MVS/SP JES2
 - Version 1.3.6
 - Version 2.1.5 and its subsequent releases
 - Version 2.2.3 and its subsequent releases
 - Version 3.1.1 and its subsequent releases
 - Version 3.1.3 and its subsequent releases
- MVS/SP JES3
 - Version 1.3.4
 - Version 2.1.5 and its subsequent releases
 - Version 2.2.3 and its subsequent releases
 - Version 3.1.1 and its subsequent releases
 - Version 3.1.3 and its subsequent releases
- RSCS
 - Version 1.3 and its subsequent releases
 - Version 2.2 and its subsequent releases
 - Version 2.3 and its subsequent releases
- VSE/POWER
 - Version 2.1 and its subsequent releases
 - Version 2.2 and its subsequent releases
- MVS/BDT Version 2

Product-Related Publications

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.

Networking - General

- *Job Networking (IBM Systems Journal V17 N3, 1978)*, G321-5071
- *Job Networking Facilities*, GG22-9042
- *NJE Concepts and Protocols Overview*, GG66-0224

Systems Network Architecture (SNA)

- *Concepts and Products*, GC30-3072
- *Technical Overview*, GC30-3073
- *Reference Summary*, GA27-3136
- *Format and Protocol Reference Manual: Architectural Logic*, SC30-3112
- *IBM Synchronous Data Link Control General Information*, GA27-3093

ACF/VTAM

- *General Information*, GC27-0608
- *Planning and Installation Reference*, SC27-0610
- *Programming*, SC27-0611

Binary Synchronous Communications (BSC)

- *General Information*, GA27-3004

Channel-to-Channel Adapter (CTC)

- *S/370 Special Feature Description: CTCA*, GA22-6983

JES2

MVS/SP JES2 - Version 1 (5740-XYS)

- *JES2 Initialization and Tuning*, SC23-0046
- *JES2 Commands*, SC23-0048
- *MVS JCL User's Guide*, GC28-1349
- *MVS JCL Reference*, GC28-1350
- *JES2 Logic*, LY24-6006

MVS/SP JES2 - Version 2 (5740-XC6)

- *JES2 Initialization and Tuning*, SC23-0065
- *JES2 Commands*, SC23-0064
- *MVS JCL User's Guide*, GC28-1351
- *MVS JCL Reference*, GC28-1352
- *JES2 Logic*, LY24-6008

MVS/ESA JES2 - Version 3 (5685-001) Release 1.0

- *JES2 Initialization and Tuning*, SC23-0065
- *JES2 Commands*, SC23-0064
- *MVS JCL User's Guide*, GC28-1830
- *MVS JCL Reference*, GC28-1829
- *JES2 Logic*, LY28-1006

MVS/ESA JES2 - Version 3 (5685-001) Release 1.1 and Subsequent Releases

- *JES2 Initialization and Tuning*, SC28-1038
- *JES2 Commands*, SC23-1039
- *MVS JCL User's Guide*, GC28-1830
- *MVS JCL Reference*, GC28-1829
- *JES2 Logic*, LY28-1006

JES3

MVS/SP JES3 - Version 1 (5740-XYN)

- *JES3 Initialization and Tuning*, SC23-0041
- *JES3 Operator's Library*, SC23-0045
- *MVS JCL User's Guide*, GC28-1349
- *MVS JCL Reference*, GC28-1350
- *JES3 Modifications and Macros*, LC28-1371
- *JES3 Diagnosis*, LC28-1369

MVS/SP JES3 - Version 2 (5665-291)

- *JES3 Initialization and Tuning*, SC23-0059
- *JES3 Commands*, SC23-0063
- *MVS JCL User's Guide*, GC28-1351
- *MVS JCL Reference*, GC28-1352
- *JES3 Modifications and Macros*, LC28-1372
- *JES3 Diagnosis*, LC28-1370

MVS/ESA JES3 - Version 3 (5685-002) Release 1.0

- *MVS/XA: JES3 Initialization and Tuning*, SC23-0059
- *MVS/XA: JES3 Commands*, SC23-0063
- *MVS JCL User's Guide*, GC28-1830
- *MVS JCL Reference*, GC28-1829
- *JES3 Customization*, LY28-0076
- *JES3 Diagnosis*, LY28-1012

MVS/ESA JES3 - Version 3 (5685-002) Release 1.1 and Subsequent Releases

- *JES3 Initialization and Tuning*, SC23-0073
- *JES3 Commands*, SC23-0074
- *MVS JCL User's Guide*, GC28-1830
- *MVS JCL Reference*, GC28-1829
- *JES3 Customization*, LY28-0076
- *JES3 Diagnosis*, LY28-1012

VM

- *Evolution of Virtual Machine Subsystem (IBM Systems Journal, V18, N1, 1979)*, G321-5089

RSCS Networking Version 1 (5749-DMT)

- *Program Logic Manual*, LY24-5203
- *Operator's Reference Summary Card*, SX24-5119

RSCS Networking Version 2 (5664-188)

- *Operation and Use*, SH24-5058
- *Diagnosis Reference*, LY24-5228

VSE/SP

VSE/POWER Version 2 (5666-273)

- *Networking User's Guide*, SC33-6140
- *Installation and Operations Guide*, SH12-5329
- *Networking Design Guide*, GG24-1570
- *Networking*, SC33-6180

Chapter 1. Introduction to Network Job Entry (NJE)

Units of Work	1-3
Functions of a Node	1-3
Types of Nodes	1-4
Types of NJE Users	1-6
NJE Addressing	1-6
NJE Protocols	1-7
Formats of NJE Transfer Units	1-8
Format of an NJE Job	1-8
Format of NMRs	1-9
Presentation Services	1-10
Preparing Transmission Buffers	1-11
Spanning Data	1-11
Compressing Data	1-11
Format of a Transmission Buffer for BSC and CTC Protocols	1-12
Format of a Transmission Buffer for the SNA Protocol	1-13
Stream Control	1-14
Multileaving for the SNA and BSC Protocols	1-14
NJE Connection Control	1-15
NJE Connection Control Without NPM	1-15
NJE Connection Control With NPM	1-15
Link Protocols	1-16
Initiating a BSC or CTC Session	1-16
Terminating a BSC or CTC Session	1-17
Initiating an SNA Session	1-17
Terminating an SNA Session	1-18

An NJE network is a group of two or more complexes or systems that communicate with each other. An NJE network is comprised of a home node and all the other nodes that it can transmit a unit of work to or receive a unit of work from. The nodes in an NJE network use protocols to communicate with each other. Protocols are rules a node uses to:

- Become part of an NJE network
- Receive a unit of work
- Send a unit of work
- Indicate it was removed from the network.

For an illustration of an NJE network, refer to Figure 1-1.

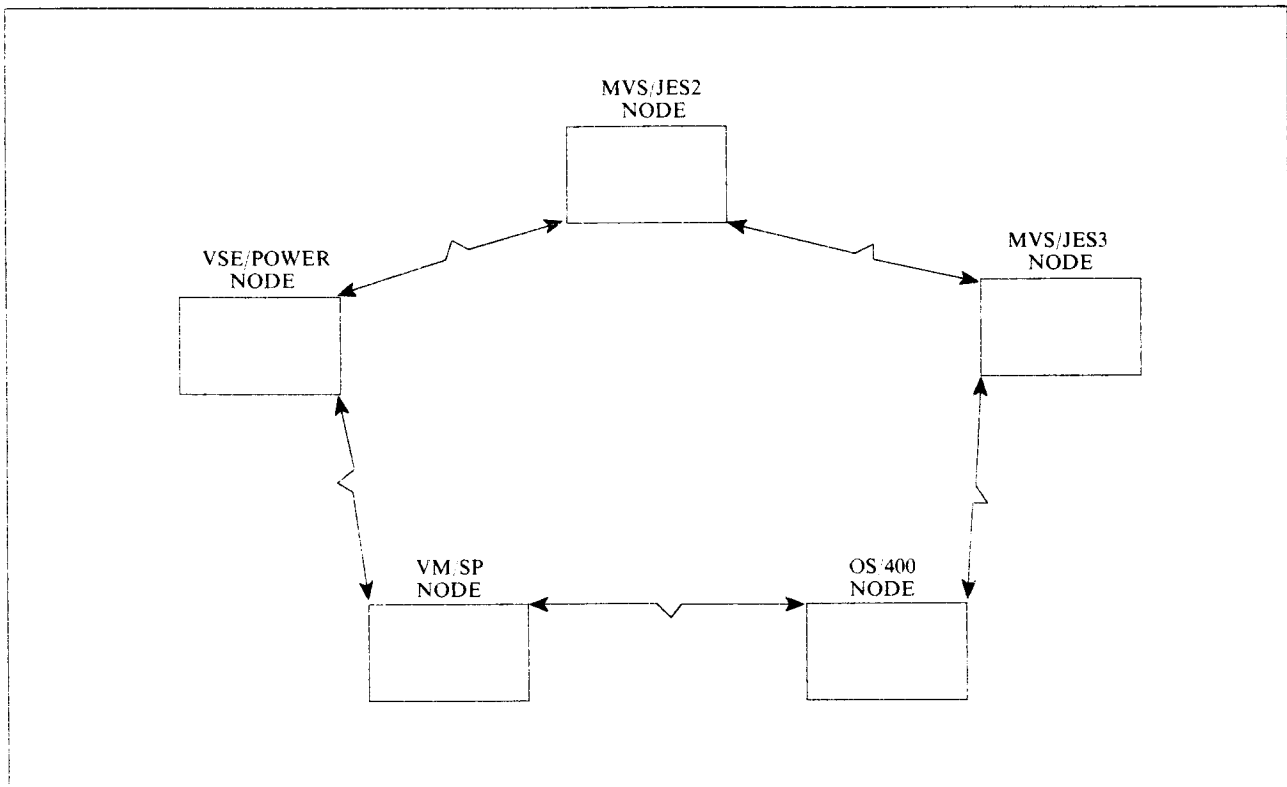


Figure 1-1. An NJE Network

To become part of an NJE network, an installation must install an NJE network facility that uses the protocols documented in this manual. Any of the following network facilities can be installed:

- JES2
- JES3
- RSCS
- POWER
- OS/400

Units of Work

A unit of work that is transmitted across the network is called an NJE transfer unit. An NJE transfer unit can be either an NJE job or a nodal message record (NMR).

An **NJE job** is a transfer unit that contains data to be processed at another node in the NJE network. It begins with a job header, is followed by data, and ends with a job trailer. The type of data contained in the NJE job further defines the type of NJE job. The data between the job header and job trailer can be either SYSIN or SYSOUT data. An NJE SYSIN job is an NJE job that contains JCL for a job and may have one or more SYSIN data sets. An NJE SYSOUT job is an NJE job that contains one or more SYSOUT data sets. Each SYSOUT data set is preceded by a data set header.

A **nodal message record (NMR)** is a unit of work that begins with an NMR header and is followed by message text. The message text can be either a message or a system command.

Functions of a Node

A node is a system or complex that is defined to an installation. A node in the network can be another complex or system within a single location or it can be a complex that resides in a remote location. Each node that a complex can access must be identified to other complexes by a unique NJE node name.

Note: If a node uses SNA protocols, the node has two names -- an LU name (as defined to VTAM) and the NJE node name. The NJE node name is created during initialization processing and appears in job headers, data set headers, and NMRs. You should not confuse the LU name with the node name, they are two separate entities.

Each node in the network can do the following with an NJE transfer unit:

Transmit

The node packages the NJE transfer unit and transmits it to another node.

Receive

The node recognizes the NJE transfer unit, receives, and stores it.

Store-and-forward

The node accepts the NJE transfer unit, stores it, and schedules it to be forwarded to another node.

Types of Nodes

NJE uses the following terminology for the nodes that comprise an NJE network.

Originating Node

The node where the user submitted the request to *transmit* the data to another complex.

Intermediate Node

A node that lies in the path of either the:

- Originating node and execution node
- Execution node and the destination node

It *receives* and *transmits* the NJE transfer unit to the next node in the path of the target node.

Target Node

Is the node where a NJE job or NMR is *received* and will either be *executed* or be *processed*. The target node can be either a:

- **Destination Node** is a node that receives and processes:
 - An NJE SYSOUT job. A node processes an NJE SYSOUT job by printing or punching the SYSOUT data set.
 - A message contained in an NMR.

Note:

When an NJE transfer unit 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.

- **Execution Node** is the node where:
 - JCL contained in an NJE SYSIN job executes. The node packages the SYSOUT data sets created by the SYSIN in an NJE SYSOUT job and sends the NJE job to the destination node.
 - A command contained in an NMR is processed. The node packages the messages that are a result of the command in an NMR and sends the NMR to the destination node.

The execution node may not necessarily be the destination node. If, for example, a user submits a job specifying that the job execute at one complex and job's output be printed at a different complex then the complex where the job runs is the execution node and the complex where the SYSOUT prints is the destination node. If no SYSOUT destination is specified, then the origin node and destination node are the same by default.

Figure 1-2 illustrates the different types of nodes in an NJE network. The network is composed of 3 nodes. If a user submits a job at node 1 to be executed at node 3:

- Node 1 is the originating node because that is where the user submitted the request.
- Node 2 is the intermediate node, because it is in the path of the destination node. Node 2 receives the data, stores it, then forwards it to the next node in the path of the target node.
- Node 3 is both the execution and destination node because it is the node that the user specified as the target.

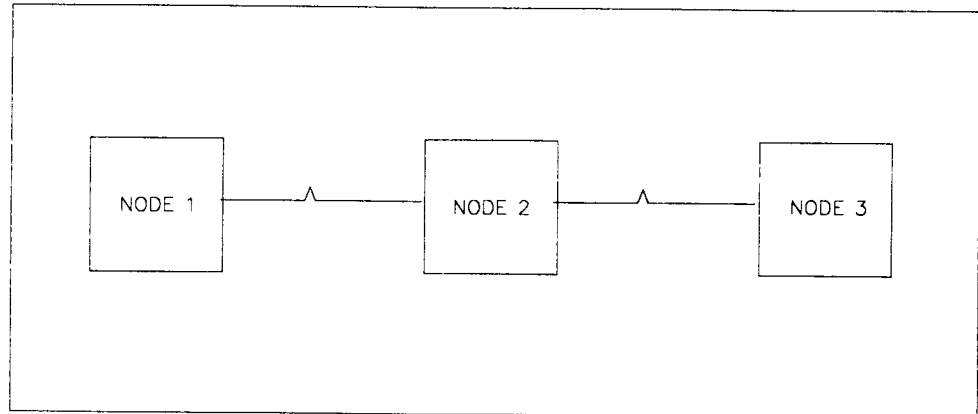


Figure 1-2. Sample NJE Network

To transmit an NJE transfer unit to a complex other than the user's installation (a remote node), the user issues a command or submits a job specifying a destination node name. The destination node can be either directly- or indirectly-connected to the originating node. In the network depicted in Figure 1-2, if NODE1 is the originating node, then:

- NODE2 is a directly-connected node to NODE1
- NODE3 is an indirectly-connected node to NODE1.

Types of NJE Users

Originating User

Is the user that submits an NJE transfer unit at the originating node. The originating user submits the NJE transfer unit at an operator console, terminal, or an RJE workstation. An NJE transfer unit may originate from another NJE transfer unit.

Destination User

Is a user or device (printer or punch) that is the target of an NJE SYSOUT job.

Notification User

Is the user who receives messages that notify the user of the status of the NJE transfer unit.

Accounting User

Is the user that receives the notification of the amount or cost of system resources used in processing an NJE transfer unit.

NJE Addressing

NJE protocols are concerned with routing an NJE job to the specified destination. NJE protocols obtain the destination from the NJE job's JCL. If the originating user identifies an RJE workstation or a userid that should receive the output or notification messages, the additional information is appended to the end of the destination node. Figure 1-3 identifies the different methods of addressing nodes.

Destination = Node
Destination = (Node. Remote id)
Destination = (Node. User id)

Figure 1-3. NJE Addressing

It is up to the destination node to route the job to the remote or user at the destination node.

NJE Protocols

Nodes in an NJE network are connected by lines such as SNA sessions, channel-to-channel adapters, leased or dial-up telephone lines, microwave links, or by satellite. An **NJE connection** is the use of NJE protocols between two nodes. NJE protocols are rules that direct the logical structure, format and operational sequence for transmitting NJE transfer units between two nodes in an NJE network. The three NJE protocols are:

- Binary synchronous communications (BSC)
- LU type 0 Systems network architecture (SNA)
- Channel-to-channel (CTC).

Channel-to-channel protocols are similar to BSC protocols and are treated the same in this manual.

Figure 1-4 illustrates the NJE protocol layers:

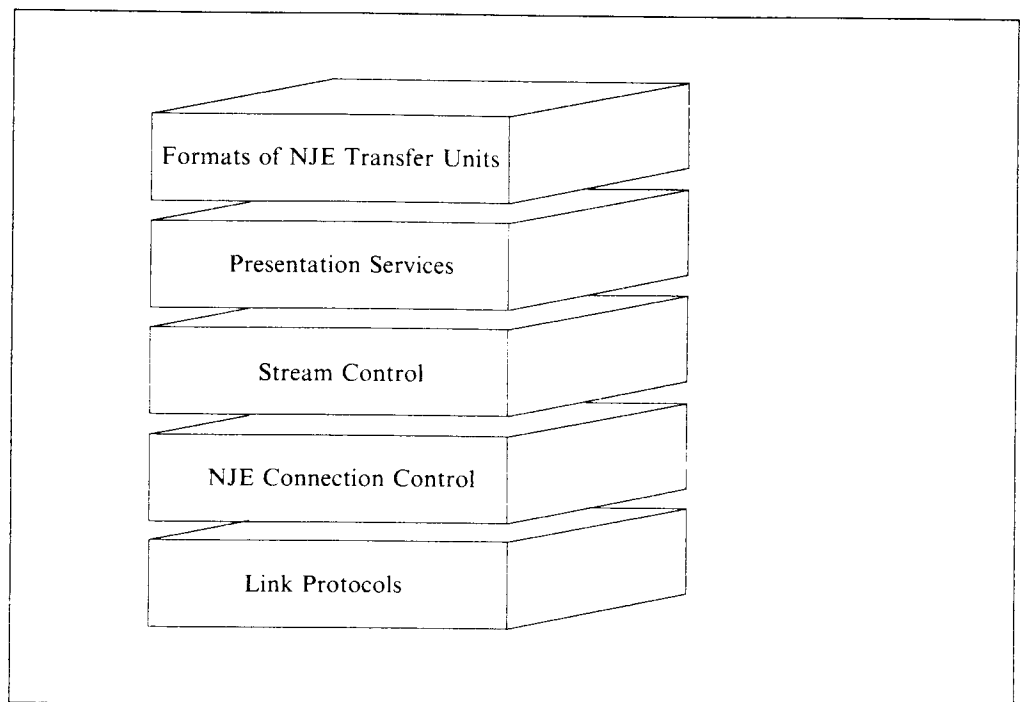


Figure 1-4. NJE Protocol Layers

The remaining information in this manual describes these layers:

- **Chapter 2, “Formats of NJE Transfer Units”** is the first layer of the NJE protocol and describes the format of data and control blocks used for:
 - An NJE job
 - Transmitting a command or message (NMR).
- **Chapter 3, “Presentation Services”** is the second layer of the NJE protocol and discusses the preparation of transmission buffers.
- **Chapter 4, “Stream Control”** is the third layer of the NJE protocol and describes the multileaving function of NJE.
- **Chapter 5, “NJE Connection Control”** is the fourth layer of the NJE protocol and describes the protocols needed to establish and terminate communication between nodes.

- **Chapter 6, “Link Protocols”** is the fifth layer of the NJE protocol and describes the protocols used to transmit an NJE job to a directly-connected node using BSC, CTC, or SNA protocols. It also discusses error recovery processing for each of the protocols.

Formats of NJE Transfer Units

This section discusses the data format for NJE jobs and nodal message records (NMRs). For detailed information about the control blocks contained in an NJE transfer unit, refer to Chapter 2, “Formats of NJE Transfer Units” on page 2-1.

Format of an NJE Job

An NJE job contains either SYSIN or SYSOUT data. The data must be packaged in a format that is recognizable by the sending and receiving nodes. An NJE job must contain:

1. A job header
2. Data
3. A job trailer.

An NJE job containing SYSIN data begins with a job header record followed by the job and SYSIN data, and ends with the job trailer record. Figure 1-5 shows the format of an NJE SYSIN job.




Figure 1-5. Format of an NJE SYSIN Job

An NJE job containing SYSOUT data begins with a job header record followed by one or more data set header records and SYSOUT data sets, and ends with a job trailer record. Figure 1-6 shows the format of an NJE SYSOUT job.




Figure 1-6. Format of an NJE SYSOUT Job

The following describes the header, data, and trailer fields that comprise an NJE SYSIN or SYSOUT job.

Job Header

The job header record is a variable length record that contains one or more variable length sections. The job header record contains the:

- Job identifier
- Network jobid
- Execution address
- Notification address
- Origination address
- Security information
- Accounting information
- Execution defaults
- Estimated resources required.

Data

The networking protocols assume the work specification record length (SYSIN record length) is 80 bytes fixed. Trailing blanks may be truncated before transmission. All data is transmitted in transparent mode (unprintable characters are not translated).

- **Job/SYSIN data** is a single job or work specification. The syntax for work specification is that of the node where the job will execute. NJE protocols do not standardize work specification syntax. The work specification can contain more than one SYSIN data set, each with a different record length. If the record length varies, a record change characteristics section (RCCS) containing the record length of the data (LRECL) and the record format (U, F, or V) will precede each SYSIN data set.

Note: For JES3 nodes, the work specification will be MVS JCL or JES3 JECL. JES3 does not recognize the RCCS, but uses the record length information at the beginning of each record to determine the record's length.

- The **Data Set Header** and **SYSOUT Data** are created either when a:
 - Job is sent to a node to execute and the job creates SYSOUT data sets that must be sent to another node
 - SYSOUT data set is sent to another node

An NJE SYSOUT job can contain one or more unique SYSOUT data sets preceded by one or more data set headers. The data set header contains the following information:

- Data set name (unique only within the creating job)
- Destination address
- Source data record format and logical record length
- Security information
- Destination processing attributes such as SYSOUT class, FCB, and forms identifier.

Job Trailer

The job trailer record is a variable length record that contains one or more variable length sections. The job trailer is designed to contain execution time statistics and, although present at the end of the SYSIN data, contains no usable information until after job execution. For SYSOUT data, it may contain execution statistics for the job that created the SYSOUT data set.

Format of NMRs

Commands and messages are transmitted as individual nodal message records (NMRs). NMRs that contain a command or message cannot exceed 256 bytes. The NMR header fields specify the:

- Originating address
- Destination address
- Identifier of the console where the command was issued.

All NMRs contain a header followed by a message or command (formatted or unformatted) as shown in Figure 1-7.

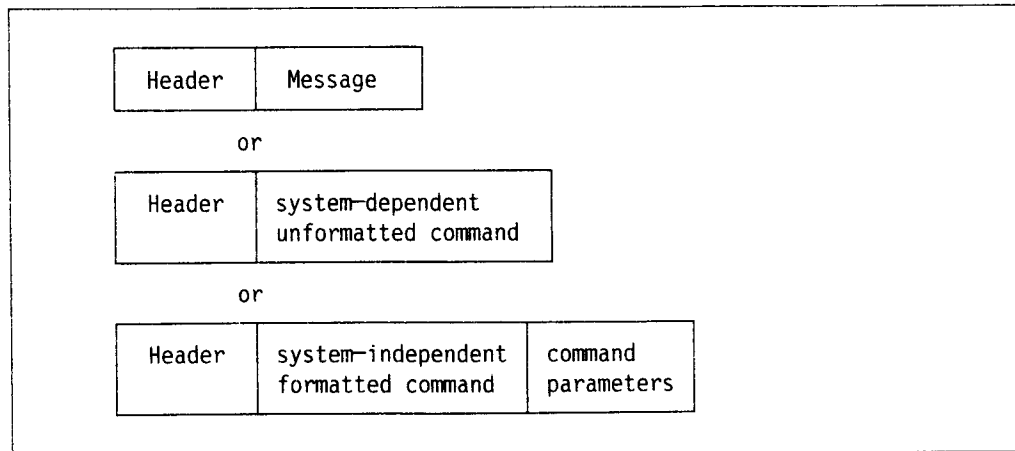


Figure 1-7. Format of an NMR

Command NMR

When a command reaches the execution node, the execution node retains the header portion of the NMR. If there is a response to the command, the header information is used by the execution node to create a header for a response NMR.

If an NMR has text that is **unformatted**, then the NMR contains commands that have system-dependent syntax. If an NMR has text that is **formatted**, then the NMR contains one of the following commands, which have system-independent syntax:

- DISPLAY
- CANCEL
- HOLD
- RELEASE
- REROUTE

Message NMR

For messages, the NMR header may also contain an:

- Importance level
- Output priority for unsolicited status messages that are not command responses.

The message text can begin with a time stamp and/or originating userid. An indicator in the NMR header defines whether any of these fields are present at the start of the text. The originating userid in the message text is used at the destination node to issue the messages to the user that submitted the job. If the message is in response to a command, then the message is transmitted to the node where the command was issued.

Presentation Services

Presentation services is the second layer of NJE protocols. It describes how an NJE transfer unit is packaged into one or more transmission buffers so that it can be transmitted across the network.

Before transmitting an NJE transfer unit, the networking facility must:

1. Obtain the logical record from spool, if necessary
2. Segment the logical record into transmission buffers, if needed
3. Compress the transmission buffers within transmission blocks
4. Prepare NJE transmission buffers.

Preparing Transmission Buffers

Each transmission buffer contains:

- Routing information used by the protocol
- Part of the NJE transfer unit.

Usually, more than one NJE transmission buffer is needed to transmit an NJE job, while an NMR can be contained within a single buffer.

The format of the transmission buffer depends on the type of protocols the node uses. A transmission buffer always contains:

- Control information used to transmit the buffer
- Transmission blocks.

Spanning Data

NJE facilities partition or span data into lengths it can easily manage for purposes of writing data to spool or transmitting data across the network. However, for the following reasons, the data in an NMR or a SYSIN data set is not spanned:

- The length of an NMR never exceeds the length of a transmission block.
- NMRs are not usually written to spool.
- The length of a logical record for SYSIN data never exceeds the length of a transmission block.

In order to transmit data across the network, the networking protocols documented in this manual require data to be divided into record segments that are no longer than 256 bytes. When an NJE network facility is spanning data, trailing blanks may be truncated. Each record segment contains a control segment and no more than 256 bytes of data. A control segment contains:

- The function control sequence (FCS) stream identifier (two function control sequence bytes indicating whether the data is SYSIN or SYSOUT).
- The type of record (job header, data set header, data, or job trailer).
- The length of the data.

Compressing Data

While preparing data for transmission, the networking facility compresses it removing blanks and duplicate characters. Compressing data reduces the amount of time required to transmit a transmission buffer.

When the networking facility encounters a string of duplicate letters or blanks, it places a string control byte (SCB) in the record to indicate the number of blanks or repeating non-blank characters that were deleted from the transmission buffer. The receiving node uses the information in the SCB to restore the record. An SCB can indicate up to 31 repeating characters and up to 63 non-repeating characters.

Format of a Transmission Buffer for BSC and CTC Protocols

Figure 1-8 illustrates the format of a buffer to be transmitted using BSC and CTC protocols.

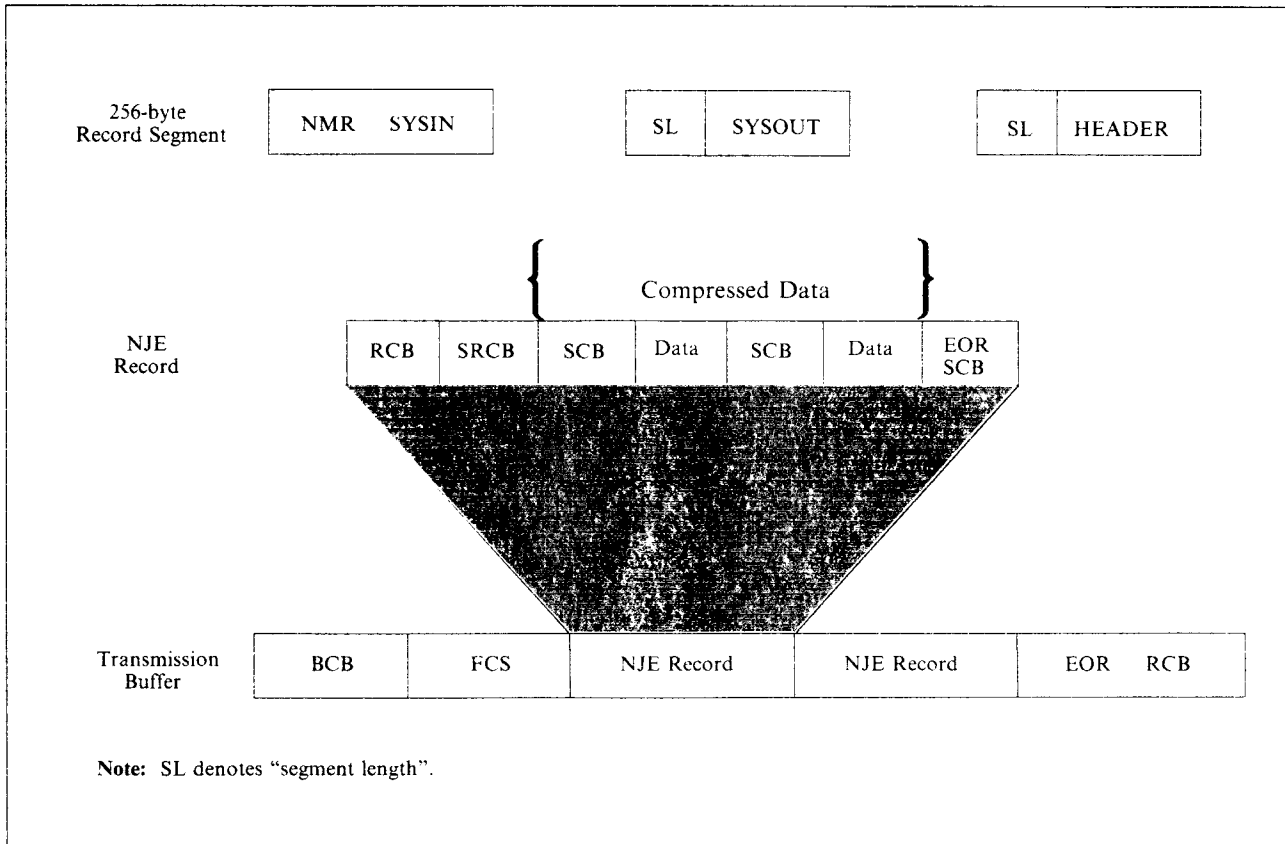


Figure 1-8. Format of a Transmission Buffer for BSC and CTC Protocols

Control and End-of-Buffer Information

Control information is placed at the beginning of each buffer that will be transmitted using BSC or CTC protocols. Every BSC transmission buffer begins with a block control byte (BCB) containing the outbound buffer sequence number (modulo 16). The sending node and the receiving node maintain an outbound and inbound buffer sequence counter.

Following the BCB are two bytes used by each receiver to control the inbound flow. These function control sequence (FCS) bytes permit nodes using BSC protocols to allow 8 SYSIN and SYSOUT streams to flow concurrently in each direction. Transmission blocks are placed after the FCS bytes. The format of a BSC transmission buffer does not prevent NJE from placing transmission blocks from different NJE transmission units in the same transmission buffer, but the common protocol precludes it.

When a compressed transmission block cannot fit into the buffer, a special stand-alone record control byte (RCB) of X'00' is placed after the last transmission block and the buffer is truncated.

NJE Records

Each record segment is compressed into an NJE record. BSC compression results in string control bytes (SCBs) interspersed in the NJE record with duplicate character strings compressed. The SCB defines how many blanks or repeating non-blank characters are to be inserted by the receiver (up to 31 characters) or the length of strings without any duplicate characters (up to 63 characters). The resulting format will always begin with an initial SCB and may have additional SCBs interspersed.

The BSC protocol places a record control byte (RCB) and a subrecord control byte (SRCB) at the beginning of an NJE record. The RCB contains the FCS stream identifier and number; the SRCB defines the type of NJE record contained in the transmission buffer.

A stand-alone SCB with a value of X'00' (sometimes referred to as an "EOR SCB") is always placed at the end of each NJE record. In the BSC protocol, an end-of-file is represented by an RCB with the FCS stream identifier, and an SRCB of X'00'.

Format of a Transmission Buffer for the SNA Protocol

The SNA protocol compresses, and optionally compacts, the information in the NJE record before placing it in a transmission buffer. Therefore, not only does NJE compress the information in the transmission buffer, but it also compresses the control information.

After the data in the NJE record is compressed, including the record identifier (RID), it is placed in the transmission buffer. A transmission buffer for SNA protocols is also called a request unit (RU). There is no control information placed before the first transmission NJE record. Figure 1-9 illustrates the format of a transmission buffer to be transmitted using the SNA protocol.

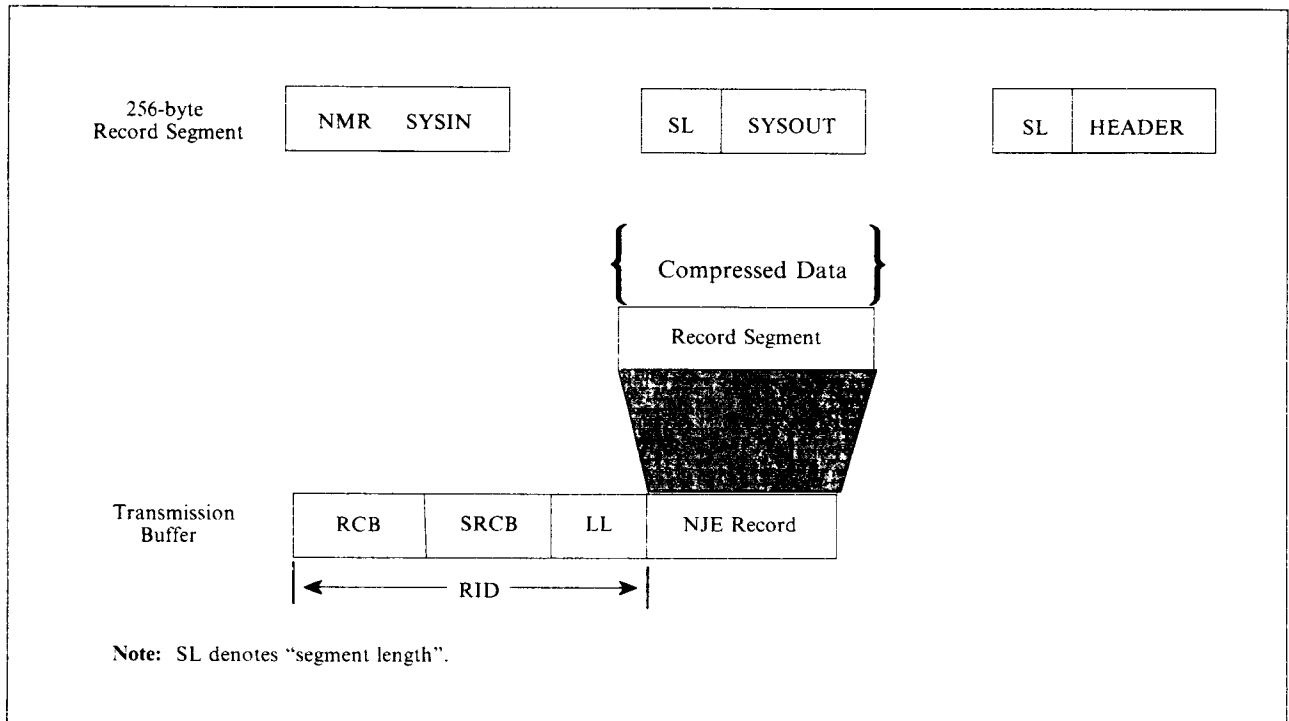


Figure 1-9. Format of a Transmission Buffer for SNA Protocol

Transmission Buffer

A transmission buffer contains a record control byte (RCB), a subrecord control byte (SRCB), and a logical record length (LL). These are placed at the front of the NJE transmission buffer. For the SNA protocol, The RCB, SRCB, and the LL comprise the record identifier. The RCB contains the stream identifier and number; the SRCB defines the type of NJE record contained in the transmission buffer; the LL contains the length of the data prior to compaction and compression, minus 1. The length byte is used by presentation services at the receiving node to locate the next RCB in the inbound RU.

Stream Control

Stream control is the third layer of NJE protocols. After NJE's presentation services have packaged the NJE record into transmission buffers and created a stream of buffers, the stream can be transmitted. The stream control layer describes a node's ability to concurrently transfer multiple streams within the same BSC connection or SNA session. This is known as "multileaving".

Multileaving for the SNA and BSC Protocols

NJE protocols permit the identification of 7 SYSIN streams, 7 SYSOUT streams, an NMR stream, and a control stream in both directions. The capability of a node to send a stream in one direction while receiving unrelated data in another direction is referred to as "full duplexing". Figure 1-10 compares the SNA and BSC protocols with respect to stream control.

<i>Figure 1-10. Stream Control Comparison Between SNA and BSC Protocols</i>		
	SNA (LU type 0)	BSC
Duplexing	<ul style="list-style-type: none">• True full duplex. Each end concurrently can be in send and receive states. The two nodes operate independently.• Is handled by VTAM at the session level.	<ul style="list-style-type: none">• Not true full duplex. Each node must flip-flop between the send and receive states. Both ends are synchronized.
Individual Stream Control	<ul style="list-style-type: none">• Does not allow for "individual" control of streams	<ul style="list-style-type: none">• Allows for "individual" control of streams -- receiving node tells the sending node to hold one or more streams, thus letting the remaining streams continue.
Transmission failures	<ul style="list-style-type: none">• VTAM retries most transmission errors.	<ul style="list-style-type: none">• The receiving node must notify the sending node of temporary link failures so that the transmission can be retried.

NJE Connection Control

NJE Connection Control is the third layer of NJE protocols.

The type of networking facility at an installation (JES2, JES3, RSCS, OS/400, or POWER) determines the way in which two nodes establish a connection. This is because JES2 uses the Network Path Manager (NPM) to establish a connection, while the others do not. If one of the connecting nodes is a JES2 complex, then NPM influences the protocols used to establish the connection. This section briefly discusses NJE connection control for complexes establishing a connection:

1. Without NPM (no JES2 nodes)
2. With NPM
 - Between a JES2 Node and Another Networking Facility
 - Between Two JES2 Nodes.

After the primary and secondary node have established communications, signon records are exchanged. The primary node sends an NJE initial signon record to the secondary node. The NJE initial signon record contains the:

- Name of the primary node. (A member number is also included for multi-CPU complexes.)
- Line and node passwords.
- Size of the BSC buffer.
- Signon concurrence flags, which allow the two nodes to determine the extended capabilities each node supports.

The secondary node responds to the initial signon record with a response signon record.

NJE Connection Control Without NPM

An exchange of control records establishes a connection between two nodes that do not use the network path manager. The initial signon record, sent by the "primary" node, is called an "I record". The receiving node, or "secondary" node, sends a response signon record, called a "J record". Extended NJE features exist and are acknowledged by the two nodes by setting bits in their respective I and J records.

Normal disconnection of an NJE node (using BSC or CTC protocols) is performed when the primary node sends a final signoff record, called a "B record". In the SNA protocol, disconnection takes place when the primary node sends a CLSDST or when the secondary node sends a TERMSESS record.

NJE Connection Control With NPM

A node that uses NPM uses the Network Path Manager (NPM) to establish and manage a connection. Among other things, NPM notifies other nodes about connections and disconnections to the network and maintains line status information. In controlling network traffic, NPM relies upon a (programmer-specified) line resistance value and the use of connection event sequence (CES) values taken from the system TOD clock.

Between a JES2 Node and Another Networking Facility

In connecting a JES2 node to another networking facility, pre-defined connections are required before I and J records can be sent. Once defined, either node can initiate an NJE network connection. The non-JES2 facility must discard NPM records other than I, J, and B sent by the JES2 node.

Between Two JES2 Nodes

The EBCDIC translation of the installation-supplied node name determines which node is "high" and which is "low". The high end is the primary node; the low end is the secondary node. Either end may initiate a connection or reset a connection if the connection is determined to be invalid. Connection status information is traded between NPMs through Add Connection and Subtract Connection records.

Link Protocols

The lowest layer of the NJE protocols is the link protocols layer, which describes establishing and terminating BSC, CTC, and SNA sessions. Installations using BSC lines or CTC adapters use a link protocol to establish and terminate a session.

SNA sessions use a link protocol to:

- Initiate a session
- Identify the capabilities of each node
- Identify the compaction tables
- Terminate a session.

The following introduces these two separate link protocols.

Initiating a BSC or CTC Session

To initiate a BSC or CTC session, the nodes exchange signon records as shown in Figure 1-11.

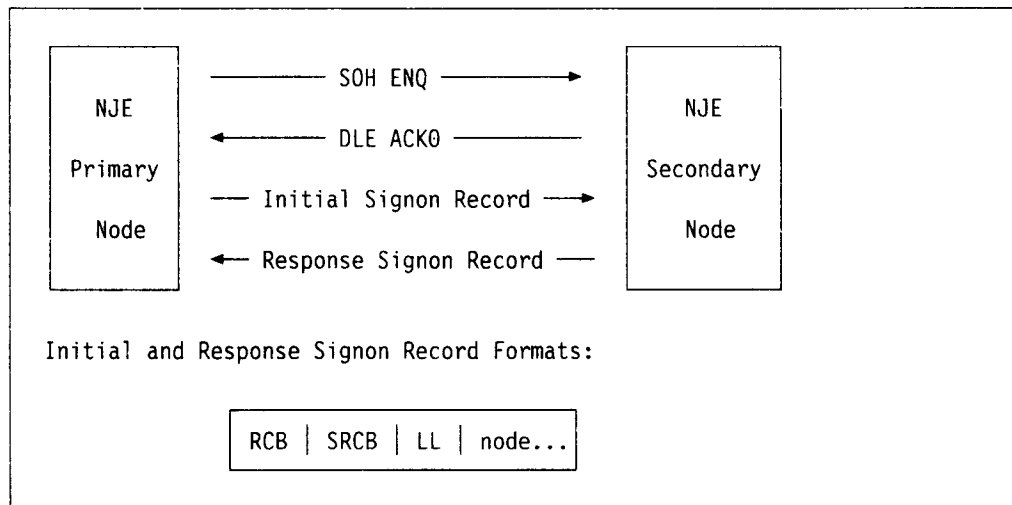


Figure 1-11. Establishing a BSC or CTC Session

A node using BSC lines or CTC adapters initiates a session by sending the SOH ENQ request. Because both nodes can simultaneously request a session, the node that receives the SOH ENQ request first becomes the **secondary node**. The secondary node acknowledges the request for the session by sending a DLE ACK0 response to node that initiated the session. The node that receives the DLE ACK0 response is called the **primary node**.

Terminating a BSC or CTC Session

The operator at either the primary or secondary node can issue an operator command to stop the transmission of an NJE job. The termination request is sent as part of a transmission block. After the request is received, the node can quiesce the transmission buffers, then terminate the connection by sending a sign-off record.

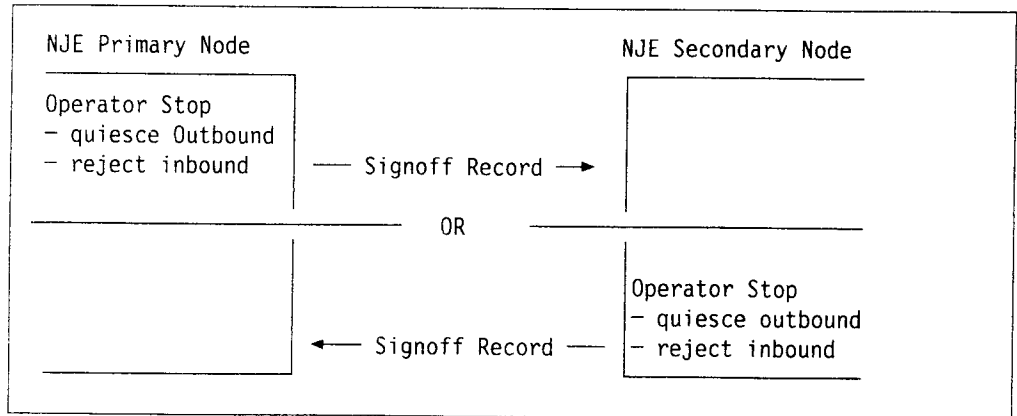


Figure 1-12. Terminating a BSC or CTC Session

Initiating an SNA Session

To initiate an SNA session, the nodes exchange signon records as shown in Figure 1-13.

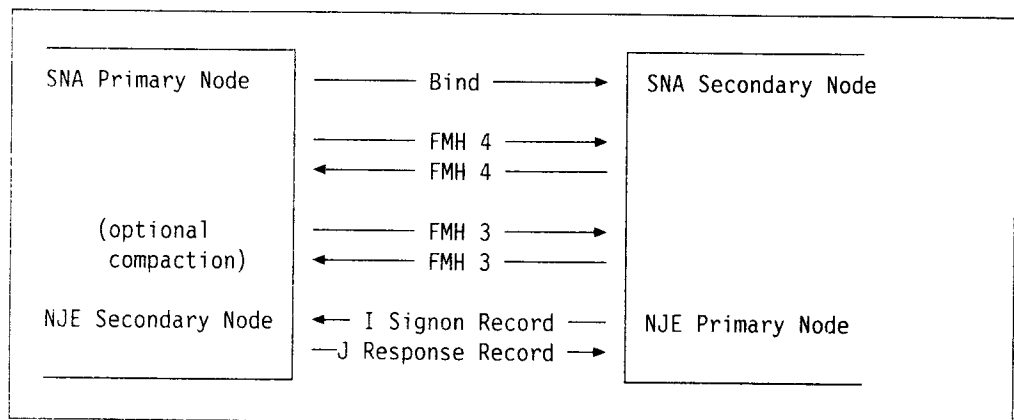


Figure 1-13. Establishing an SNA Session

To initiate an SNA session, a node sends a BIND request. The node that issues the BIND request is the SNA primary node and the node that receives the BIND request becomes the SNA secondary node.

After the session is initiated, the SNA primary and SNA secondary nodes must identify their capabilities. The nodes exchange private NJE Function Management Header type 4 (FMH4) records to describe:

- The size of the transmission buffer or request unit (RU) for the node.
- Whether SNA primary node will compact the data in the NJE transmission unit.
- Network topology records.

If the FMH4 records indicate the data in the transmission buffers will be compacted, compaction translation tables are exchanged in FMH3 records. FMH3 and FMH4 records are sent with "definite responses" requested. If a negative response is received by either of the nodes, the session terminates. All other records are sent as "exception response only" records.

After the exchange of FMH records, the nodes are ready to signon. The SNA node with the higher node name (the NJE primary node) sends the signon record. The NJE secondary node responds as in BSC/CTC protocol.

Terminating an SNA Session

The operator can terminate an SNA session at either the SNA primary or the SNA secondary node but only the SNA secondary node can request the termination. Figure 1-14 illustrates termination processing when an SNA secondary node terminates the session. Figure 1-15 illustrates termination processing when the SNA primary node terminates the session.

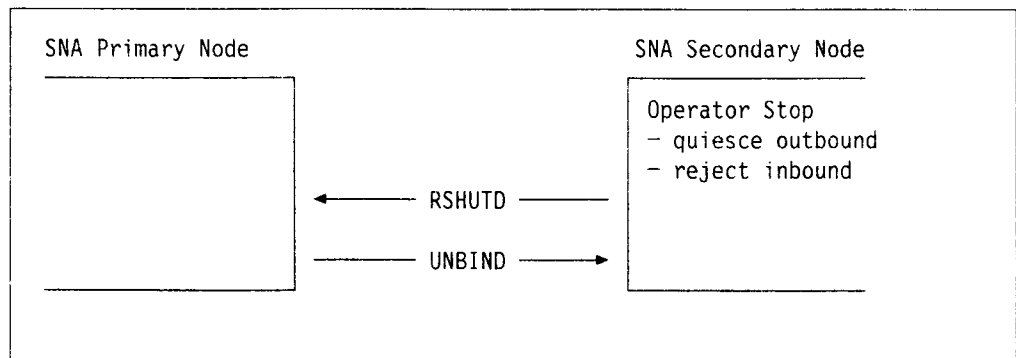


Figure 1-14. Terminating Processing by an SNA Secondary Node

The SNA secondary node requests the SNA primary node to shut down by sending a RSHUTD request. When the primary receives the request, it quiesces the outbound transmission buffers. When traffic has ended, the SNA primary node sends an UNBIND request. The SNA secondary node should not send any other data following the RSHUTD request.

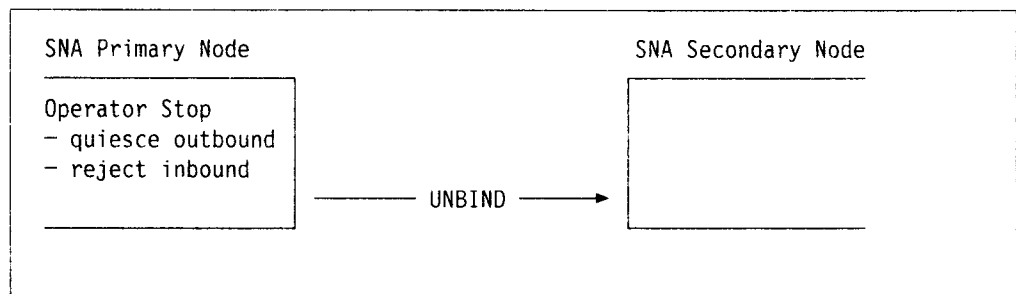


Figure 1-15. Terminating Processing by an SNA Primary Node

If the operator at the SNA primary node wants to terminate a session, the termination command generates an UNBIND request. If the SNA secondary node initiates a new transmission while the SNA primary node is quiescing the current transmission, an UNBIND request is sent to the secondary to deny the transmission.

Chapter 2. Formats of NJE Transfer Units

NJE Job	2-2
Contents of an NJE Job	2-2
Job Header	2-3
Job Header Prefix	2-4
Job Header General Section	2-4
JES2 Section of the Job Header	2-16
POWER Section of the Job Header	2-17
Job Scheduling Section of the Job Header	2-18
Security Section of the Job Header	2-19
Installation-Defined Section	2-20
Data Set Header	2-21
Data Set Header Prefix	2-22
Data Set Header General Section	2-22
3800 Section	2-32
RSCS Section	2-36
POWER Section	2-37
Record Characteristics Change Section	2-38
Output Processing Section	2-39
OPTB Structure	2-40
OPTU Key Definitions	2-42
Detailed OPTU Key Specifications	2-47
Error Handling	2-51
Security Section of the Data Set Header	2-52
Installation-Defined Section	2-53
Job Trailer	2-54
Job Trailer Prefix Section	2-54
Job Trailer General Section	2-54
Accounting Section	2-57
Installation-Defined Section	2-57
NJE Nodal Message Records (NMR)	2-58
Identifying the Contents of an NMR	2-58
Interpreting the Contents of the NMROUT Field	2-58
Format of an NMR that Contains an Unformatted Command	2-59
Format of an NMR that Contains a Formatted Command	2-61
Format of an NMR that Contains a Message	2-63
Interpreting the Contents of Field NMRMSG	2-63
Networking Connection Control Records (NCCR)	2-67
Initial Signon Record	2-67
Response Signon Record	2-68
Reset Signon Record	2-69
Concurrence Signon Record	2-69
Add/Subtract Connection Record	2-70
Signoff Records	2-70

This chapter describes the format of the control data used to transmit a NJE transfer unit across a network. The types of transfer unit a node can transmit are:

- An NJE job, either SYSIN or SYSOUT.
- An NJE Nodal Message Record (NMR) which is a record that contains a command or message.
- An NJE Connection Control Record (NCCR).

NJE Job

Contents of an NJE Job

An NJE job contains:

- either SYSIN or SYSOUT data
- the control records used to identify the data being transmitted

A node uses these control records to transmit an NJE job:

- **a job header.** See "Job Header" on page 2-3 for additional information on the job header.
- **a data set header** if the data is SYSOUT or optionally SYSIN data. If the NJE job contains a job and a SYSIN data set, a Record characteristics change section (RCCS) may be included in the NJE job to indicate a change in the length of the records. See "Data Set Header" on page 2-21 for additional information on the data set header.
- **a job trailer.** See "Job Trailer" on page 2-54 for additional information on the job trailer.

All NJE jobs contain both a job header and a job trailer. NJE SYSOUT jobs contain one or more data set headers before every SYSOUT data set.

The following sections describe the control records that are found in an NJE job. Each chart contains the:

- **Field name** of the fields found in the control record
- **Offset** or position of the field in the control record
- **Length** of the field or bit. A whole number represents a byte of data. If the length is a fraction, field describes a bit. For example, field NJHGFLG1 is a flag byte that contains 8 bits. The first two bits of field NJHGFLG1 is 2 bits and is reserved for IBM's use.
- **Type of Data** indicates if the data contains characters (char), binary data, or is a bit indicator.

If the field contains character data, the valid characters are A-Z, 0-9, #, @, and \$.

If the field is a flag byte, information in the type of data column can also indicate the position of the bit in the field. Each flag byte is composed of 8 bits. The high-order half word of the byte is represented by a 'x0' and the low-order half word of the byte is represented by a '0x', where 'x' is a hexadecimal digit between 0 and F. Refer to Figure 2-1 for an illustration of the bit notation used in the type of data column.

1...
Refers to bit 0. This bit is represented by a 80.

1.1.
Refers to bits 0 and 2. These bits are represented by a A0.

.... ...1
Refers to bit 8. This bit is represented by a 01.

.... ..11
Refers to bits 7 and 8. These bits are represented by a 03.

Figure 2-1. Example of the Bit Notation Used in the Type of Data Column

- **Default Value** indicates the default value of the field for an NJE job.
- **Valid Range** identifies the range of data the field can be in.
- **Description** indicates the use of the field.

Job Header

An NJE job must contain a job header to indicate the start of the job. The NJE job can contain either SYSIN or SYSOUT data.

The job header may include several sub-sections depending on the NJE product where the NJE job originated. All NJE products require a prefix section and a general section in the job header.

The following lists the additional sections that can be included in the job header for individual NJE products:

Figure 2-2 describes which sections in the header are built and used by each subsystem.

Figure 2-2. Job Header Sections used by Each Subsystem

Section	JES2	JES3	RSCS	POWER
General Section	built and used	built and used	built and used	built and used
Scheduling Section	built and used	built and used	not built or used	not built or used
Security Section	built and used with JES2 3.1.3	built and used with JES2 3.1.3	not built or used	not built or used
Product-specific Sections	JES2	- none -	- none -	POWER

Additionally, all NJE products may include an installation-defined section. Products may ignore sections that do not apply, but these sections are not be discarded.

Job Header

Job Header Prefix

Figure 2-3 describes the fields in the prefix section of the job header.

<i>Figure 2-3. Fields in the Prefix Section of a Job Header</i>						
Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJHLEN	0	2	binary	none	8-32764	<p>The entire length of the job header (including all sections).</p> <p>Note that although the valid length ranges for NJE headers (and individual sub-sections) are specified as being up to 32,764 bytes, the actual lengths for each (sub-)section supported by the NJE systems are defined as the offset of the last field defined in the section plus its length. The NJE systems are designed to receive sections with actual lengths greater than those defined here without error, but data beyond the defined lengths are ignored.</p> <p>Also, the NJE subsystems may have internal restrictions that limit the actual length of the header sections that can be received. For example, JES2 can not receive a job header that is bigger than a JES2 spool buffer minus the size of the JES2 JCT.</p>
NJHFLAGS	2	1				Reserved for IBM's use.
NJHSEQ	3	1	binary	0	none	<p>The sequence indicator. It contains valid information after the header is segmented for transmission specifying:</p> <ul style="list-style-type: none"> The position of the segment within the record Whether more record segments follow.

Job Header General Section

Figure 2-4 describes the fields in the general section of the job header.

<i>Figure 2-4 (Page 1 of 13). Fields in the General Section of a Job Header</i>						
Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGLN	0	2	binary	none	4-32764	Length of the job header general section.
NJHGTYPE	2	1	binary	none	none	Contains a X'00' to indicate a job header.
NJHGMOD	3	1	binary	0	none	Contains a X'00' to indicate the general section of the job header.
NJHGJID	4	2	binary	0	1-65535	<p>The job number assigned to the NJE job at the originating node.</p> <p>JES2 and JES3 use numbers 1-32767.</p> <p>RSCS uses numbers 1-9999.</p> <p>POWER uses numbers 1-65535.</p>
NJHGJCLS	6	1	char	A	A-Z 0-9	<p>The execution class associated with the job.</p> <p>JES2 sets the class from the job card, JECL statement, or input device class.</p> <p>JES3 and RSCS set the value to A for jobs that originate at the specified subsystem. The field is unused.</p> <p>POWER sets the job class from the job's JECL statements.</p>

Figure 2-4 (Page 2 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGMCLS	7	1	char	A	A-Z 0-9	Message class associated with the job. JES2 sets the message class from the job card. JES3 , RSCS , and POWER set the value to A but do not use the field.
NJHGFLG1	8	1	bits	0	none	Flag byte.
bit NJHGF1PR		.1	80 bit	0		The recompute selection priority bit. JES2 uses the bit if the submitter specifies the job's priority on either the /*PRIORITY JECL or JOB statement. If JES2 is receiving an NJE job and the bit is: <ul style="list-style-type: none">• off, JES2 gives the NJE job the installation's default priority.• on, JES2 uses the priority in NJHGPRIO as the NJE job's priority. Reserved for IBM's use
		.3	70 bits	0		Reserved for IBM's use
bit NJHGF1CF		.1	08 bit	0		Indicates whether notification of transmission forwarding is to be sent to the notify user (identified by NJHGNTYN/NJHGUSID). Forwarding is defined as the successful transmission of the job or SYSOUT to another node. Notification consists of a status message NMR indicating successful forwarding. NJHGF1CF set to 0 indicates that notification of forwarding may be sent from any intermediate node that processes this transmission. NJHGF1CF set to 1 indicates that notification of forwarding is <i>not</i> to be sent. JES2 does not use this flag. JES2 does not send notifications of transmission forwarding. JES3 does not use this flag. JES3 sends notifications of transmission forwarding on BSC links. RSCS sets the flag according to the SENTMSG parameter on the TAG command. POWER does not use this flag. POWER does not send notifications of transmission forwarding.

Figure 2-4 (Page 3 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
bit NJHGF1CA		.1	04 bit	0		<p>Indicates whether notification of transmission acceptance is to be sent to the notify user (identified by NJHGNTYN/NJHGUSID). Acceptance is defined as successful receipt (and storage) of the transmission at the destination node. Notification consists of a status message NMR indicating acceptance.</p> <p>NJHGF1CA set to 0 indicates that notification of acceptance may be sent from the destination node. NJHGF1CA set to 1 indicates that notification of acceptance is <i>not</i> to be sent.</p> <p>JES2 and JES3 do not use this flag. Notification of transmission acceptance is always sent if a notify user is specified.</p> <p>RSCS sets this flag according to the FINALMSG parameter on the TAG command.</p> <p>POWER does not use this flag. POWER always sends a notification of transmission acceptance.</p>
bit NJHGF1PE		.1	02 bit	0		<p>Indicates that the value in NJHGPASS has been encrypted by the security product at the originating node.</p> <p>JES2 and JES3 in MVS/SP 3.1.3 with RACF 1.9 use this (and the next) bit.</p> <p>RSCS and POWER do not use or set this (or the next) bit.</p>
bit NJHGF1NE		.1	01 bit	0		<p>Indicates that the value in NJHGNPAS has been encrypted by the security product at the originating node.</p> <p>Note:</p> <ol style="list-style-type: none"> 1. NJHGF1PE = 1 and NJHGF1NE = 0 indicates that NJHGPASS is encrypted and that NJHGNPAS is zeros. 2. NJHGF1PE = 0 and NJHGF1NE = 1 is not permitted. 3. See the security product for documentation on the encryption algorithm used.

Figure 2-4 (Page 4 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGPRIO	9	1	binary	0	0-F	<p>The selection/transmission job priority. 0 is the lowest priority and 15 is the highest job priority.</p> <p>JES2 uses this byte as the selection priority. If the JES2 node is receiving the NJE job and NJHGF1PR is on, JES2 uses the indicator as the execution priority. If the bit is off, JES2 ignores the indicator. If the JES2 node is transmitting an NJE job, JES2 uses NJHGPRIO to indicate the execution priority.</p> <p>JES3 does not use this byte because NJE jobs are transmitted FIFO.</p> <p>RSCS translates its priorities 99-0 to 0-15 on transmission and translates them again when the stream is received. Store-and-forward jobs are never altered, even if their priority is changed while on a VM/370 system.</p> <p>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</p> <p>POWER translates NJHGPRIO to the POWER priority (0-9) on received jobs and translates NJHGPRIO from the POWER priority on transmitted jobs. The priority is never changed for store-and-forward jobs, even if the operator alters the value on the POWER node.</p> <p>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</p>

Figure 2-4 (Page 5 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGORGQ	A	1	binary	0	none	<p>The system qualifier for the originating node in a loosely coupled multi-processor complex. The system qualifier is used to return status messages to the node where the NJE job was submitted.</p> <p>JES2 uses the indicator as the member number.</p> <p>JES3 uses the indicator as the index to a local main for TSO submitted jobs. If the job originated at a JES3 node and was submitted from a TSO user, the proper index value is set. A qualifier of 0 indicates the job is not from TSO. On input, the field is saved for generating TSO NOTIFY messages.</p> <p>RSCS does not set the field, but uses it for status messages.</p> <p>POWER sets the field from the POWER SYSID field. POWER uses the field to define the shared spooling system when running in a shared spooling complex. The field may contain X'40' or X'F1'-X'F9'.</p>
NJHGJCPY	B	1	binary	1	none	<p>The number of output copies for the entire NJE job as indicated before the job is executed.</p> <p>JES2 sets the field from the /*JOBPARM statement. At the destination node, this count is multiplied by the data set header copy count for each SYSOUT data set in the job.</p> <p>JES3 sets the field to zero at the origin node or defaults it if the execution node equals the origin node. Not used.</p> <p>RSCS sets the field to the default value. The field is unused.</p> <p>POWER sets the field to zero. The field is unused.</p>
NJHGLNCT	C	1	binary	0	none	<p>The number of lines per page for NJE SYSOUT jobs. A value of X'00' or X'FF' causes the system not to count lines or to use the default value at the node that prints the SYSOUT data set. Any other value causes a page eject after the specified number of lines.</p> <p>JES2 sets the lines per page for SYSOUT files from the /*JOBPARM or JOB statement. The special values X'00' and X'FF' are used as follows:</p> <ul style="list-style-type: none"> • X'00' -- Use the default value of the destination node. • X'FF' -- Do not force any page ejects. Let the skipping be solely determined by the carriage control (if present) in the SYSOUT data. <p>JES3 sets the field to the default value. The field is not used.</p> <p>RSCS sets the value to X'FF' for printing and sets the value to the default for punching. The field is not used.</p> <p>POWER uses the default value.</p>
	D	1				Reserved for IBM's use

Figure 2-4 (Page 6 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGHOPS	E	2	binary	0	1-65535	<p>This field contains a count of the number of nodes (which support this field) that have processed (i.e., stored and forwarded) this transmission. It may be used to detect looping of a transmission in a network.</p> <p>Set to zero at the origin node (and not reset for NJE SYSOUT jobs generated from this job) and incremented at intermediate nodes.</p> <p>JES3 sets the field to zero at origin nodes for SYSIN jobs and increments the field at intermediate nodes, but does not use the field.</p> <p>JES2 and POWER set the default value, but do not increment nor use the field.</p>
NJHGACCT	10	8	char	blanks	none	<p>The job's network accounting information.</p> <p>JES2 sets this field from the /*NETACCT statement or converts the local account number using the NETACCT translation tables when transmitting data. JES2 uses this field as input to the NETACCT translation tables. Blanks can be accepted as the default when data is being received.</p> <p>JES3 obtains accounting information from the /*NETACCT control statement.</p> <p>RSCS does not set nor use the field.</p> <p>POWER sets, but does not use the field.</p>
NJHGJNAM	18	8	char	none	none	<p>The name of the job.</p> <p>JES2 initializes the field with the job name, which can contain blanks.</p> <p>JES3 initializes the field with the job name. If a job name is not supplied, NJEJOB is used as the default.</p> <p>RSCS initializes the field with the job name. For networking jobs, RSCS uses RSCSnnnn where nnnn is the spool file number of the job.</p> <p>POWER initializes the field to the POWER job name. If the field is blank when POWER is receiving a job, POWER sets the field to JOB appended with the value specified in NJHGJID.</p>

Figure 2-4 (Page 7 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGUSID	20	8	char	blanks	none	<p>This field contains the destination userid at the NJHGNTYN node for status notification message NMRs associated with the job or SYSOUT transmission.</p> <p>JES2 obtains the information from either the NOTIFY parameter on the JOB statement or the userid specified on the /*NOTIFY JECL statement. Set to the origin userid if none of the above is specified.</p> <p>JES3 obtains the information either from the userid of the person who submitted the request, if the request was submitted using TSO, or from the /*NETACCT control statement. Otherwise, sets the field to blanks.</p> <p>RSCS uses this field as the origin userid if NJHGORGR is blanks or zeros. Set from the origin CMS userid.</p> <p>POWER sets the field with the notify userid specified on the NTFY control statement.</p>
NJHGPASS	28	8	char or binary	0	none	<p>The password for the job used to validate the user (specified in NJHGORGR) at the execution node.</p> <p>JES2 obtains the value from the JOB statement. JES2 blanks the field when transmitting an NJE SYSOUT job and does not use the field if receiving an NJE SYSOUT job.</p> <p>JES3 obtains the value from the job's JOB JCL statement.</p> <p>RSCS sets the default to binary zeros but does not use the field.</p> <p>POWER sets the field to the value specified on the job's JCL JOB statement and uses the password for security purposes.</p>
NJHGNPAS	30	8	char or binary	0	none	<p>The new password.</p> <p>JES2 and JES3 obtain the new password from the JOB statement. Not used or set for SYSOUT jobs.</p> <p>RSCS uses the binary zeros default value for the new password.</p> <p>POWER does allow the user to specify a new password for the job. It sets the field to zeros if not specified.</p>
NJHGETS	38	8	binary	0	none	<p>The time and date the job entered the system in S/370 STCK (store clock) format. The time stamp reflects the Greenwich Mean Time (GMT).</p> <p>JES2 and JES3 set and use the time stamp for the System Management Facility (SMF).</p> <p>POWER sets the time stamp but never uses it.</p>
NJHGORGN	40	8	char	none	none	<p>The name of the node where the job originated (where the user submitted the job).</p> <p>NJHGORGN must never be modified during the life of the NJE job. It may be used as the true origin of the job for security authorization and auditing purposes.</p>

Figure 2-4 (Page 8 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGORGR	48	8	char	blanks	none	<p>The identifier of the user who submitted the job or the name of the remote workstation from which the job was submitted (at node NJHGORGN).</p> <p>JES2 and JES3 set the field as follows for NJE SYSIN jobs:</p> <ul style="list-style-type: none"> Set to submitter's userid if submitted via internal reader (e.g., TSO userid if TSO internal reader), Set to remote name if submitted via RJE workstation, Set to input device group name if submitted via input device (JES3 only), Set to default otherwise. <p>JES2 and JES3 set the field as follows for NJE SYSOUT jobs:</p> <ul style="list-style-type: none"> Propagated from job header of job creating SYSOUT by JES2, Set to zeros by JES3.
NJHGXEQN	50	8	char	blanks	none	<p>The name of the node where an NJE SYSIN job is sent for execution.</p> <p>If an operator issues a command to reroute the NJE SYSIN job, the value in this field will be changed to indicate the new node name.</p> <p>JES2 sets the execution node name for an NJE SYSIN job from the:</p> <ul style="list-style-type: none"> • /*ROUTE XEQ statement • /*XEQ statement • /*XMIT statement <p>JES3 sets the execution node name for an NJE SYSIN job from the:</p> <ul style="list-style-type: none"> • //ROUTE XEQ statement • // XMIT statement
NJHGXEQU	58	8	char	blanks	none	<p>The execution userid.</p> <p>JES2 sets the execution userid from the /*XEQ or /*XMIT JECL statement. JES2 does not use the execution userid.</p> <p>JES3 sets the execution userid from the //ROUTE XEQ or // XMIT JECL statement. JES3 does not use the execution userid.</p> <p>RSCS uses the execution userid as the destination virtual machine for SYSIN jobs, and sets it to zeros. For NJE SYSOUT jobs, RSCS sets this field as the origin userid.</p>

Figure 2-4 (Page 9 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGPRTN	60	8	char	blanks		<p>The default destination where the system sends print files that are not specifically routed by the job to be printed.</p> <p>JES2 obtains the default destination from the job's JCL or from the statement used to define the reader where the job was submitted.</p> <p>JES3 sets the field from the job's JECL and uses it. It uses the originating node as the default destination for NJE SYSIN jobs. To cause the default for all SYSOUT from an NJE SYSIN job to be the default print node, the user can insert a JES3 /*FORMAT statement into the job's JECL. However, the default destination node must not be blanks or the same as the originating node.</p> <p>RSCS sets the field to the originating node and uses this value as the print destination node for NJE SYSIN jobs. For NJE SYSOUT jobs, RSCS sets the field to the destination node and uses this value as the print destination node.</p> <p>POWER sets the field to the originating node or to the value specified on the LDEST parameter.</p>
NJHGPRTR	68	8	char	blanks		<p>The default userid or workstation for the output created by the job.</p> <p>JES2 obtains the default destination from the READER initialization statement or from the PRDEST parameter on the statement that defines the reader where the job was submitted. JES2 uses the destination at the execution node to route the output to if a node is not specified. JES2 uses this field to initialize the destination (field NDHGRMT) in the data set header.</p> <p>JES3 uses the originating node as the default destination for NJE SYSIN jobs. To cause the default for all SYSOUT from an NJE SYSIN job to be the default print node, the user can insert a JES3 /*FORMAT statement into the job's JECL. However, the default destination node must not be blanks or the same as the originating node.</p> <p>POWER sets the field from the originating userid (or remote workstation) or sets it to the value specified on the LDEST parameter.</p>

Figure 2-4 (Page 10 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGPUNN	70	8	char	blanks		<p>The default punch destination node for punch files that are not specifically routed to another punch destination node.</p> <p>JES2 sets the default punch destination node from either the job's JECL or the PUNODE parameter on the reader where the job was submitted. If neither is specified, JES2 sets the punch destination node to the originating node.</p> <p>JES3 sets default punch destination node to the originating node for NJE SYSIN jobs. To cause the default for all SYSOUT from an NJE SYSIN job to be the default print node, the user can insert a JES3 // *FORMAT statement into the job's JECL. However, the default destination node must not be blanks or the same as the originating node.</p> <p>RSCS sets the default punch destination node to the originating node for NJE SYSIN jobs and to the destination node for NJE SYSOUT jobs.</p> <p>POWER sets the default punch destination node from the originating node or from the node specified on the PDEST parameter.</p>
NJHGPUNR	78	8	char	blanks		<p>The default punch output destination (either a userid or a remote workstation).</p> <p>JES2 sets the default punch output destination from either the:</p> <ul style="list-style-type: none"> • /*ROUTE PUNCH statement • PUDEST parameter on the Rnnnn.RDm statement. <p>JES2 uses the default punch output destination at the execution node to:</p> <ul style="list-style-type: none"> • Route SYSOUT if a destination is not specified • Set field NDHGRMT in the the data set header. <p>JES3 sets the default punch output destination to the remote originating identifier when transmitting an NJE SYSIN job. If JES3 is receiving an NJE SYSIN job and the default punch output destination is not the same as the originating workstation, the user can insert a JES3 // *FORMAT statement into the JECL to cause the default for all punch output to be the default punch node and remote workstation.</p> <p>RSCS sets the field to the destination userid or the destination remote workstation. If specified destination is SYSTEM, X'00' is set in the field. For NJE SYSIN jobs, RSCS sets the field to the originating userid or X'00' if the NJE SYSIN job originated from a remote workstation.</p> <p>POWER sets the default punch output destination to:</p> <ul style="list-style-type: none"> • The originating userid • The originating remote workstation • The values specified for the PDEST parameter on the JOB statement.

Figure 2-4 (Page 11 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGFORM	80	8	char	blanks		<p>The default form characteristics for any output created by the job during execution. It is copied into field NDHGFORM in the data set header during execution.</p> <p>JES2 sets the default form characteristics at the destination node from the job's JECL or from initialization statements.</p> <p>JES2 does not use the field.</p> <p>JES3 and RSCS do not use the field but set it to zeros.</p> <p>POWER does not use or set the field.</p>
NJHGICRD	88	4	binary	0		<p>The number of statements in the NJE SYSIN job.</p> <p>JES3 sets the field to the default value and does not use it.</p> <p>RSCS sets the field to the number of records in the file for NJE SYSIN jobs but does not use the field.</p> <p>POWER sets the field to the transmitted record count for the NJE job but does not use the field.</p>
NJHGETIM	8C	4	binary	0		<p>The estimated CPU time it will take to execute the job.</p> <p>JES2 sets the time from the job's JECL and uses the field for SMF records.</p> <p>JES3 sets the time from the job's JECL and uses it in SMF type 26 records.</p> <p>RSCS and POWER set the field to the default but does not use it.</p>
NJHGELIN	90	4	binary	0		<p>The estimated number of output print lines.</p> <p>JES2 sets the field from the job's JECL and uses it for SMF accounting.</p> <p>JES3 sets the field from the job's JECL and uses it for SMF accounting.</p> <p>RSCS does not use or set the field.</p> <p>POWER does not use the field but sets it to the default.</p>
NJHGECRD	94	4	binary	0		<p>The estimated punch card output.</p> <p>JES2 sets the field from the job's JECL and uses it for SMF accounting.</p> <p>JES3 sets the field from the job's JECL and uses it for SMF accounting.</p> <p>RSCS and POWER do not use the field but set it to the default.</p>
NJHGPRGN	98	20	char	blanks		<p>The name of the programmer specified on the JOB statement.</p> <p>JES3 sets the name of the submitter from the <code>/*NETACCT</code> statement.</p> <p>RSCS sets the name of the submitter for NJE SYSOUT jobs. It is set to X'00' for NJE SYSIN jobs. RSCS does not use the field.</p> <p>POWER uses the field for information that is printed on the separator pages for the NJE job's SYSOUT data.</p>

Figure 2-4 (Page 12 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGROOM	AC	8	char	blanks		<p>The room number of the programmer as specified on the JOB statement.</p> <p>JES2 sets the room number from the JOB statement or the /*JOBPARM statement.</p> <p>JES3 sets the room number from the /*NETACCT statement.</p> <p>RSCS sets the field from the VM/SP distribution code.</p> <p>POWER uses the field for information printed on the separator pages for the NJE job's SYSOUT data.</p>
NJHGDEPT	B4	8	char	blanks		<p>The programmer's department number as specified on the JOB statement.</p> <p>JES2 and RSCS set the field with the default but do not use it.</p> <p>JES3 sets the field with information on the /*NETACCT statement.</p> <p>POWER uses the department number on the separator pages of the NJE job's SYSOUT.</p>
NJHGBLDG	BC	8	char	blanks		<p>The programmer's building number as specified on the JOB statement.</p> <p>JES2 and RSCS set the field with the default but do not use it.</p> <p>JES3 sets the field with information on the /*NETACCT statement.</p> <p>POWER uses the department number on the separator pages of the NJE job's SYSOUT.</p>
NJHGPREC	C4	4	binary	0		<p>The record count for an NJE SYSOUT job. It is not set for an NJE SYSIN job.</p> <p>JES2 sets the field to the sum of record counts for each data set after the job header. The sum does not include the job headers or data set headers in the NJE SYSOUT job. It is used by the node receiving the NJE SYSOUT job to issue a message acknowledging receipt of the job.</p> <p>JES3 sets the field to the default and does not use it.</p> <p>RSCS sets the field from the number of records in the file for NJE SYSOUT jobs. RSCS does not use the field for NJE SYSOUT jobs. For NJE SYSIN jobs, RSCS does not set or use the field.</p> <p>POWER uses the field as the line count for the output file. Records containing spanned data count as 1 record.</p>
	C8	4	binary	0		Reserved for IBM's use.

Job Header

Figure 2-4 (Page 13 of 13). Fields in the General Section of a Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHGNTYN	CC	8	char	blanks		<p>This is the node name of the user to notify about job/SYSOUT transmission or processing status</p> <p>JES2 sets this field to the node name specified via the /*NOTIFY JECL statement.</p> <p>JES3 sets this field to the node name specified via the /* MAIN JECL statement.</p> <p>RSCS does not set this field.</p> <p>POWER sets this field to the node name specified via the NTFY statement.</p>

JES2 Section of the Job Header

Figure 2-5 describes the format of the JES2 section of the job header.

Figure 2-5. Fields in the JES2 Section of the Job Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJH2LEN	0	2	binary	none	4-32764	Length of the JES2 section of the job header.
NJH2TYPE	2	1	binary	none	none	The major type of this section. The JES2 section of the job header should contain a X'84'.
NJH2MOD	3	1	binary	0	none	Modifier for the major type of this section (NJH2MOD). The modifier code for the JES2 section of the job header is X'00'.
NJH2FLG1 Bit NJH2FJOB	4	.2	bits 03 bit	0	none	Job level flags. The NJE job was submitted as a batch job if the field is X'00'.
Bit NJH2FSTC		.1	01 bit			The NJE job was submitted under a started task if this bit is on.
Bit NJH2FTSU		.1	02 bit			A time-sharing user submitted the NJE job if this bit is on.
NJH2USE		.1	04 bit			SMF Userid present in NJH2USID.
NJH2ACCT	8	4	char	0	none	The submitter's account number.
NJH2USID	C	8	char	blanks	none	Address of the SMF record. NJE does not use this field; it is used by the JES2 spool offload facility.
NJH2USR	14	8	char	binary zeros	none	The userid specified on the JCL used to submit the NJE job. NJE does not use this field; it is used by the JES2 spool offload facility.
NJH2GRP	1C	8	char	binary zeros	none	The group identifier specified on the JCL used to submit the NJE job. NJE does not use this field; it is used by the JES2 spool offload facility.
NJH2SUSR	24	8	char	binary zeros	none	The userid of the user who submitted the NJE job. NJE does not use this field; it is used by the JES2 spool offload facility.
NJH2SGRP	2C	8	char	binary zeros	none	The group identifier of the user who submitted the NJE job. NJE does not use this field; it is used by the JES2 spool offload facility.

POWER Section of the Job Header

Figure 2-6 describes the format of the POWER section of the job header.

Figure 2-6 (Page 1 of 2). Fields in the POWER Section of the Job Header						
Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHPLEN	0	2	binary	none	4-32764	Length of the POWER section of the job header.
NJHPTYPE	2	1	binary	none	none	The major type of this section. The POWER section of the job header contains a X'86'.
NJHPMOD	3	1	binary	0	none	Modifier of the major type of this section (NJH2MOD). The modifier code for the POWER section of the job header is X'00'.
NJHPFLG1	4	1	bits	0	none	Reserved for IBM's use
NJHPDISP	5	1	char	D	none	The job's disposition. POWER obtains the job's disposition from the * \$\$ JOB statement.
	6	1	binary	0		Reserved for IBM's use.
NJHPSYID	7	1	char	blank	none	The receiving node's system qualifier. POWER uses the system qualifier only if the node uses a shared spool environment.
NJHPUSER	8	16	char	blanks	none	The user information specified on the * \$\$ JOB statement.
NJHPDSKT	18	2	binary	0	none	The diskette address.
	1A	2	binary	0		Reserved for IBM's use.
NJHPDD	1C	15	binary	0	none	Due date information indicating the date and time when a SYSIN job is to be scheduled. For a SYSOUT job, this field contains zeros. This field includes the remaining fields in the section.
NJHPDGP1 Bit NJHPDG1R	1C	1	binary	0	none	Due date flags. The job must be rerun if the due date expired during the time the system was down. Reserved for IBM's use. The job should run daily. Weekdays are specified as due dates. Numbers are specified as due dates. Reserved for IBM's use.
Bit NJHPDG1T		.3	70 bits	0	none	
Bit NJHPDG1W		.1	08 bit	0	none	
Bit NJHPDG1W		.1	04 bit	0	none	
Bit NJHPDG1D		.1	02 bit	0	none	
		.1	01 bit	0	none	
NJHPDGP2	1D	1	binary	0	none	Due date flags. Reserved for IBM's use.
Bit NJHPDG2X		.7	FE bits	0	none	
		.1	01 bit	0	none	Information about the due date exists.
NJHPDCY	1E	6	binary	0	none	This field indicates if the job has to be scheduled more than once (for example, every Monday or every first of every month). This field includes the following two fields.
NJHPDMD	1E	2	binary	0	none	The months in which a job is to be scheduled. X'80' = January, X'40' = February, ...
NJHPDDM	20	4	binary	0	none	Days of a month on which a job is to be scheduled: X'80' = first, X'40' = second, ... or, if weekdays have been specified: X'80' = Monday, X'40' = Tuesday,
NJHPDN	24	6	binary	0	none	The date and time at which the job is to be scheduled (may be zero, if not yet calculated). If a job is to be scheduled more than once, a new scheduling date is calculated according to the information in NJHPDCY. This field includes the next five fields.

Job Header

Figure 2-6 (Page 2 of 2). Fields in the POWER Section of the Job Header

Field Name	Offset in Hex	Length in Dec.	Type of Data	Default Value	Valid Range	Description
NJHPDNY	24	2	binary	0	1988-2087	The year in which the job is to be scheduled (in packed decimal format without sign).
NJHPDNM	26	1	binary	0	1-12	The month in which the job is to be scheduled (in packed decimal format without sign).
NJHPDND	27	1	binary	0	1-31	The day in which the job is to be scheduled (in packed decimal format without sign).
NJHPDNTH	28	1	binary	0	0-23	The hour at which the job is to be scheduled (in packed decimal format without sign).
NJHPDNM	29	1	binary	0	0-59	The minute at which the job is to be scheduled (in packed decimal format without sign).
	2A	1	binary	0		Reserved for IBM's use.
	2B	1	binary	0		Reserved for IBM's use.

Job Scheduling Section of the Job Header

Figure 2-7 describes the format of the job scheduling section of the job header. The job scheduling section is used by JES2 and JES3 only.

Figure 2-7. Fields in the Job Scheduling Section of the Job Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJHELEN	0	2	binary	none	4-32764	Length of the job scheduling section of the job header.
NJHETYPE	2	1	binary	none	none	The major type of this section. The job scheduling section of the job header should contain X'8A'.
NJHEMOD	3	1	binary	0	none	Modifier of the major type of this section (NJH2MOD). The modifier code for the job scheduling section of the job header is X'00'.
NJHEPAGE	4	4	binary	0	none	The estimated number of "begin page" structured fields for page-mode SYSOUT data sets. Each "begin page" record starts with the sequence X'D3A8AF'. JES3 initializes the field (after execution) to the number of actual pages created by the job.
NJHEBYTE	8	4	binary	0	none	The estimated number of output bytes. This includes all the bytes in the job's SYSOUT data set. JES3 initializes the field (after execution) to the number of actual bytes written to spool by the job.

Security Section of the Job Header

Figure 2-8 describes the format of the Security Section of the Job Header.

The section is identified by an identifier field of X'8C' and a modifier field of X'00'. It may be built as part of a created NJE work element (job or SYSOUT), or may be added to a forwarded work element.

The security section contains information related to the identification of the owner or submitter user associated with the NJE job (SYSIN or SYSOUT). The information in this section is to be treated as a token that is to be interpreted by the security facility component of the trusted computer base.

Figure 2-8 (Page 1 of 2). Fields in the Security Section of the Job Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJHTLEN	0	2	binary	84	4-32764	Length of the job security section.
NJHTTYPE	2	1	binary	X'8C'	none	The major type identifier.
NJHTMOD	3	1	binary	0	none	Modifies the major type identifier.
NJHTLENP	4	2	binary	4	4-32764	Length of the prefix sub-section.
NJHTFLG0 Bit NJHTF0JB	6	1 .1	bits 80 bit	0 0	none none	This flag byte defines the following flags: Indicates that the security section represents the owner of the NJE job. The bit not set indicates that the security section represents the submitter of the NJE job.
RESERVED	7	1		0	none	Reserved for IBM's use.
	8					End of prefix sub-section
NJHTLENT	0	1	binary	80	80-255	Length of the security data sub-section.
NJHTVERS	1	1	binary	X'01'	1-255	Security section version.
NJHTFLG1 Bit NJHT1EN Bit NJHT1EXT	2	1 .1 .1	bits 80 bit 40 bit	0 0 1	none none none	This flag byte defines the following flags: Indicates that the security section (from NJHTSTYP on) is encrypted. Encryption algorithm is private to the security product. Indicates the format of the security section. B'1' indicates external format which is used exclusively in NJE.
NJHTSTYP	3	1	binary	X'07'	1-255	Security session type. '7' indicates batch job type.
NJHTFLG2 Bit NJHT2DFT Bit NJHT2MLO Bit NJHT2SHI Bit NJHT2TRS Bit NJHT2SUS Bit NJHT2RMT	4	1 .1 .1 .1 .1 .1	bits 80 bit 20 bit 10 bit 08 bit 04 bit 02 bit	1 0 0 0 0 1	none none none none none	This flag byte defines the following flags: Indicates that the security section data has not been verified by the security product at the node which created the security section. If not set, the security data may be assumed to have been verified. Indicates that all required ML options were active at the node which created the security section. Indicates that the security data could not be verified at the node creating the security section. If set, the user to which this section applies is to be considered as part of the trusted computer base. If set, the security data was obtained via a surrogate user at the node of creation. Indicates that the job or data set was originated remotely.

Figure 2-8 (Page 2 of 2). Fields in the Security Section of the Job Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJHTPOEX	5	1	binary	X'03'	1-255	Indicates the port of entry class index. The following port of entry class indices are defined: 1 = Terminal Class 2 = Console Class 3 = JESINPUT Class
RESERVED	6	2				Reserved for IBM's use.
NJHTSECL	8	8	char	none	A-Z,0-9	The security label associated with the user represented by this section.
NJHTCNOD	10	8	char	none	A-Z,0-9	Node at which security section was generated.
NJHTSUSR	18	8	char	none	A-Z,0-9	Userid of submitting user.
NJHTSNOD	20	8	char	none	A-Z,0-9	Node of job submission.
NJHTSGRP	28	8	char	none	A-Z,0-9	Groupid of submitting user.
NJHTPOEN	30	8	char	none	A-Z,0-9	Port of entry name at the creation node for this job.
RESERVED	38	8				Reserved for IBM's use.
NJHTOUSR	40	8	char	none	A-Z,0-9	Security session owner userid.
NJHTOGRP	48	8	char	blanks	A-Z,0-9	Security session owner groupid.

Installation-Defined Section

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

Beyond the four-byte definition, this section is the responsibility of the installation. The section length is limited to 32764 bytes. The combined length of all sections in the header is limited to 32764 bytes, 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. The following depicts a sample eight-byte field:

Figure 2-9. Fields in the Job Header Installation Extension Section

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJHULEN	0	2	binary	none	4-32764	Length of the section.
NJHUTYPE	2	1	binary	none	C0-FF	The major type of this section.
NJHUMOD	3	1	binary	none	none	Modifier for the major type identifier.
NJHUCODE	4	4	char	none	none	This could be used for a GUIDE/SHARE installation code, or some other unique identifier.

Data Set Header

An NJE job begins with a job header, is followed by data, and ends with a job trailer. The data between the job header and the job trailer can be either SYSIN or SYSOUT data.

An NJE SYSIN job contains JCL for a job and one or more SYSIN data sets. It may or may not contain data set headers. If the SYSIN data set has record characteristics other than 80-byte fixed records, then it is preceded by a data set header containing only a record characteristics change section. See "Record Characteristics Change Section" on page 2-38.

An NJE SYSOUT job contains one or more SYSOUT data sets. Each SYSOUT data set is preceded by one or more data set headers. A data set header contains a prefix section and a basic general section, which may be followed by any of the following sections:

- **3800 Section** -- for output to be printed on a 3800 printer.
- **RSCS Section** -- for output created at a VM/RSCS node.
- **POWER Section** -- for output created at a POWER node.
- **Output Processing Section** -- for advanced function printing output.
- **Security Section** -- for resource authorization control.

Figure 2-10 describes which sections in the header are built and used by each subsystem.

Figure 2-10. Data Set Header Sections used by Each Subsystem

Section	JES2	JES3	RSCS	POWER
General Section	built and used for SYSOUT only	built and used for SYSOUT only	built and used for SYSOUT only	built and used for SYSOUT only
Record Change Characteristics Section	built and used for SYSIN only if the record length (LRECL) is not greater than 80 bytes, or if the record format (RECFM) is not fixed format.	not built, used, nor forwarded	not built, but used for forwarding SYSIN	built and used
3800 Section	built and used for 3800 attributes	built and used for 3800 attributes	built and used for 3800 attributes	built and used for 3800 attributes
Output Processing Section	built and used for AFP attributes	built and used for AFP attributes	built and used for AFP attributes	built and used for AFP attributes
Security Section	built and used with JES2 3.1.3	built and used with JES2 3.1.3	not built nor used	not built nor used
Product-specific Sections	- none -	- none -	RSCS	POWER

Data Set Header

Data Set Header Prefix

The data set header begins with the standard four-byte control record prefix containing the combined length of all sections. Figure 2-11 describes the format of the prefix section of the data set header.

Figure 2-11. Fields in the Prefix Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHLEN	0	2	binary	none	4-32764	The Combined length of all sections in the data set header.
NDHFLAGS	2	1	binary	none	none	Reserved for IBM's use.
NDHSEQ	3	1	binary	none	none	The sequence indicator. It contains valid information after the header is segmented for transmission. The information specifies: <ul style="list-style-type: none"> The position of the segment within the record. More record segments follow.

Data Set Header General Section

The basic general section is identified by an identifier field of X'00' and a modifier field of X'00'. Figure 2-12 describes the format of the basic general section of the data set header.

Figure 2-12 (Page 1 of 10). Fields in the General Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHGLEN	0	2	binary	none	4-32764	Length of the data set header general section.
NDHGTYPE	2	1	binary	none	none	The major type of this section. General sections are identified by a X'00'.
NDHGMOD	3	1	binary	0	none	Modifier for the major type identifier. A modifier code of X'00' is the modifier for the general section.
NDHGNODE	4	8	char	none	N/A	The destination node for the data set. This field may be set at execution time from either NJHGPRTN or NJHGPUNN.
NDHGRMT	C	8	char	blanks	N/A	Either the destination userid, remote workstation, or printer identifier. This field may be set at execution time from either NJHGPRTN or NJHGPUNN. JES2 and JES3 also use this field for the TSO/E TRANSMIT and RECEIVE commands. If this field and the NDHGXWTR field are the same, this field is used as the TSO destination userid. RSCS sets the field to the destination userid or printer (or X'00' when SYSTEM is the destination). POWER uses the field to route the NJE job to a remote workstation, userid, or program.
NDHGPROC	14	8	char	blanks	none	The JCL procedure that was executing when this data set was created. RSCS sets the field from the information in FILENAME. RSCS does not use the field. POWER defaults this field and does not use it.

Figure 2-12 (Page 2 of 10). Fields in the General Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHGSTEP	1C	8	char	blanks	none	<p>The step name that was executing when the data set was created.</p> <p>RSCS sets the field from the information in FILETYPE. RSCS does not use the field.</p> <p>POWER defaults this field and does not use it.</p>
NDHGDD	24	8	char	blanks	none	<p>The DDNAME that referenced the data set when it was created.</p> <p>JES3 If blank on input, JES3 uses the default DDNAME of 'NJEWKnn', where nn is an internal index value.</p> <p>RSCS sets the field to the VM distribution code and does not use the field.</p> <p>POWER defaults the field and does not use it.</p>
NDHGDSNO	2C	2	binary	0	0-32767	<p>A counter incremented each time a SYSOUT data set is allocated within the same job. This field maintains the uniqueness of data sets, while allowing them to be spun off.</p> <p>JES2 uses and sets the field as follows:</p> <ol style="list-style-type: none"> 1. Ignores it on input if it is greater than 100 (the range for JES2 reserved keys). 2. Sets it from the PDBDSKEY field if less than 32768. If larger than 32767, then it will be set to zero. <p>JES3 defaults and does not use this field.</p> <p>RSCS defaults the field and does not use it.</p> <p>POWER defaults the field and does not use it.</p>
	2E	1				Reserved for IBM's use.
NDHGCLAS	2F	1	char	A	A-Z 0-9	<p>The job's SYSOUT class.</p> <p>JES2 sets the SYSOUT class to A if the class specified in the job's JCL is not alphanumeric.</p> <p>JES3 sets the field to the SYSOUT class. If receiving an NJE SYSOUT job, this field is used if the class and type (indicated by NDHGF2PR and NDHGF2PU) do not conflict with the local system's use of class/type.</p> <p>If a conflict exists, a default class is used for the appropriate type.</p> <p>RSCS uses the field for the VM output class and sets it from this class unless overridden by the SYSOUT keyword on the TAG.</p> <p>POWER sets the SYSOUT class to the output class. Only A-Z are valid (set to A if not in the range A-Z when received).</p>

Data Set Header

Figure 2-12 (Page 3 of 10). Fields in the General Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHGPREC	30	4	binary	0	none	<p>The count of the number of records in the data set.</p> <p>JES2 uses the field as the line count of the data set. A spanned record counts as 1, regardless of the number of segments.</p> <p>JES3 sets the number of records in the data set but does not use the information.</p> <p>RSCS uses the information in a message to the user. RSCS sets the field from the record count in the spool file block.</p> <p>POWER sets and uses the field as the line count.</p>
NDHGFLG1	34	1	bits	0	none	<p>A flag byte containing the following flags:</p>
Bit NDHGF1SP		.1	80 bit	0	none	<p>Is the spin data set flag.</p> <p>JES2 sets the bit from information in the job's JCL. JES2 sometimes sets this bit for the TSO/E TRANSMIT and RECEIVE commands according to the data set disposition.</p> <p>JES3 does not use this bit but sets it for spin data sets and data sets generated by the TSO/E TRANSMIT and RECEIVE commands.</p> <p>RSCS uses the bit to separate data sets into individual spool files, but does not set this bit.</p> <p>POWER sets the bit only if the data set is segmented and uses the bit to determine the job number. If the flag is on, the job number is set from the job header (NJHGJID), otherwise a new job number is assigned.</p>
Bit NDHGF1HD		.1	40 bit	0	none	<p>The destination node is to hold the data set.</p> <p>JES2 sets the bit from information in the JCL and uses it to place the data set on the system output hold queue.</p> <p>JES3 sets the bit and uses it to place the data set on the system output hold queue.</p> <p>RSCS uses the bit to indicate to the destination node that it is to hold the data set. The bit is set when specified by the originator on the TAG command using the HOLD keyword.</p> <p>POWER uses the bit and sets it if the data set disposition is 'H' or 'L'.</p>
Bit NDHGF1LG		.1	20 bit	0	none	<p>The job log indicator bit.</p> <p>JES2 uses the bit to indicate that the data set is the JES2 job log.</p> <p>JES3, RSCS, and POWER do not use the bit but set it to the default value.</p>

Figure 2-12 (Page 4 of 10). Fields in the General Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description																																								
Bit NDHGF1OV		.1	10 bit	0	none	<p>The page overflow indicator. If the bit is on, the data set is allowed to print over the page fold. If off, the printer should start at the top of each page.</p> <p>JES2 and RSCS do not use this bit but set it with the default value.</p> <p>JES3 sets the bit if OVFL = OFF is specified on the // *FORMAT statement in the job's JCL. JES3 uses the bit as described above.</p> <p>POWER uses the field but does not set it.</p>																																								
Bit NDHGF1IN		.1	08 bit	0	none	<p>The interpreter feature is to be used for punched output if punched on a device that has the interpreter feature installed.</p> <p>Note: For further interpretation of this field, see the NDHGF1DF field (01 bit).</p> <p>JES2 and JES3 use the information from the job's JCL to set the bit.</p> <p>RSCS and POWER do not use this bit but set it to the default value.</p>																																								
Bit NDHGF1LC		.1	04 bit	0	none	<p>The NDHGLNCT field (offset X'3A') contains information if this bit is on.</p> <p>JES2 sets and uses the bit.</p> <p>JES3, RSCS, and POWER set this bit to the default value, but do not use it.</p>																																								
Bit NDHGF1ST		.1	02 bit	0	none	<p>If on, the job log data set contains statistics for the job.</p> <p>JES2 uses the bit to avoid multiple instances of the job statistics in the job log.</p> <p>JES3, RSCS, and POWER do not use this bit, but set it to the default value.</p>																																								
Bit NDHGF1DF		.1	01 bit	0	none	<p>Specifies whether the interpret bit (NDHGF1IN) was explicitly set off (1) or not (0). The NDHGF1DF is used with the NDHGF1IN bit as follows:</p> <table style="margin-left: 40px;"> <tr> <td>NDHGF1IN</td> <td>0</td> <td>1</td> <td>0</td> <td>1</td> </tr> <tr> <td>NDHGF1DF</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> </tr> <tr> <td></td> <td> </td> <td> </td> <td> </td> <td> </td> </tr> <tr> <td></td> <td>v</td> <td> </td> <td>v</td> <td> </td> </tr> <tr> <td>Meaning:</td> <td>Use</td> <td> </td> <td>Do not</td> <td> </td> </tr> <tr> <td></td> <td>default</td> <td> </td> <td>interpret</td> <td> </td> </tr> <tr> <td></td> <td></td> <td>v</td> <td></td> <td>v</td> </tr> <tr> <td></td> <td></td> <td>Interpret</td> <td></td> <td>Interpret</td> </tr> </table> <p>JES2, RSCS, and POWER default this bit and do not use it.</p>	NDHGF1IN	0	1	0	1	NDHGF1DF	0	0	1	1							v		v		Meaning:	Use		Do not			default		interpret				v		v			Interpret		Interpret
NDHGF1IN	0	1	0	1																																										
NDHGF1DF	0	0	1	1																																										
	v		v																																											
Meaning:	Use		Do not																																											
	default		interpret																																											
		v		v																																										
		Interpret		Interpret																																										

Data Set Header

Figure 2-12 (Page 5 of 10). Fields in the General Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHGRCFM	35	1	binary	0	none	<p>Specifies the record format of the data set.</p> <p>11.. undefined format 10.. fixed format 01.. variable format 00. no carriage control 10. ASA control characters 01. machine control characters</p> <p>JES2 uses the RECFM in the DCB specified in the job's JCL to set the field. The DCB defines other settings that are not defined by this field. However, JES2 allows all flag settings and transmits them across the network.</p> <p>JES3 sets the carriage control to ASA, machine, or none (bits 5 and 6 off). Undefined format is set (bits 0 and 1 on). Only carriage control bits are used on input.</p> <p>RSCS sets this field to X'42' for print files, and X'80' for punch files.</p> <p>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 a 3211 file. (For further details, see the following field, NDHGLREC.)</p> <p>For punch files, this field is used 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.</p> <p>This field is modified for certain types of store-and-forward data when RSCS acts as an intermediate node.</p> <p>POWER does not use this field, but sets it to X'Cx' (undefined). At store-and-forward nodes, files having ASA carriage control will not have the CC converted to machine carriage control. POWER supports all values of the field.</p>

Figure 2-12 (Page 6 of 10). Fields in the General Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description								
NDHGLREC	36	2	binary	0	1-32760	<p>The maximum logical record length (LRECL) of any record that appears within the data set. It includes the carriage control character if specified for the data set.</p> <p>JES2 obtains the LRECL from the DCB. It includes the carriage control character.</p> <p>JES3 uses the field if it is not 0. The following shows how JES3 defaults the field according to the type of data set and the presence of carriage control:</p> <pre>Print with CC 133 (132 without CC) Punch with CC 81 (80 without CC)</pre> <p>RSCS uses this field as follows:</p> <p>For incoming print files, this determines whether they are defined as 1403, 3211, or 3800 files.</p> <table> <tr> <td>Record Length</td> <td>Device</td> </tr> <tr> <td>1 - 132</td> <td>1403</td> </tr> <tr> <td>133 - 150</td> <td>3211</td> </tr> <tr> <td>151 - 32760</td> <td>3800</td> </tr> </table> <p>punch records greater than 80 (greater than 81 with carriage control) are truncated at the destination node. In setting the field, RSCS obtains the LRECL from TAGRECLN (with one added to each record to account for the CCTL byte) as follows:</p> <ol style="list-style-type: none"> 1. This field is set to 80 for punch files. 2. Virtual 3800 files that contain any load CCWs have this field set to 8192 (regardless of the actual maximum record length in the file). <p>RSCS does not modify this field as an intermediate node.</p> <p>POWER uses the field and sets it to the maximum length of any record in the data set, up to a maximum record length of 32K-1.</p>	Record Length	Device	1 - 132	1403	133 - 150	3211	151 - 32760	3800
Record Length	Device													
1 - 132	1403													
133 - 150	3211													
151 - 32760	3800													
NDHGDSCT	38	1	binary	0	0-255	<p>The data set copy count. This field defines the number of copies to print or punch at the destination node and 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.</p> <p>JES2 sets the data set copy count to the number of copies for the data set.</p> <p>JES3 sets and uses the data set copy count. If the value in the field is 0, one copy of the data set is produced.</p> <p>RSCS sets the field from the copy count specified on the SPOOL command. RSCS uses the field to set the spool copy count.</p> <p>POWER sets the field with the information supplied from the COPY parameter of the '* \$\$ LST' or '* \$\$ PUN' statement. It is used as the copy count (defaulted to 1).</p>								

Data Set Header

Figure 2-12 (Page 7 of 10). Fields in the General Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHGFCBI	39	1	binary	0	-31 to 31	<p>Used as an index byte when loading the FCB on a 3211. It allows the user to print each line with the data shifted right or left by (up to) 31 characters. A negative value indicates left indexing; a positive value indicates right indexing. Right indexing adds leading blanks; left indexing removes data characters.</p> <p>JES2 supports both right and left indexing.</p> <p>JES3 and POWER do not use this byte but set the byte to the default value.</p> <p>RSCS does not use the field. Set from the INDEX keyword on the TAG statement. (Only positive values may be specified.)</p>
NDHGLNCT	3A	1	binary	0	none	<p>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.</p> <p>JES2 obtains the value from the job's JCL. The special values X'00' and X'FF' are used as follows: X'00' (default): Use the default value of the printing subsystem. X'FF': Do not force any page ejects. Let the skipping be solely determined by the carriage control (if present) in the SYSOUT data.</p> <p>JES3, RSCS, and POWER do not use this field, but set it to the default value.</p>
	3B	1				Reserved for IBM's use.
NDHGFORM	3C	8	char	blanks	N/A	<p>The name of the form to use when printing or punching the data set.</p> <p>JES2 and JES3 use only the first four bytes, followed by blanks. Forms are set to zero if not specified. If received, either zero or blanks indicate the standard (STD) forms are desired.</p> <p>RSCS sets the field with the value supplied by CP FORM unless that value is overridden by the value supplied by TAG. The default is only set when the CP SPOOL command specifies NULL for the FORM name. This field is used for the CP FORM when not defaulted. When the header field contains the default (blanks), the CP FORM name is set to the installation's default.</p> <p>POWER uses only the first four bytes because POWER has only four-character forms names. For a 3800 SYSOUT data set, a value of X'00' indicates that hardware defaults are to be used. A value of X'40' indicates that software defaults are to be used. When the job originated on a non-POWER system (POWER section was not present), X'00' is used as the software default.</p>

Figure 2-12 (Page 8 of 10). Fields in the General Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHGFCB	44	8	char	blanks	N/A	<p>The name of the FCB to be used if this data set is printed.</p> <p>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:</p> <pre>'FCB1' 1403 'FCB2' 3211 and 3203-5 'FCB3' 3800 'FCB4' 4245 and 4248</pre> <p>JES3 sets and uses the field, but the default is zero.</p> <p>RSCS this field is used for FCB on spool file. It is set from FCB on the SPOOL command. Only the first four characters are used. The field may be overridden by the CP TAG command (up to eight characters).</p> <p>POWER uses the field as the FCB name. The field is set from the FCB name if one was specified. The last four characters may be set to '\$\$\$\$' for device-independence. For a 3800 data set, a value of X'00' indicates that hardware defaults are to be used; a value of X'40' indicates that software defaults are to be used. When the job originates from a non-POWER system (POWER section was not present), X'00' is used as the software default.</p>
NDHGUCS	4C	8	char	blanks	N/A	<p>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".</p> <p>JES2 and JES3 use only the first four bytes, followed by blanks. UCS is set to zero if not specified. If received, either zero or blanks indicate that the default UCS (****) is desired. NDHATAB1 is used if UCS is specified as "*****", blanks, or zeros. JES2 prefixes the first four bytes with:</p> <pre>'UCS1' 1403 'UCS2' 3211 'UCS3' 3203-5 'UCS5' 4245 'UCS6' 3262-5 and 4248</pre> <p>RSCS does not use the field. RSCS sets the field with the value specified on the UCS keyword in the TAG.</p> <p>POWER sets the field with the UCS name specified on the * \$\$ LST statement.</p>

Data Set Header

Figure 2-12 (Page 9 of 10). Fields in the General Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHGXWTR	54	8	char	blanks	N/A	<p>The name of the external writer to process this data set. If equal to the NDHGRMT field, this field is also set and used by the TSO/E TRANSMIT and RECEIVE commands and the CMS SENDFILE command to allow files to be sent to individual MVS users.</p> <p>JES2 and JES3 set the field to the name of the external writer. A default of zero is set; either zero or blanks are accepted as defaults for input.</p> <p>RSCS does not use this field but sets it to the value specified of the EXTWTR keyword if specified on TAG statement.</p> <p>POWER does not use the field but sets it to the external writer name as specified on the DEST parameter.</p>
NDHGNAME	5C	8	char	blanks	A-Z,0-9	<p>This is the data set name qualifier associated with the data set at the creation node. It may be user specified or generated by the system.</p> <p>JES2 and JES3 set from the DSNAME = keyword on the DD JCL statement, or default it to a single question mark '?'. It is used to construct a resource name associated with this data set for purposes of resource access authorization.</p> <p>RSCS does not use the field and sets it to the default.</p> <p>POWER does not use the field and sets it to the default.</p>
NDHGFLG2 Bit NDHGF2PR	64	1 .1	bits 80 bit	0 0	none none	<p>Flag byte containing the following flags: If on, the data set is to be printed.</p> <p>JES2 does not use the bit, but sets it if the output class is a print class at the output execution node.</p> <p>JES3 uses this bit to denote print output. If there is print output, JES3 sets this bit.</p> <p>RSCS uses the bit to cause the data set to be spooled to a printer. RSCS sets the bit if the input file is a print file.</p> <p>POWER uses the bit to denote print output. It is set if the input file is a print file.</p>
Bit NDHGF2PU		.1	40 bit	0	none	<p>If on, the data set is to be punched.</p> <p>JES2 does not use the bit, but sets it if the output class is a punch class at the output execution node.</p> <p>JES3 uses this bit to denote punch output. If there is punch output, JES3 sets this bit.</p> <p>RSCS uses the bit to cause output to be spooled to a punch device. RSCS sets the bit if the input file is a punch file.</p> <p>POWER uses the bit to denote punch output. If there is punch output, POWER sets this bit.</p>
		6	3F bits			Reserved for IBM's use.

Figure 2-12 (Page 10 of 10). Fields in the General Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHGUCSO	65	1	bits	0	none	The UCS option byte. JES2, JES3, and RSCS do not use this field but set it to the default value. POWER uses the field and sets it if UCS is specified.
Bit NDHGUCSD		.1	80 bit	0	none	If on, indicates the UCS will be loaded with the block data check option. JES2, JES3, and RSCS do not use this bit but set it to the default value. (JES2 always blocks data checks.)
Bit NDHGUCSF		.1	40 bit	0	none	If the bit is on, the UCS will be loaded with the FOLD option. JES2, JES3, and RSCS do not use this bit but set it to the default value. 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.
		6	3F bits			Reserved for IBM's use.
	66	2				Reserved for IBM's use.
NDHGPMDE	68	8	char	blanks	none	The process mode for the data set specified by the user. The process mode does not necessarily imply the presence of a particular type of record; it only indicates a request by the user for the preferred type of output processing. This field is used for SYSOUT selection and scheduling. The IBM-defined process modes are: LINE PAGE SOS!1 SOS!2 JES2 and JES3 set the field from the PRMODE keyword on the OUTPUT JCL statement. If PRMODE is not specified, then the data itself is examined to determine the process mode. RSCS moves this field into an XAB and sets it from an XAB (if existing) for VM/PSF POWER does not use this field, but sets it to the default value.
	70	4	binary	0		Reserved for IBM's use.

3800 Section

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

Figure 2-13 (Page 1 of 5). Fields in the 3800 Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHALEN	0	2	binary	none	4-32764	The length of the 3800 section of the data set header.
NDHATYPE	2	1	binary	none	none	Defines the major type of this section. General sections are identified by a X'00'.
NDHAMOD	3	1	binary	0	none	Modifier for the section identifier. For the 3800 section of the data set header, there must be a modifier code of X'80'.
NDHAFLG1	4	1	byte	0	none	Flag byte containing the following flags:
Bit NDHAF1J		.1	80 bit	0	none	<p>If on, the data set contains table reference characters. The bit is set if OPTCD = J is specified.</p> <p>RSCS uses this bit as follows:</p> <ol style="list-style-type: none"> 1. If this flag is on, the incoming file is made a virtual 3800 print file. 2. For Version 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 stored and forwarded. Store-and-forward files have the TRC bytes re-inserted and the SELECTs removed when forwarded. The TRC bytes are only re-inserted when the file is forwarded to a non-VM/SP NJE system. 3. Version 2 typically does not strip off the TRC bytes. <p>This flag is set as follows:</p> <ol style="list-style-type: none"> 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 CAT CCWs in these files will be removed. 2. For all other files, the flag is turned on only if OPTCD = J is specified on the TAG command. <p>POWER uses this flag and sets it when a Spool Access Support (SAS) user indicates that the first character is a TRC.</p>
Bit NDHAF1BR		.1	40 bit	0	none	<p>If on, the 3800 burster is to be used.</p> <p>JES2 and JES3 use this bit and set it from the job's JCL.</p> <p>RSCS does not use the bit but sets it if BURST = Y is specified in the TAG.</p> <p>POWER uses the bit and sets it from information on the * \$\$ LST statement or from the default printer setup.</p>

Figure 2-13 (Page 2 of 5). Fields in the 3800 Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
Bit NDHAF1BN		.1	20 bit	0	none	<p>If on, the 3800 burster should not be used.</p> <p>JES2 does not use the bit, but sets it to the default value.</p> <p>JES3 sets the bit if STACKER = C is specified on the JES3 // *FORMAT statement to indicate continuous forms stacking.</p> <p>RSCS does not use the bit but sets it if BURST = N is specified on the TAG command.</p> <p>POWER uses the bit and sets it from the * \$\$ LST statement or from the default printer setup.</p>
Bit NDHAF1BN		.5	1F bits	0	none	Reserved for IBM's use.
NDHAFLCT	5	1	binary	0	0-255	<p>The 3800 flash count defining the number of copies of this data set to be flashed. If this field is not specified, but NDHGFLSH was specified, then all copies are flashed.</p> <p>JES2 and JES3 use and set this field from the job's JCL.</p> <p>RSCS sets the field using the CP SPOOL command or it may be overridden by the CP TAG command. RSCS uses the flash count with the spool file.</p> <p>POWER uses the field and sets it from information on the * \$\$ LST statement or from the default printer setup.</p>
NDHATREF	6	1	binary	0	0-3	<p>The table reference character that specifies which of the four translate table entries to use when printing the copy modification.</p> <p>JES2 and JES3 obtain the information from the job's JCL.</p> <p>RSCS uses this field and sets it using the CP SPOOL command or it may be overridden by the MODTRC keyword on the TAG command.</p> <p>POWER uses the field and sets it from information on the * \$\$ LST statement or from the default printer setup.</p>
	7	1				Reserved for IBM's use

Figure 2-13 (Page 3 of 5). Fields in the 3800 Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHATAB1	8	8	char	blanks	N/A	<p>The name of the first translate table. A translate table defines an index to a font that is defined for a 3800 printer. This 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 OPCTD=J is specified. Only the first four characters are used and set.</p> <p>JES2 obtains the name of the translate table from the job's JCL or JECL. If received, zeros or blanks indicate that JES2 should default the table name to zero. For NDHATAB1 only, if the field is set to ****, blanks, or zeros, JES2 uses the name specified in NDHGUCS.</p> <p>JES3 defaults the field to zeros.</p> <p>RSCS uses this field and sets it using the CHARS operand on the CP SPOOL command. It may be overridden using the CHARS keyword on the TAG command.</p> <p>POWER uses the field and sets it from the * \$\$ LST statement. For a 3800 SYSOUT data set, a value of X'00' indicates the hardware defaults should be used. A value of X'40' indicates the software defaults should be used.</p> <p>When the job originates on a non-POWER system (POWER section is not present), X'00' is used as a software default.</p>
NDHATAB2	10	8	char	blanks	none	<p>The name of the second translate table. Only the first four characters are used and set.</p> <p>JES2, JES3, RSCS, and POWER: See the individual product descriptions in the NDHATAB1 field.</p>
NDHATAB3	18	8	char	blanks	none	<p>The name of the third translate table. Only the first four characters are used and set.</p> <p>JES2, JES3, RSCS, and POWER: See the individual product descriptions in the NDHATAB1 field.</p>
NDHATAB4	20	8	char	blanks	none	<p>The name of the fourth translate table. Only the first four characters are used and set.</p> <p>JES2, JES3, RSCS, and POWER: See the individual product descriptions in the NDHATAB1 field.</p>

Figure 2-13 (Page 4 of 5). Fields in the 3800 Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHAFLSH	28	8	char	blanks	none	<p>The flash cartridge identifier used to produce pre-printed forms. This overlay will be printed on every page before the data is placed into position. Only the first four bytes are used.</p> <p>JES2 and JES3 use the field and set it from the job's JCL or JECL. If the user does not specify a value when submitting the job, it is set to zero.</p> <p>A value of NONE in the first four bytes indicates that no flashing is to be performed, regardless of device or installation default.</p> <p>RSCS uses the field and sets it from the CP flash name. It may be overridden by the FLASH keyword on the TAG command.</p> <p>POWER uses the field and sets it from the * \$\$ LST statement. For a 3800 SYSOUT data set, a value of X'00' denotes the use of hardware defaults; a value of X'40' denotes the use of software defaults. When a job originates on a non-POWER system (POWER section is not present), X'00' is used as the software default.</p>
NDHAMODF	30	8	char	blanks	none	<p>The copy modification id. It contains the name of a module placed on every page of the output data set when it is being printed. Only the first four bytes are used.</p> <p>JES2 uses the field and sets it from information in the job's JCL or JECL. If the user does not specify a value when the job is submitted, the copy modification id is defaulted to 0.</p> <p>JES3 uses and sets the copy modification id as specified in the job's JCL or JECL. The last four bytes are zeros. (If copy mod is not specified, the entire field consists of zeros.)</p> <p>RSCS uses and sets the field from the CP MOD name. It may be overridden by the MODIFY keyword on the TAG command.</p> <p>POWER uses and sets this field from the * \$\$ LST statement. In the VSE/SP 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. If the user does this, POWER defaults to the first CHAR. For a 3800 SYSOUT data set, a value of X'00' denotes the use of hardware defaults; a value of X'40' denotes the use of software defaults.</p>

Data Set Header

Figure 2-13 (Page 5 of 5). Fields in the 3800 Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHACPYG	38	8*1	binary	0	0-255	<p>identifies the copy groups. It is actually eight one-byte fields. The fields define the number of times each page of the data set will be copied when the data set is sent to the 3800 printer.</p> <p>The first field refers to the first transmission, the second to the second transmission, etc. In this case, NDHGDSCCT is not used. The number of transmissions is determined by the number of copy groups defined. The sum of all copy groups, or a single copy group, cannot exceed 255.</p> <p>RSCS uses the sum of all bytes as the CP copy count and turns on the CP copy group flag. 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 command.</p> <p>POWER uses the field and sets it from information specified on the * \$\$ LST statement or by a program.</p>

RSCS Section

Figure 2-14 describes the format of the RSCS section of the data set header.

Figure 2-14. Fields in the RSCS Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHVLEN	0	2	binary	none	4-32764	Length of the RSCS section of the data set header.
NDHVTYPE	2	1	binary	none	none	The major type of this section. The RSCS section of the data set header is identified by X'87'.
NDHVMOD	3	1	binary	0	none	Modifier for the major type identifier. A modifier code of X'00' is the modifier for the RSCS section.
NDHVFLG1	4	1	byte			Reserved for IBM's use
NDHVCLAS	5	1	char	none	A-Z 0-9	The CP spool file class.
NDHVDEV	6	1	binary	none	none	The originating CP device type. (See the CP DEVTYPES macro for valid values.)
NDHVPGLE	7	1	binary	none	none	The 3800 virtual page length.
NDHVDIST	8	8	char	blanks	none	The CP distribution code.
NDHVFNAM	10	12	char	blanks	none	The CP file name.
NDHVFTYP	1C	12	char	blanks	none	The CP file type.
NDHVPRIO	28	2	binary	50	0-99	The RSCS transmission priority as specified on the TAG command.
NDHVRSN	2A	1	binary	none	none	The version number of the RSCS system that created the header.
NDHVRELN	2B	1	binary	none	none	The release number of the RSCS system that created the header.
NDHVTAGR	2C	136	char	none	none	The TAG record as specified on the TAG command.

POWER Section

Figure 2-15 describes the format of the POWER section of the data set header.

<i>Figure 2-15. Fields in the POWER Section of a Data Set Header</i>						
Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHPLEN	0	2	binary	none	4-32764	The length of the POWER section of the data set header.
NDHPTYPE	2	1	binary	none	none	The major type of this section. The POWER section of the data set header is identified by a X'86'.
NDHPMOD	3	1	binary	0	none	Modifier for the major type identifier. A modifier code of X'00' indicates the POWER section of the data set header.
NDHPFLG1	4	1	byte			Reserved for IBM's use.
NDHPIDEV	5	1	binary	0	none	The POWER device type.
NDHPPRIO	6	1	char	none	none	The output priority. This field is defaulted to the priority specified in the POWER generation.
NDHPDISP	7	1	char	D	none	The output disposition. Valid values are D (delete after printing), K (keep), H (hold), and L (leave).
NDHPUSER	8	16	char	blanks	none	This field contains any specified user information. It is used for printing on a separator page.
NDHPJSBF	18	1	binary	0	none	The job suffix number, which is created for segmented output. It is similar to the spin data set number.
NDHPSYID	19	1	char	zero	none	The system qualifier for a shared spool configuration. It is used for printing on a separator page.
NDHPNSEP	1A	1	binary	0	none	The number of separator pages that should be printed for this output.
NDHPOPTN	1B	1	bits	0	none	The COPYSEP parameter specified for the output.
NDHPPART	1C	2	char	blank	none	The partition identifier where the job executed. It is used for printing on a separator page.
	1E	2				Reserved for IBM's use
NDHPRCFM	20	1	bits	0	none	The record format (such as SCS, BMS, or 3270).
	21	1	byte			Reserved for IBM's use.
NDHPJNUM	22	2	binary	0	none	The job number.
NDHPCOMP	24	4	char	blank	none	The compaction table name used for output destined to RJE/SNA.
NDHPPASS	28	8	char	blank	none	The password that has been specified for this output.
NDHPSETP	30	68	char	blank	none	The default SETPRT parameter list.
NDHPSTRT	74	8	char	blank	none	The time (in STCK format) when output spooling started.

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 is not required if all SYSIN data is RECFM F and LRECL 80. When present, it is sent as the only section in the data set header, without the "basic" general section, having a modifier of X'00'.

JES3 does not use, send, or store-and-forward these sections.

Figure 2-16. Fields in the Record Characteristics Change Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHCLEN	0	2	binary	none	4-32764	Length of the record characteristics change section of the data set header.
NDHCTYPE	2	1	binary	none	none	The major type of this section. The record characteristics change section of the data set header is identified by an identifier field of X'00'.
NDHCMOD	3	1	binary	0	none	Modifier for the major type identifier. A modifier code of X'40' indicates this is the record characteristics change section of the data set header.
NDHCFLG1	4	1	byte	0	none	Reserved for IBM's use.
NDHCRCFM	5	1	binary	0	none	The record format of the SYSIN data set. The bit settings and associated record formats are: 11.. undefined format 10.. fixed format 01.. variable format RSCS does not create this section of the header. It uses this field at a store-and-forward node to determine if a SYSIN file should be stored as coded NOPs (done whenever NDHCRCFM indicates that the file is not in fixed format). POWER always has an undefined record format.
NDHCLREC	6	2	binary	0	0-252	The maximum record length of records in the SYSIN data set that follows. RSCS uses this field at a store-and-forward node to determine if a SYSIN file should be stored as coded NOPs (done whenever NDHCLREC indicates that the record length is not 80). POWER allows a maximum value of 128.

Output Processing Section

The Output Processing Section is used for transmitting output processing parameters and is identified by an identifier field of X'89' and a modifier field of X'00'. This section's length is variable and depends on the type and number of output processing parameters being transmitted.

The individual output processing parameters are represented by self defining data structures contained within the Output Processing Text Block (OPTB) which follows a fixed length portion of the section. The parameter representations are termed Output Processing Text Units (OPTUs or TUs) and their structure is defined in "OPTB Structure" on page 2-40. The set of NJE defined OPTUs are defined in "OPTU Key Definitions" on page 2-42.

The 3800 section of the data set header is still built if the user specifies any 3800 attributes. Some fields are duplicated in the OPTB. However, if any of the 3800 fields are changed by operator command, those changes need only be reflected in the 3800 section, not in the OPTB. The 3800 section fields override the OPTB at the SYSOUT destination.

Figure 2-17 describes the format of the output processing section of the data set header. The output processing section includes several fields that describe the *length* of parts of the output processing section. All length fields include the length of their own field. For example, field NDHSLEN is the length of the entire output processing section including the length of field NDHSLEN.

Note: JES2 and JES3 call this section the "data stream" section in their control block references.

Figure 2-17 (Page 1 of 2). Fields in the Output Processing Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHS	0	0	label	none	none	The beginning of the output processing section.
NDHSLEN	0	2	binary	none	none	The length of the entire output processing section. It includes the two bytes used to hold the length.
NDHSTYPE	2	1	binary	X'89'	none	The identifier for an output processing section.
NDHSMOD	3	1	binary	X'00'	none	The modifier for the output processing section, which must be X'00' to indicate OUTPUT OPTBs. The values from X'01' to X'FF' are reserved.
NDHSFLEN	4	2	binary	none	none	The length of the fixed area of the section, down through the field NDHSGPID, but not including NDHSOPTB or any of the OPTBs. It includes these two bytes used to hold the length.
NDHSFLG1 Bit NDHSCPDS	6	1 .1	bits 80 bit	X'00' 0	none n/a	Flag byte. If set, this flag indicates the presence of at least one CPDS data stream record in the data set. Note: The presence of "stream mode" data (SRCB = X'B0') and "CPDS" data are semantically identical in the NJE protocols. JES3 sets, but does not use this flag.
	7	1				Reserved for IBM's use.
	8	8				Reserved for IBM's use.

Figure 2-17 (Page 2 of 2). Fields in the Output Processing Section of a Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHSNSTR	10	4	binary	none	none	The number of "begin page" structured fields in the data set (for example, the number of data records that begin with the sequence X'D3A8AF'). This field is not adjusted for multiple copies. RSCS does not set or use this field.
NDHSGPID	14	8	char	blanks	none	The 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 sets the field from the user's output JCL or generates it at job execution time. JES3, RSCS, and POWER do not use this field and set it to the default value.
NDHSOPTB	1C	0	mixed	none	none	The start of the variable area containing Output Processing Text Units (OPTUs). The variable area consists of a prefix area followed by the OPTUs.
NDHSPRID	1C	4	char	'SJPF'	none	A required constant that identifies the OPTB prefix. RSCS and POWER do not use this field.
NDHSVERS	20	1	binary	X'02'	none	A required constant that identifies the version of the prefix. This constant should be set, but may be ignored by non-MVS systems. The settings indicate the following MVS systems: 01 = MVS/SP 1.3.1 (FMID JBB1327) 02 = MVS/SP 1.3.3 (FMID JBB1329) and later
NDHSPLEN	21	1	binary	X'1C'	none	The prefix length, which is used to point to the beginning of variable-length OPTUs (start of OPTBs = NDHSOPTB + NDHSPLEN). This field is fixed at X'1C' for the NDHSVERS = 02 version.
NDHSDLEN	22	2	binary	X'0000'	none	The length of the data (all OPTUs) following the prefix.
NDHSVERB	24	8	char	'OUT-PUT'	none	A constant required for compatibility with MVS systems. This constant identifies the JCL statement normally used to specify output processing parameters in the MVS environment.
	2C	8	char	blanks	none	Reserved for IBM's use.
NDHSFLG2	34	1	bits	X'80'	none	Flag byte reserved for IBM's use.
	35	1				Reserved for IBM's use.
	36					The start of the OPTUs. (The previous field is the end of the OPTB prefix.)

OPTB Structure

The OPTB is structured as a fixed area followed by a sequence of Output Processing Text Units (OPTUs). The number of OPTUs (and their sequencing) is arbitrary; hence, the appearance of any specific OPTU is optional. When OPTUs do not appear that would have supplied values needed at the destination to process the data, installation-defined or product-defined defaults may be used.

OPTB Structure Definition

The structure of an OPTB can be defined in terms of a set of rules for building it. 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 at the right (:= definition). The sequence of elements allowed is indicated by explicit sequences of bracketed items. Alternative sequences are delimited by the "|" symbol, which means "OR".

```

<DSHOPT> := Output processing section of the data set header
           := <FIXED> <OPTB>
  <FIXED> := Fixed portion of the output processing section
           starting with NDHSLEN through NDHSGPID
<OPTB>   := <PREFIX> <OPTUS>
<PREFIX> := Fixed portion of the OPTB starting with NDHSPRID
           through the RESERVED field at offset X'36'
<OPTUS>  := <OPTU> | <OPTUS> <OPTU>
<OPTU>   := <KEY> <COUNT> <PARMS>
  <KEY>   := Two-byte key defined in keyword definition table
  <COUNT> := Two-byte count of number of parms (1-255)
  <PARMS> := <PARM> | <PARM> <PARM>
<PARM>   := <LL> <DATA>
  <LL>    := Two-byte data length (1-128)
  <DATA>  := Sequence of exactly <LL> bytes

```

OPTU Structure Definition

OPTUs are identified by a 2-byte key that is registered and unique within NJE. Each OPTU 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 OPTU is specified in a 2-byte count field that follows the key and precedes the parameter list. The following 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 IBM's use and 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. For a count of 0, no data elements should follow the count field.

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

Note: For compatibility with MVS systems, the parameter length is 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

parameter value associated with the key.

OPTU Key Definitions

The suggested defaults listed in the following table may be used when the OPTU is not contained in the OPTB.

Figure 2-18 (Page 1 of 5). OPTU Key Definition Table

Key	Maximum Count	Length Range	External Keyword	Data Values Allowed and Description
0001	1	1	Reserved	Reserved - private JES2/JES3.
0002	4	4	Reserved	Reserved - private JES2/JES3.
0003	1	2	CKPTLINE	Integer: range 0-32767. 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. (Not supported by VSE or VM.)
0004	1	2	CKPTPAGE	Integer: range 1-32767. 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 an SNA work station. (Not supported by VSE.)
0005	1	2	CKPTSEC	Integer: range 1-32767. Number of seconds that may elapse between printer checkpoints. Installation defaults may be used. (Not supported by VSE or VM.)
0006	1	1	Reserved	Reserved - private JES2/JES3.
0007	1	8	COMPACT	Symbolic name from 1-8 alphameric characters long, used to determine the compaction table when sending the SYSOUT data set described by this control statement to an SNA remote terminal. This specification overrides any remote device-defined compaction table. Installation defaults may be used. (Not supported by VSE or VM.)
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. Installation defaults may be used. (Not supported by VSE or VM.)
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. Indicates the data set indexing print position offset (to the right) for the 3211 printer. (Not supported by VSE or VM.)
0013	1	1	Reserved	Not used at output nodes.

Figure 2-18 (Page 2 of 5). OPTU Key Definition Table

Key	Maximum Count	Length Range	External Keyword	Data Values Allowed and Description
0014	1	1	LINDEX	Integer: range 1-31. Indicates the data set indexing print position offset (to the left) for the 3211 printer. (Not supported by VSE or VM.)
0015	1	1	Reserved	Reserved - private JES2/JES3.
0016	1	4	Reserved	Reserved - private JES2/JES3.
0017	1	1	Reserved	Reserved - private JES2/JES3.
0018	1	8	Reserved	Reserved - private JES2/JES3.
0019	1	1	PRTY	Integer: range 0-255. Represents the priority of the output data set for output queuing. (Not supported by VSE or VM.)
001A	1	1	Reserved	Reserved - private JES2/JES3.
001B	1	4	Reserved	Reserved - private JES2/JES3.
001C	1	8	WRITER	1-8 character name (consisting of any combination of alphanumerics, #, @, or \$) of an installation-written program in the system library that is to write the output data set. (Not supported by VSE or VM.)
001D	1	1-6	FORMDEF	1-6 character member name (consisting of any combination of alphanumerics, #, @, or \$) of the SYS1.IMAGELIB partitioned data set containing information used by the 3800-3 printer print a data set. The first two characters of the member name are pre-defined by installation conventions and are prefixed to the name specified here. The members can contain the following information: <ul style="list-style-type: none"> Overlays to be invoked during output processing. Location on the page where the overlays are placed. Suppression to be activated for specified page formats. (Not supported by VSE.)
001E	-	-	Reserved	Reserved for IBM's use.
001F	1	1-6	PAGEDEF	1-6 character name (consisting of any combination of alphanumerics, #, @, or \$) of a member in SYS1.IMAGELIB containing the information used by the 3800-3 printer to print a data set. The first two characters of the member name are pre-defined by installation conventions and are prefixed to the name specified here. The members can contain the following information: <ul style="list-style-type: none"> Logical page size and width. Desired fonts and page segments. Definition of multiple page types or formats. Definition of lines within a page (for example, line origin, carriage controls, spacing, fonts, and suppression control). Definition of multiple logical pages on a physical page and line-skipping instructions. (Not supported by VSE.)
0020	-	-	Reserved	Reserved for IBM's use.

Data Set Header

Figure 2-18 (Page 3 of 5). OPTU Key Definition Table

Key	Maximum Count	Length Range	External Keyword	Data Values Allowed and Description
0021	2	1-3	PIMSG MESSAGES	Mixed data format. Level of printing system-generated messages by a functional subsystem. See "PIMSG" on page 2-49 for details. (Not supported by VSE.)
0022	1	4	THRESHLD	Integer: range 1-9999999. 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 multiplied by the number of copies. (Not supported by JES2, VM, or VSE.)
0023	1	1	DPAGELBL	Binary: range X'40' - X'80' Specifies whether or not data pages are labeled with the security label associated with the data set. X'40' - (YES) Specifies that security labels are to be printed. X'80' - (NO) Specifies that security labels are not to be printed (the application may generate its own security labels). (Not supported by VM or VSE.)
0024	1	1	SYSAREA	Binary: range X'40' - X'80' Specifies whether a System Defined User Printable Area (UPA) is to be in effect when printing output. The System Defined UPA is a subset of the physical page and only Security Overlays (used for B1 security labeling) can print outside of it. X'40' - (YES) Specifies that the System Defined UPA is in effect for printed output. X'80' - (NO) Specifies that the System Defined UPA is not in effect for printed output.
0025 to 200A	-	-	Reserved	Reserved for IBM's use.
200B	2	1-41	FILEID	This option specifies where the file to be printed or associated resource files can be found. A print file can be either line or stream mode records. Stream mode data sets must have each record preceded by X'5A'. See "FILEID" on page 2-47 for syntax details. (Not supported by MVS or VSE.)
200C 2019	-	-	Reserved	Reserved for IBM's use.
201A	1	1	BIN	Integer: range 1-2. Specifies the media bin from which sheets are fed for printing as follows: 1 - indicates that all sheets are to be fed from the primary input source. 2 - indicates that all sheets are fed from an alternate input source. Overrides the value in the active FORMDEF. (Not supported by MVS or VSE.)

Figure 2-18 (Page 4 of 5). OPTU Key Definition Table

Key	Maximum Count	Length Range	External Keyword	Data Values Allowed and Description
201B	1	1	DUPLEX	<p>Binary: range X'00'-X'02'</p> <p>Specifies whether printing is to be done in simplex or duplex mode. This parameter overrides the value in the active FORMDEF and applies during the printing of the entire print data set.</p> <p>X'00' - (NODUPLEX) Specifies standard simplex printing; printing will occur on only one side of the physical page.</p> <p>X'01' - (DUPLEX) Specifies that printing is to occur on both sides of the physical page. This permits the print data set to be bound on its side.</p> <p>X'02' - (DUPLEX TUMBLE) Specifies that printing will occur on both sides of the physical page such that the top of the back side is oriented with the bottom of the front side. Tumble duplexing permits the print data set to be bound at the top.</p> <p>(Not supported by MVS or VSE.)</p>
201C	1	1	OPTIMIZE	<p>Binary: range X'00'-X'02'</p> <p>Specifies whether resource (page segments, fonts, etc.) loading is to be optimized to determine the best time to load resources into the printer. This will improve the performance of the loading and printer processes at the cost of more storage and the CPU used during the optimization analysis.</p> <p>X'00' - (NOOPTIMIZE) Indicates that resources are to be loaded into the printer on a page-by-page basis.</p> <p>X'01' - (OPTIMIZE) Specifies that all possible resources are to be loaded prior to the printing of the data set.</p> <p>X'02' = (OPTIMIZE FONTS) Specifies, that in addition to resource optimization, only those characters of a font that are referenced in the data set are loaded in the printer. This option is not defined for double-byte fonts.</p> <p>The optimization function is not defined for line-mode or mixed-mode data sets.</p> <p>(Not supported by MVS or VSE.)</p>
201D	1	1	VERIFY	<p>Binary: range X'00'-X'01'</p> <p>Specifies that the data stream contained in the data set be syntax checked at the destination.</p> <p>0 = NO, 1 = YES. For example, each CPDS structured fields are checked for fields containing values within ranges, reserved fields not used, mandatory values being correct, etc.</p> <p>(Not supported by MVS or VSE.)</p>
201E	4		TRACE	<p>Mixed Format.</p> <p>Allows for the tracing of the internal operation of the print server as a particular spool file is being processed. See "TRACE" on page 2-49 for additional details.</p> <p>(Not supported by MVS or VSE.)</p>

Data Set Header

Figure 2-18 (Page 5 of 5). OPTU Key Definition Table

Key	Maximum Count	Length Range	External Keyword	Data Values Allowed and Description
201F	32	2	SHEETS	<p>Integer: range 1-32767</p> <p>This option overrides a FORMDEF specification and will cause only the sheets specified to be printed. A sheet represents the physical page, not the logical page. The numbering begins with the first sheet (sheet 1) and may or may not have a relationship with the page number printed on the pages. The sheet numbers are specified in begin/end pairs. Each number and each range must be listed in ascending order from left to right.</p> <p>(Not supported by MVS or VSE.)</p>
2020	3	1-8	DUMP	<p>Mixed format.</p> <p>Allows for the DUMPing of the internal operation of the printing subsystem (e.g., Print Services Facility). See "DUMP" on page 2-47 for additional details.</p> <p>(Not supported by MVS or VSE.)</p>
2021	-	-	Reserved	Reserved for IBM's use.
2022	1	1	DATAACK	<p>Binary: range X'00' - X'82'</p> <p>This option specifies whether the printer will block print-positioning and invalid-character errors. The term "block" means the printer will not return error messages to PSF, even if data is lost.</p> <p>X'00' - (BLOCK) Specifies that the printer will block (not report) print-positioning and invalid-character errors.</p> <p>X'80' - (UNBLOCK) Specifies that the printer will not block (will report) both print-positioning and invalid-character errors.</p> <p>X'81' - (BLKCHAR) Specifies that the printer will block invalid-character errors. The printer will report print-positioning errors.</p> <p>X'82' - (BLKPOS) Specifies that the printer will block print-positioning errors. The printer will report invalid-character errors.</p> <p>(Not supported by VSE.)</p>
2023 to 3EFF	-	-	Reserved	Reserved for IBM's use.
3F00	1	1-128	CICSDATA	<p>Character: alphameric and national.</p> <p>A sequence of from 1 to 128 bytes passed by the CICS application that created the output request. This data is then associated with the output data set within the spool subsystems.</p> <p>Mapped by CICS. (Not supported by MVS or VM.)</p>
3F01 to FFFF	-	-	Reserved	Not used in NJE.

Notes:

1. Except where otherwise noted, the OPTU Keys are supported on VM, MVS, and VSE systems.
2. 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 defined herein.
3. Alphameric characters are A-Z, a-z, and 0-9.

Detailed OPTU Key Specifications

The tables in this section use captions that are explained by the following:

Offset	Offset into OPTU
Length	OPTU sub-field length in bytes
Type	Type of value for parameter 'N' = binary number 'C' = characters 'B' = Bit 'X' = Hexadecimal value
Values	Parameter values
Key	OPTU identifier
Count	Number of parameters
Parm Length	Length of parameter in bytes

DUMP: Allows for the DUMPing of the internal operation of the print server.

Figure 2-19. Print Server Dumping

Parameter	Offset	Length	Type	Values
KEY	0	2	N	X'2020'
COUNT	2	2	N	3
PARM LENGTH	4	2	N	0
REASON CODE	-	-	-	-
PARM LENGTH	6	2	N	7
MESSAGE ID	8	7	C	MMMxxxa
PARM LENGTH	15	2	N	1
MESSAGE COUNT	17	1	N	0-255

REASON CODE

A 4 byte hex reason code. This subparameter is currently not supported for NJE and the length for the subparameter is always set to zero indicating that the parameter is not specified.

MESSAGE ID

Identifies a particular message generated by the printing subsystem.

In VM systems, a PSF message identifier is represented by a 10-character field with the form CCCMMMxxxa. Message id is specified as MMMxxxa, dropping the three letter PSF component id, where 'MMM' is the PSF module id, 'xxx' is the PSF message number and 'a' is the action code of the message. Messages will be monitored for the issuance of the specified message id, and PSF will terminate the print data set with a DUMP if it is encountered.

MESSAGE COUNT

A binary count between 0 and 255 which identifies the occurrence of the message. If message count is specified as 12, then PSF will terminate the data set with a DUMP on the 12th occurrence of the specified message. This parameter is only applied to monitoring for messages. A message count of zero with no reason code indicates no dump.

FILEID: Specifies VM file identifiers for various files used during output processing.

Figure 2-20. File Identifiers used during Output Processing

Parameter	Offset	Length	Type	Values
KEY	0	2	N	X'200B'
COUNT	2	2	N	2
PARAM LENGTH	4	2	N	0
FILECODE	6	1	N	1-11
PARAM LENGTH	7	2	N	40
FILEID	9	36		(as follows)
FILENAME		8	C	CMS file name
FILETYPE		8	C	CMS file type
SEND CODE		1	N	0-2
USERID		8	C	CMS userid
DISK		3	C	000-FFF
PASSWORD		8	C	
EDISK		4	C	0000-FFFF

FILE CODE

- 01 = VM FORMDEF file
- 02 = VM PAGEDEF file
- 03 = VM CODED FONT file
- 04 = VM OVERLAY file
- 05 = VM PAGESEG file
- 06 = VM PRINT file, with CC (no TRC)
- 07 = VM CODE PAGE file
- 08 = VM FONT CHARACTER SET file
- 09 = VM PRINT file, with CC and TRC
- 10 = VM PRINT file (no CC, no TRC)
- 11 = VM PRINT file, with TRC (no CC)

SEND CODE

0 = SYSDISK

Read the specified file from the SFCM resource disk. When SYSDISK or SEND are specified, the userid, disk address and password do not need to be specified.

1 = SEND

Send the file to the SPOOL.

2 = LINK

File will be read by the SFCM from the specified disk. When LINK is specified, the userid, disk address and password must also be specified.

USERID

CMS userid that owns the file.

DISK

Disk address of the minidisk which contains the file.

PASSWORD

Password to the minidisk.

EDISK

Extended disk address of the minidisk which contains the file. (Optional.)

PIMSG: Level of printing system-generated messages by a functional subsystem.

Figure 2-21. Functional Subsystem Messages

Parameter	Offset	Length	Type	Values
KEY	0	2	N	X'0021'
COUNT	2	2	N	1-2
PARAM LENGTH	4	2	N	2
PIMSG/MESSAGES	6	1	X	X'80', X'40'
PARAM LENGTH	7	2	N	2
MESSAGE COUNT	9	2	N	0-999

PIMSG/MESSAGES**X'40'**

specifies that the system is to print the messages generated by a functional subsystem. This is the default specification.

X'80'

specifies that the system is not to print the messages generated by a functional subsystem.

MESSAGE COUNT

Number of messages that are to be printed. A zero value means that all messages should be printed.

TRACE: Allows for the tracing of the internal operation of the SFCM and PDM virtual machines as a particular spool file is being processed. More than one of the parameters: INPUT, OUTPUT, CALL, IOCMD, STATE, etc. may be specified. Use of this option is restricted to those users identified to the printing subsystem as privileged. Specification by non-privileged users will result in a warning message issued with the printed document.

Figure 2-22. Trace ids

Parameter	Offset	Length	Type	Values
KEY	0	2	N	X'201E'
COUNT	2	2	N	4
PARAM LENGTH	4	2	N	1
TRACE CODE	6	1	X	X'80', X'40'
PARAM LENGTH	7	2	N	7
MESSAGE ID	9	7	C	MMMxxxa
PARAM LENGTH	16	2	N	1
MESSAGE COUNT	18	1	N	1-255
PARAM LENGTH	19	2	N	3
STATE	21	2	B	(see the following page)

TRACE CODE

X'80' = NO

Specifies no tracing should be done.

X'40' = YES

Specifies tracing should be done

MESSAGE ID

Identifies a particular PSF message identifier as documented in the PSF messages. A PSF message identifier is represented by a 10-character field with the form 'CCMMMyxxa'. Message is specified as MMMxxxa, dropping the three letter PSF component id, where 'MMM' is the PSF module id, 'xxx' is the PSF message number and 'a' is the action code from the message. Messages will be monitored for the issuance of the specified message id and tracing started with the specified options if it is encountered. If specified as X'FFFFFFF', tracing is done for the entire SFCM and PDM operation.

MESSAGE COUNT

A 1 to 255 binary count which identifies the occurrence of the message. If message count is specified as 12, PSF will start tracing on the 12th occurrence of the specified message. This parameter is only applied to monitoring for messages.

STATE

(bits numbered left to right)

Bit 0 = 1 (INTOKEN) - Each CPDS structured field will be recorded to the TRACE output file before being processed (SFCM parameter).

Bit 1 = 1 (INREAD) - Each record of CPDS or line data read from disk, spool, or created by the SPM will be recorded. In the PDM, each record of CCWs or IPDS read from disk will be returned (SFCM, PDM).

Bit 2 = 1 (OUTTOKEN) - Each CCW or IPDS control created will be recorded (SFCM).

Bit 3 = 1 (OUTWRITE) - Each record of IPDS or CCWs created will be recorded (SFCM).

Bit 4 = 1 (CALL) - The name of the caller and the return address will be recorded for each module invocation made (SFCM,PDM).

Bit 5 = 1 (IOCMD) - Each CCW or IPDS command will be recorded before being issued to the printer (PDM).

Bit 6 = 1 (STATE) - Records each state change in the SFC or SPM state machines (SFCM).

Bit 7 = 1 (ALLOCATE) - Records each new block of storage allocated and its allocation pointers (SFCM).

Bit 8 = 1 (FREE) - Records each freed block of storage and its allocation pointers (SFCM).

Bit 9 = 1 (PAGE) - Records each page record read or written (SFCM, PDM).

Bit 10 = 1 (FONTTAB) - Records contents of internal font table whenever changed. Done before each page is printed. Also, records device font table information (PDM).

Bit 11 = 1 (OVLITAB) - Records contents of internal overlay table whenever changed. Done before each page is printed (PDM).

Bit 12 = 1 (PSEGTAB) - Records contents of internal segment table whenever changed. Done before each page is printed (PDM).

Bit 13 = 1 (CHKPTTAB) - Records contents of internal checkpoint table whenever changed. Done after each page is printed, at the beginning and end of the document, and in error recovery situations. Also records device page number counter information (PDM).

Bit 14 = 1 (LIST) - May be specified along with the INTOKEN, INREAD, OUTTOKEN, OUTWRITE, ALLOCATE, FREE, PAGE, and/or IOCMD parameters. When specified with the ALLOCATE or FREE parameters, LIST will cause all blocks on the allocated or free lists, respectively, and their storage allocation pointers to be recorded to the TRACE output file. When specified with the INTOKEN, INREAD, OUTTOKEN, OUTWRITE, PAGE, or IOCMD parameters, LIST will record the entire input record.

Bit 15 = 1 (PPCC) - After each SNA LU 6.2 communications verb issued to the PPCC, the PPCC internal trace of module calls and returns for that PPCC verb will be recorded. These modules are not traced by the "CALL" keyword because PPCC is a separate component (3820 PDM).

Bit 16 = 1 (EMSGSNAP) - Records a single snap of the 3820 internal trace table, the SNA driver I/O control block and the 3820 global work area when message type E or S occurs (3820 PDM).

Bit 17 = 1 (3820TBLS) - Records the contents of the data set information table and the data set tracker table whenever they are modified (3820 PDM).

Bits 18-23 are reserved for future use (set to 0).

Error Handling

Destination nodes that do not recognize OPTUs received in an Output Processing Section may either reject the SYSOUT (X'BO' 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 NJHNTYN node and NJHNUSID in the job header. Destination nodes that process OPTUs must adhere to the following rules for error cases:

Invalid text block format

A system-dependent action may be taken, such as rejecting the file or substituting an entire set of default parameters. A message must be sent 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.

Note: Not all products support error handling.

Security Section of the Data Set Header

Figure 2-23 describes the format of the Security Section of the Data Set Header.

The section is identified by an identifier field of X'8C' and a modifier field of X'00'. It may be built as part of a created NJE work element (job or SYSOUT), or may be added to a forwarded work element.

The security section contains information related to the identification of the owner or submitter user associated with the data set described by the corresponding data set header. The information in this section is to be treated as a token that is to be interpreted by the security facility component of the trusted computer base.

Figure 2-23 (Page 1 of 2). Fields in the Security Section of the Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHTLEN	0	2	binary	84	4-32764	Length of the data set security section.
NDHTTYPE	2	1	binary	X'8C'	none	The major type identifier.
NDHTMOD	3	1	binary	0	none	Modifies the major type identifier.
NDHTLENP	4	2	binary	4	4-32764	Length of the prefix sub-section.
NDHTFLG0 bit NDHTF0JB	6	1 .1	bits 80 bit	0 0	none none	This flag byte defines the following flags: Indicates that the security section represents the owner of the NJE job. The bit not set indicates that the security section represents the submitter of the NJE job.
RESERVED	7	1		0	none	Reserved for IBM's use.
	8					End of prefix sub-section
NDHTLENT	0	1	binary	80	80-255	Length of the security data sub-section.
NDHTVERS	1	1	binary	X'01'	1-255	Security section version.
NDHTFLG1 bit NDHT1EN bit NDHT1EXT	2	1 .1 .1	bits 80 bit 40 bit	0 0 1	none none none	This flag byte defines the following flags: Indicates that the security section (from NDHTSTYP on) is encrypted. Encryption algorithm is private to the security product. Indicates the format of the security section. B'1' indicates external format which is used exclusively in NJE.
NDHTSTYP	3	1	binary	X'07'	1-255	Security session type. '7' indicates batch job type.
NDHTFLG2 bit NDHT2DFT bit NDHT2MLO bit NDHT2SHI bit NDHT2TRS bit NDHT2SUS	4	1 .1 .1 .1 .1	bits 80 bit 20 bit 10 bit 08 bit 04 bit	1 0 0 0 0	none none none none none	This flag byte defines the following flags: Indicates that the security section data has not been verified by the security product at the node which created the security section. If not set, the security data may be assumed to have been verified. Indicates that all required ML options were active at the node which created the security section. Indicates that the security data could not be verified at the node creating the security section. If set, the user to which this section applies is to be considered as part of the trusted computer base. If set, the security data was obtained via a surrogate user at the node of creation.

Figure 2-23 (Page 2 of 2). Fields in the Security Section of the Data Set Header

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
bit NDHT2RMT		.1	02 bit	1	none	Indicates that the job or data set was originated remotely.
NDHTPOEX	5	1	binary	X'03'	1-255	Indicates the port of entry class index. The following port of entry class indices are defined: 1 = Terminal Class 2 = Console Class 3 = JESINPUT Class
RESERVED	6	2				Reserved for IBM's use.
NDHTSECL	8	8	char	none	A-Z,0-9	The security label associated with the user represented by this section.
NDHTCNOD	10	8	char	none	A-Z,0-9	Node at which security section was generated.
NDHTSUSR	18	8	char	none	A-Z,0-9	Userid of submitting user.
NDHTSNOD	20	8	char	none	A-Z,0-9	Node of job submission.
NDHTSGRP	28	8	char	none	A-Z,0-9	Groupid of submitting user.
NDHTPOEN	30	8	char	none	A-Z,0-9	Port of entry name at the creation node for the job creating this data set.
RESERVED	38	8				Reserved for IBM's use.
NDHTOUSR	40	8	char	none	A-Z,0-9	Security session owner userid.
NDHTOGRP	48	8	char	blanks	A-Z,0-9	Security session owner groupid.

Installation-Defined Section

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

Beyond the four-byte definition, this section is the responsibility of the installation. The section length is limited to 32764 bytes. The combined length of all sections in the header is limited to 32764 bytes, 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. The following depicts a sample eight-byte field:

Figure 2-24. Fields in the Data Set Header Extension Section

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NDHULEN	0	2	binary	none	4-32764	Length of the section.
NDHUTYPE	2	1	binary	none	C0-FF	The major type of this section.
NDHUMOD	3	1	binary	none	none	Modifier for the major type identifier.
NDHUCODE	4	4	char	none	none	This could be used for a GUIDE/SHARE installation code, or some other unique identifier.

Job Trailer

A job trailer indicates the end of an NJE job. The job trailer is divided into several sections depending on the networking facility where the NJE job originated. All networking facilities require a job trailer prefix section and a job trailer general section.

Figure 2-25 describes which sections in the trailer are built and used by each subsystem.

Figure 2-25. Job Trailer Sections used by Each Subsystem

Section	JES2	JES3	RSCS	POWER
General Section	built and used	built and used	built and used	built and used
Accounting Section	built and used	not built or used	not built or used	not built or used
Product-specific Sections	- none -	- none -	- none -	- none -

Job Trailer Prefix Section

The job trailer begins with the standard four-byte control record prefix containing the combined length of all sections. Figure 2-26 describes the format of the job trailer prefix section.

Figure 2-26. Fields in the Job Trailer Prefix Section

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJTLEN	0	2	binary	none	8-32K	Length of all sections in the job trailer.
NJTFLAGS	2	1	binary	0	none	Reserved for IBM's use.
NJTSEQ	3	1	binary	0	none	The sequence indicator, which contains valid information after the header is segmented for transmission. The information specifies: <ul style="list-style-type: none"> The position of the segment within the record. More record segments follow.

Job Trailer General Section

The job trailer general section is identified by an identifier field of X'00' and a modifier field of X'00'. Figure 2-27 describes the format of the job trailer general section.

Figure 2-27 (Page 1 of 3). Fields in the Job Trailer General Section

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJTGLEN	0	2	binary	none	4-32764	Length of the job header general section.
NJTGTYPE	2	1	binary	none	none	The type of section. This value must be X'00' for the job trailer general section.
NJTGMOD	3	1	binary	0	none	Modifier for the type of section. This field contains X'00' to identify this as the job trailer general section.
NJTGFLG1	4	1	bits	0	none	Reserved for IBM's use.

Figure 2-27 (Page 2 of 3). Fields in the Job Trailer General Section

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJTGXCLS	5	1	char	A	A-Z 0-9	The actual execution class of the job. JES3 sets this field to zero and does not use it. RSCS and POWER do not use the execution class.
	6	2				Reserved for IBM's use.
NJTGSTRT	8	8	binary	0	none	The time of day the job started execution in System/370 TOD clock format (GMT). JES2 sets the time of day and uses this field in SMF type 26 records. JES3 and RSCS do not use the field but set the field to the default value. POWER sets the field but does not use it. POWER sets the field to zero or the default value.
NJTGSTOP	10	8	binary	0	none	The time the job completed execution in System/370 TOD clock format (GMT). JES2 sets the field to zero or the default value for spin data sets. JES3 and RSCS do not use the field but set the field to the default value. POWER sets the field but does not use it. For spin output, all trailers except the last are zero.
	18	4				Reserved for IBM's use.
NJTGALIN	1C	4	binary	0	none	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 sets the field to zero or the default value if SYSOUT is sent before the job finishes executing. JES3 sets the default value for this field, but does not use it. RSCS sets the field to the number of records in the file (from TAGRECNM) if the file is a print file. RSCS does not use this field. POWER sets the field to the number of records produced by the job. For NJE jobs that contain a spin-off data set, all trailers except for the last are zero.
NJTGACRD	20	4	binary	0	none	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 sets the field to X'00' for NJE jobs that contain spin-off data sets. JES3 uses the default value for the field. JES3 does not use the field. RSCS does not use this field. RSCS sets the field to the number of records in the file (from TAGRECNM) if it is a punch file (SYSOUT) or SYSIN. POWER does not use this field. It is set to the number of records produced by the total job. For spin output, all trailers except the last are zero.
	24	4				Reserved for IBM's use.

Figure 2-27 (Page 3 of 3). Fields in the Job Trailer General Section

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJTGIXPR	28	1	binary	0	0-15	<p>The initial requested execution selection priority.</p> <p>JES2 uses this field but does not set it. The low-order 4 bits are shifted to the high-order four bits of JCTIPRIO by the SYSOUT receiver.</p> <p>JES3 defaults this field but does not use it.</p> <p>RSCS does not use this field. It is set to the RSCS transmission priority as shown in NJHGPRIO.</p> <p>POWER sets this field from the PRI = parameter on the * \$\$ JOB statement.</p>
NJTGAXPR	29	1	binary	0	0-15	<p>The actual execution selection priority used.</p> <p>JES2 sets this field, but does not use it.</p> <p>JES3 sets this field to the default, but does not use it.</p> <p>RSCS does not use this field. It is set to the RSCS transmission priority as shown in NJHGPRIO.</p> <p>POWER defaults this field, but does not use it.</p>
NJTGIOPR	2A	1	binary	0	0-255	<p>The initial job priority for output selection.</p> <p>JES2 sets the range (0-15) when the SYSOUT job is selected for transmission. (JES2 sets the field to 1 for spin data sets if the job is still executing.) JES2 does not use this field.</p> <p>JES3 uses the greater value of NJTGIOPR or NJTGAOPR for the job's priority for received output (unless greater than 15, in which case 15 is used). This field is set from the job's priority for outgoing SYSOUT.</p> <p>RSCS does not use this field. It is set to the RSCS transmission priority as shown in NJHGPRIO.</p> <p>POWER defaults this field but does not use it.</p>
NJTGAOPR	2B	1	binary	0	0-255	<p>The actual output selection priority used.</p> <p>JES2 defaults this field, but does not use it.</p> <p>JES3 uses the greater value of NJTGIOPR or NJTGAOPR for the job's priority for received output (unless greater than 15, in which case 15 is used). This field is set from the job's priority for outgoing SYSOUT.</p> <p>RSCS does not use this field. It is set to the RSCS transmission priority as shown in NJHGPRIO.</p> <p>POWER defaults the field but does not use it.</p>

Accounting Section

The following defines the accounting section. This section is identified by an identifier field of X'89' and a modifier field of X'00'.

Note: These statistics are for the entire job, not just for an individual SYSOUT data set.

Figure 2-28. Fields in the Job Trailer Accounting Section

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJTS	0	0	label	none	none	The start of the accounting section.
NJTSLLEN	0	2	binary	none	none	Length of the accounting section including the two bytes used to hold the length.
NJTSTYPE	2	1	binary	X'89'	none	The identifier for the accounting section.
NJTSMOD	3	1	binary	X'00'	none	The modifier for the accounting section, which must be '00'.
NJTSAPAG	4	4	binary	none	none	The actual number of the "begin page" structured fields. JES3, RSCS, and POWER do not set or use this field.
NJTSABYT	8	4	binary	none	none	The actual number of bytes. JES3, RSCS, and POWER do not set or use this field.

Installation-Defined Section

This section is identified by an identifier field of B'11xxxxxx' and a modifier field of the installation's 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 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. The following depicts a sample eight-byte field:

Figure 2-29. Fields in the Job Trailer User Section

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NJTULEN	0	2	binary	none	4-32764	Length of the section.
NJTUTYPE	2	1	binary	none	C0-FF	The major type of this section.
NJTUMOD	3	1	binary	none	none	Modifier for the major type identifier.
NJTUCODE	4	4	char	none	none	This could be used for a GUIDE/SHARE installation code, or some other unique identifier.

NJE Nodal Message Records (NMR)

A node in a network uses a nodal message record (NMR) to transmit commands and messages. Command responses are treated as messages in the nodal message record. The RCB associated with the NMR contains X'9A'.

The NMR record consists of a header containing control information followed by the message or command text. Some of the fields in the header are used differently depending on whether the record contains a message, formatted command, or an unformatted command.

The rules for sending messages and commands in a network are different than those for sending files. NMRs are not stored on spool at an intermediate node. If the NMR cannot be sent to the next node, it is discarded and only message responses are sent back to the user at the originating node. (The originating node may not receive any information about an NMR if disastrous errors occur at any of the network nodes during its transmission.)

When an NMR can be sent to a node, it is sent directly. After a node processes a command, the node generates and sends a message to the originating node to inform the person issuing the command that it has been processed.

Identifying the Contents of an NMR

To understand the contents of an NMR, you must first examine the contents of the NMRFLAGC and NMRTYPEF flags. The NMRFLAGC flag contains valid information only if the NMRTYPEC flag contains a 1. Use Figure 2-30 to determine the contents of an NMR.

NMRFLAGC	NMRTYPEF	NMR Contents:
-----	-----	-----
1	0	Unformatted command
1	1	Formatted command
0	N/A	Message

Figure 2-30. Identifying the Contents of an NMR

Interpreting the Contents of the NMROUT Field

The contents of the NMROUT field at offset X'0D' are determined by examining the following bits in NMRFLAG at offset X'00':

- NMRFLAGW -- 2 bits into NMRFLAG
- NMRFLAGT -- 3 bits into NMRFLAG
- NMRFLAGU -- 4 bits into NMRFLAG

Use Figure 2-31 to determine the contents of NMROUT.

NMRFLAGW	NMRFLAGT	NMRFLAGU	NMROUT Contents
0	0	0	MCS routing codes for the message.
0	0	1	The console identifier.
0	1	0	Userid.
1	0	0	The remote workstation identifier.

Figure 2-31. Interpreting the NMROUT Field

Format of an NMR that Contains an Unformatted Command

Unformatted commands are sent in NMRs and can be distinguished from other NMRs because the NMRFLAGC flag is set on and the NMRTYPEF flag is set off.

Figure 2-32 (Page 1 of 2). Fields in the Unformatted Command NMR

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NMRFLAG	0	1	byte	0	none	Flag byte
bit NMRFLAGC		.1	80 bit	1	none	If on, the NMR contains a command. See Figure 2-30 on page 2-58 to determine whether the command is unformatted or formatted.
bit NMRFLAGW		.1	40 bit	0	none	If on, the NMROUT field contains the remote number. Refer to Figure 2-31 on page 2-59 for information about using this bit to interpret the NMRMSG flag. JES3 sets and does not use the bit. RSCS uses but does not set the bit.
bit NMRFLAGT		.1	20 bit	0	none	POWER does not use or set the bit. If on, the NMROUT field contains the originating userid.
bit NMRFLAGU		.1	10 bit	0	none	JES3 and POWER do not use or set the bit. If on, NMROUT contains the originating console identifier. Refer to Figure 2-31 on page 2-59 to interpret the contents of the NMROUT field.
bit NMRFLAGR		.1	08 bit	0	none	POWER does not use or set this bit. JES3 sets the bit if it is transmitting an NMR that contains a command. It is not used when JES3 receives an NMR that contains a command.
bit NMRFLAGJ		.1	04 bit	0	none	RSCS does not use the bit. If on, the console is remote-authorized only.
bit NMRFLAGD		.1	02 bit	0	none	JES3 sets but does not use the bit. POWER and RSCS do not use or set this bit.
bit NMRFLAGG		.1	01 bit	0	none	If on, the console is not job-authorized.
bit NMRFLAGH		.1	01 bit	0	none	If on, the console is not device-authorized.
bit NMRFLAGI		.1	01 bit	0	none	If on, the console is not system-authorized.
NMRLEVEL	1	4 high-order bits	binary	0	none	Networking products do not use these bits for commands. The bit should be set to X'70'.

Figure 2-32 (Page 2 of 2). Fields in the Unformatted Command NMR

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NMRPRIO		4 low-order bits	binary	0	none	The output priority. Networking products do not use these bits for commands. The bit should contain X'07'.
NMRTYPE bit NMRTYPEX bit NMRTYPE4 bit NMRTYPEP bit NMRTYPEF bit NMRTYPED	2	1 .4 .1 .1 .1 .1	byte F0 bits 08 bit 04 bit 02 bit 01 bit	0 0 0 0 0 0	none none 0 0 0 none	Flag byte Reserved for IBM's use. Always 0 for unformatted commands. Always 0 for unformatted commands. Always 0 for unformatted commands. Reserved for IBM's use.
NMRML	3	1	binary	1	1-132	Length of the information in NMRMSG.
NMRTO NMRTONOD NMRTOQUL	4	9 8 1	mixed char binary	none blanks 0	none none none	Describes the destination node for the command with the following two fields. The destination node for the command. The system identifier for the destination node. It is used to identify the system that should process the command. JES3 sets but does not use the byte. RSCS and POWER sets it to the default value but does not use the byte.
NMRROUT	D	8	char	zeros	none	The originating userid, remote id, or console id depending on the setting of NMRFLAG. Refer to "Interpreting the Contents of the NMRROUT Field" on page 2-58 for additional information about interpreting where the NMR originated. JES3 sets the originating id to either the remote workstation or a JES3 console where the *SEND command was issued. Note: The format for consoles is the same as for messages: NMRDESC = 0000, NMRROUT = JES3 console number, NMRCNID = 00000000. RSCS initializes the field to the the originating userid. It always contains the originating id when either NMRFLAGT or NMRFLAGW is on.
NMRFM NMRFMNOD NMRFMQUL	15	9 8 1	mixed char binary	none blanks 0	none none none	The following two fields describe the originating node. The originating node. The system identifier of the originating node. JES3 sets the system identifier to the default and uses the field. RSCS and POWER sets it to the default value and uses the field.
NMRMSG	1E	132	mixed	none	none	The command text.

Format of an NMR that Contains a Formatted Command

Formatted commands are sent in NMRs and can be distinguished from other NMRs because flag NMRFLAGC and NMRTYPEF are both on.

JES2 is the only system that sends formatted commands. RSCS and JES3 support formatted commands for input only, and change them to equivalent RSCS or JES3 commands prior to processing. Therefore, in the following section, the RSCS and JES3 descriptions indicate only the usage of the commands.

POWER does not send out formatted commands. If POWER receives an NMR that contains a formatted command, the NMR is discarded and the user at the originating node is not notified that the command has been discarded.

Figure 2-33 (Page 1 of 3). Fields in the Formatted Command NMR

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NMRFLAG	0	1	byte	0	none	Is a flag byte
bit NMRFLAGC	0	.1	80 bit	1	none	If on, the NMR contains a command. See Figure 2-30 on page 2-58 to identify whether the command is formatted or unformatted.
bit NMRFLAGW	0	.1	40 bit	0	none	If on, the NMROUT field contains a JES2 remote number. JES3 does not use this bit.
bit NMRFLAGT	0	.1	20 bit	0	none	RSCS does not use or set this bit. If on, the NMROUT field contains a userid. JES3 does not use the bit.
bit NMRFLAGU	0	.1	10 bit	0	none	RSCS does not use the bit. If RSCS is a path node, the information is not passed with the NMR to the next node. If on, the NMROUT field contains the originating console identifier. Refer to Figure 2-31 on page 2-59 to interpret the contents of field NMROUT. JES3 does not use this bit.
bit NMRFLAGR		.1	08 bit	0	none	RSCS does not use this bit and does not pass the information to the next node in the path. If on, the console is remote-authorized only. JES3 does not use or set the bit.
bit NMRFLAGJ		.1	04 bit	0	none	RSCS does not use the bit. If on, the console is not job-authorized.
bit NMRFLAGD		.1	02 bit	0	none	If on, the console is not device-authorized.
bit NMRFLAGS		.1	01 bit	0	none	If on, the console is not system-authorized.
NMRLEVEL	1	4 high-order bits	binary	0	none	Does not apply to NMRs that contain commands.
NMRPRIO		4 low-order bits	binary	0	none	Does not apply to NMRs that contain commands.
NMRTYPE	2	1	bits		None	
bit NMRTYPEX		.4	F0 bits	0	none	Reserved for IBM's use.
bit NMRTYPE4		.1	08 bit	0	none	Always 0 for commands.
bit NMRTYPEP		.1	04 bit	0	none	Always 0 for commands.
bit NMRTYPEF		.1	02 bit	0	1	On for formatted commands.
bit NMRTYPEP		.1	01 bit	0	none	Not applicable to commands.

Figure 2-33 (Page 2 of 3). Fields in the Formatted Command NMR

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NMRML	3	1	binary	1	1-132	Length of the information in NMRMSG. The length includes NMRFM, even if it is not specified.
NMRTO NMRTONOD NMRTQUL	4	9 8 1	mixed char binary	none blanks 0	none none none	This field is comprised of the following two fields: The destination node for the command. The system identifier for the destination node. It identifies the system that should process the command. JES3 does not use or set this field. RSCS does not use or set this field.
NMRROUT	D	8	char	zeros	none	The originating userid, remote id, or console id depending on the setting of NMRFLAG. Refer to "Interpreting the Contents of the NMRROUT Field" on page 2-58 for additional information about interpreting where the NMR originated.
NMRFM NMRFMNOD NMRFMQUL	15	9 8 1	mixed char binary	none blanks 0	none none none	The following two fields describe the originating node. The originating node. The system identifier of the originating node.
NMRMSG	1E	variable	mixed	none	none	The formatted command.
NMRFOP		1	binary	none	none	The action the destination node should take with the NMR that contains the formatted command. The possible values are: 1. NMRFOPD - the NMR contains a command to display a job. 2. NMRFOPC - the NMR contains a command to cancel a job. 3. NMRFOPA - the NMR contains a command to release a job. 4. NMRFOPH - the NMR contains a command to place a job in operator hold. 5. NMRFOPR - the NMR contains a command to reroute the job.
NMRFFLG	1F	1	byte	0	none	Flags or opcode modifiers. JES2 initializes the field to zeros.
bit NMRFFLGO bit NMRFFLGD		.1 .1	80 bit 40 bit	0 0	none none	If for a cancel command, the output will be purged. If for a route command, output will be rerouted. This flag indicates that the node should cancel job execution and issue a dump. This field is mutually exclusive with NMRFFLGO when NMRFFLGO is on for cancel. JES3 does not use this bit. <u>Reserved for IBM's use.</u>
NMRFJID	20	2	binary	none	none	The job identifier of the job to be processed by the NMR. JES2 initializes the field to zeros. It is set only if specified by the \$G operator command. JES3 does not use the original job number field.
NMRFORGN	22	8	char	blanks	none	The originating node where the job was issued.
NMRFJNAM	2A	8	char	blanks	none	The job name of the job to be processed. JES3 uses the job name or number that the operator would use on an operator command. RSCS uses the first 4 characters as the spool identifier.

Figure 2-33 (Page 3 of 3). Fields in the Formatted Command NMR

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NMRFD	32	8	char	blanks	none	The destination for the route command.
NMRFR	3A	8	char	blanks	none	The remote name for the route command if not implied by NMRFD.

Format of an NMR that Contains a Message

Messages are sent in NMRs and can be distinguished from other NMRs because flag NMRFLAGC is off. The format of the NMR varies depending on the settings of bits in NMRFLAG and NMRTYPE. Figure 2-35 describes the format of an NMR that contains a message.

Interpreting the Contents of Field NMRMSG

To determine the contents of the NMRMSG field at offset X'1E', you must also examine the contents of the NMRTYPE4 field at offset X'02' and the NMRTYPE5 field at offset X'00'. Use Figure 2-34 to determine the contents of NMRMSG.

NMRTYPE4	NMRTYPE5	NMRMSG Contents
0	0	Time stamp and the text of the message which starts at NMRMSG + 8.
0	1	The text of the message, which starts at NMRMSG.
1	0	Time stamp, userid, and the text of the message, which starts at NMRMSG + X'10'.
1	1	userid and the text of the message, which starts at NMRMSG + 8.

Figure 2-34. Interpreting Field NMRMSG

Figure 2-35 (Page 1 of 4). Fields in the Messages NMR

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NMRFLAG	0	1	bits	0	none	Flag byte.
bit NMRFLAGC		.1	80 bit	0	0	0 indicates the NMR contains a message.
bit NMRFLAGW		.1	40 bit	0	none	If on, the NMROUT field contains the remote number. Refer to Figure 2-31 on page 2-59 for information on using this bit to interpret the NMRMSG flag.
bit NMRFLAGT		.1	20 bit	0	none	RSCS does not use or set the bit. If on, the NMROUT field contains the destination userid.
bit NMRFLAGU		.1	10 bit	0	none	JES3 and POWER do not use the bit. If on, the NMROUT contains the destination console identifier. Refer to Figure 2-31 on page 2-59 to interpret the contents of field NMROUT.
bit NMRFLAGR		.1	08 bit	0	none	POWER does not use this bit. RSCS does not use the bit. Not used for messages.

Figure 2-35 (Page 2 of 4). Fields in the Messages NMR

Field Name	Offset In Hex	Length	Type of Data	Default Value	Valid Range	Description
bit NMRFLAGJ		.1	04 bit	0	none	Not used for messages.
bit NMRFLAGD		.1	02 bit	0	none	Not used for messages.
bit NMRFLAGS		.1	01 bit	0	none	Not used for messages.
NMRLEVEL	1	.4	bits	0	none	The importance level of the NMR. The bit definitions are: X'10' Non-essential messages X'30' Normal messages X'50' Messages requiring delayed operator action X'70' Essential messages X'80' Messages requiring immediate operator action X'F0' Extremely important messages JES3 and POWER set the importance level to X'70', but do not use it. RSCS does not use the importance level. The importance level is set to X'70' for notify messages and sets it to X'30' for command responses.
NMRPRIO	1	.4	bits	0	none	The output priority. JES3 and POWER set the output priority to X'07', but do not use it. RSCS sets the output priority to X'07'.
NMRTYPE bit NMRTYPEX bit NMRTYPE4	2	1 .4 .1	byte F0 bits 08 bit	0 0 0	none none none	Flag byte. Reserved for IBM's use. If on, the sending userid is included in NMRMSG. Refer to Figure 2-34 on page 2-63 for additional information about using this field to interpret the contents of the NMRMSG field.
bit NMRTYPEP		.1	04 bit	0	none	If on, the timestamp is not imbedded in NMRMSG. Refer to Figure 2-30 on page 2-58 for additional information about using this field to interpret the contents of the NMRMSG and NMROUT fields.
bit NMRTYPEF bit NMRTYPEP		.1 .1	02 bit 01 bit	0 0	none none	Not used for NMRs that contain a message. If on, the NMR contains a delete operator message (DOM) request. DOM is a function of Multiple Console Support (MCS). This bit is not set or used by any system.
NMRML	3	1	binary	1	1-148	Length of the information in NMRMSG, including the userid if it is present. It does not include the timestamp if it is present.
NMRTO	4	9	mixed	none	none	The following two fields describe the destination node and system identifier for the command.
NMRTONOD		8	char	blanks	none	The destination node for the message.
NMRTOQUL	C	1	binary	0	none	The qualifier for the destination node. RSCS and POWER set but do not use the qualifier code.
NMROUT	D	8	char	zeros	none	The destination userid, remote id, or console id depending on the setting of NMRFLAG. Refer to "Interpreting the Contents of the NMROUT Field" on page 2-58 for additional information about interpreting where the NMR originated.

Figure 2-35 (Page 3 of 4). Fields in the Messages NMR

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NMRUCM	D	1	binary	0	none	This field contains the MCS console identifier when NMRFLAGU is set. JES3, RSCS, and POWER do not use or set this field.
NMRUCMA	E	1	binary	0	none	This field contains the MCS console area when NMRFLAGU is set. JES3, RSCS, and POWER do not use or set this field.
NMRLINET	F	2	binary	0	none	The line spacer for a multi-line WTO when NMRFLAGU is set. JES3, RSCS, and POWER do not use or set this field.
NMRDESC	D	2	binary	0	none	This field contains the MCS descriptor codes when NMRFLAGT and NMRFLAGU are not set. JES3 does not use or set the field for messages. RSCS RSCS does not use the descriptor codes because it does not use MCS and does not forward them to the next node. POWER sets the descriptor codes to X'0800' if a destination userid was not specified.
NMRROUT	F	2	binary	0	none	This field contains the MCS routing codes when NMRFLAGT and NMRFLAGU are not set. JES3 does not set the routing codes for messages. When receiving an NMR that contains a message, this field may contain a JES3 console number. RSCS and POWER sets the field to X'4100' if the destination userid is not specified.
NMRCNID	11	4	binary	0	none	The extended console identifier. JES3 does not set or use the field.
NMRRMT	D	8	char	none	none	This field contains the remote name when NMRFLAGW is set. JES2 uses the form 'Rnnnn' for the remote name.
NMRUSER	D	8	char	none	none	This field contains the destination userid when NMRFLAGT is set.
NMRFM NMRFMNOD NMRFMQUL	15	9 8 1	mixed char binary	none blanks 0	none none none	The following two fields describe the originating node and the qualifier. The origin node. The qualifier for the originating node. JES3 sets but does not use the qualifier. RSCS and POWER does not use or initialize the field.
NMRMSG	1E	1-148	mixed	none		May contain the following: <ul style="list-style-type: none"> • 8 bytes that contain the timestamp when the message was issued. • 8 bytes that contain the userid of the operator or user that sent the message. • Up to 132 bytes of message text. Refer to Figure 2-34 on page 2-63 for information about interpreting the contents of the NMRMSG field.

Figure 2-35 (Page 4 of 4). Fields in the Messages NMR

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NMRECSID	1E	8	char	blanks	none	<p>If NMRTYPE4 is set, this field contains the originating userid followed by the message text. Otherwise, this field does not exist.</p> <p>JES3 does not use or set the field.</p> <p>RSCS sets and uses the field for for messages that originate from a VM user.</p>

Networking Connection Control Records (NCCR)

Initial Signon Record

The primary NJE node sends an initial signon record to initiate a session or connection. The initial signon record is the only data in the transmission buffer and is not compressed or compacted. If a session between two nodes is already active and additional signon records are received, the subsequent signon records are ignored.

Figure 2-36. Fields in an NJE Initial Signon Record

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NCCRCB	0	1	binary	X'F0'	none	All path manager RCBs contain X'F0'.
NCCSRCB	1	1	char	none	none	The type of connection control record. Initial signon records contain a C'I'.
NCCIDL	2	1	char			Length of the initial signon record from the RCB to the end of the initial signon record.
NCCINODE	4	8	char	none	none	The EBCDIC name of the primary node.
NCCIQVAL	B	1	binary	0	1-7	The member number of the node.
NCCIEVNT	C	4	binary	0		Not used for initial signon records.
NCCIREST	10	2	binary	0		Partial node-to-node resistance. JES2 uses and sets this field. JES3, RSCS, and POWER do not use or set the field.
NCCIBFSZ	12	2	binary			The maximum size of the transmission buffer that the sending node can receive (must be greater than 300 bytes).
NCCILPAS	14	8	char	blanks		The password for the line for nodes using BSC protocols.
NCCINPAS	1C	8	char	blanks		The password for the node.
NCCIFLG	24	1	flag	bits		This field contains X'00' for an initial signon record.
NCCIFEAT NCCIPREP NCCITRM	25	4 .1 .1 .6	binary 80 bit 40 bit 3F bit	0 0 0 0		Defines space for up to 32 features. BSC/CTC quiesce options. SNA termination options. Reserved for IBM's use.
	26	3	binary	0		Reserved for IBM's use.

Response Signon Record

A node that receives an initial signon record responds with a response signon record if a session or connection is not already established. The response signon record is the only data in the transmission buffer and is not compressed or compacted.

Figure 2-37. Fields in an NJE Response Signon Record

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NCCRCB	0	1	binary	X'F0'		All path manager RCBs contain X'F0'.
NCCSRCB	1	1	char	none	none	The type of connection control record. Response signon records contain a 'C'.
NCCIDL	2	1	binary			Length of the response signon record from the RCB to the end of the extension.
NCCINODE	3	8	char	none	none	The EBCDIC name of the secondary node.
NCCIQUAL	B	1	binary	0	1-7	The member number of the node. (1 indicates the member is a JES2 node.)
NCCIEVNT	C	4	binary	0		Normal signon sequence is 0 if high end sends or secondary trunk of multiple trunk connection. Next higher sequence if low end sends and the trunk is primary. Predefined connection is X'FFFFFFF'.
NCCIREST	10	2	binary	0		Partial node-to-node resistance.
NCCIBUFSZ	12	2	binary			The size of the largest transmission buffer the NJE secondary node can receive.
NCCILPAS	14	8	char	blanks		The password for the line.
NCCINPAS	1C	8	char	blanks		The password for the node.
NCCIFLG NCCIFLGM	24	1 .1 .7	bits 80 bit 7F bits	0		Flag byte. On for multi-trunk responses. Reserved for IBM's use.
NCCIFEAT	25	4	binary	0		Defines space for up to 32 new features. See Figure 2-36 on page 2-67 for bit definitions.

Reset Signon Record

A reset signon record is sent to increment the CES value and/or to change the resistance. The NJE secondary node responds to the request by sending a reset signon record instead of a response signon record. The reset signon record can be included with other NJE records in the transmission buffer but is always placed at the beginning of the transmission buffer. The data in the NJE record is not compressed or compacted.

Figure 2-38. Fields in an NJE Reset Signon Record

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NCCRCB	0	1	binary	X'F0'		All path manager RCBs contain X'F0'.
NCCSRCB	1	1	char	none	none	The type of connection control record. Reset signon records contain a C'K'.
NCCCDL	2	1	binary			Length of the response signon record from the RCB to the end of the concurrence signon record.
NCCCEVNT	3	4	binary			0 if the record is sent by the high end or secondary trunk of a multi-trunk connection. (A reset record cannot be sent from a low end on a normal multi-trunk signon sequence).
	7	3				Reserved for IBM's use.
NCCIREST	10	2	binary	0		Partial node-to-node resistance.

Concurrence Signon Record

A node sends a concurrence signon record when an acknowledgement is expected from a response or reset signon record. The concurrence signon record can be included with other connection records in an NJE record. The data in the NJE record is not compressed or compacted.

Figure 2-39. Fields in a Concurrence Signon Record

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NCCRCB	0	1	binary	X'F0'		All path manager RCBs contain X'F0'.
NCCSRCB	1	1	char	none	none	The type of connection control record. Concurrence signon records contain C'L'.
NCCCDL	2	1	binary			Length of the response signon record from the RCB to the end of the concurrence signon record.
NCCCEVNT	3	4	binary			The contents of the CCCEVNT field from the concurrence signon record.
NCCCREST	7	2	binary			The total node-to-node resistance.

Add/Subtract Connection Record

Add/subtract connection records inform other members of a network of the connection status between two nodes. JES2 is the only networking product that sends add/subtract connection records. Members of a MAS configuration may have identical node names and can be distinguished by the member names. The add/subtract connection record can be included with other connection records in an NJE record. The data in the NJE record is not compressed or compacted.

Figure 2-40. Fields in an Add/Subtract Connection Record

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NCCRCB	0	1	binary	X'F0'		All path manager RCBs contain X'F0'.
NCCSRCB	1	1	char	none	none	The type of connection control record. Subtract connection records contain C'N'. Add connection records contain C'M'.
NCCADL	2	1	binary			Length of the response signon record from the RCB to the end of the add/subtract connection record.
NCCANODA	3	8	char			The name of the low-end node.
NCCAQULA	B	1	binary	1	1-7	Member number for the low-end (if shared spool).
NCCANODB	C	8	char			The name of the high-end node.
NCCAQULB	14	1	binary	1	1-7	Member number for the high-end (if shared spool).
NCCAEVNT	15	4	binary			The connection event sequence.
NCCAREST	19	2	binary		0-8000	The total node-to-node resistance.

Signoff Records

Signoff records inform the other node that the connection or session between two nodes is terminated.

Figure 2-41. Fields in a Signoff Record

Field Name	Offset in Hex	Length	Type of Data	Default Value	Valid Range	Description
NCCRCB	0	1	binary	X'F0'		All path manager RCBs contain X'F0'.
NCCSRCB	1	1	char	none	none	The type of connection control record. Signoff records contain C'B'.

Chapter 3. Presentation Services

Non-SNA Buffer Format	3-2
Transmission Buffer Size	3-2
Compressing Data for BSC and CTC Transmissions	3-2
BSC and CTC Buffers	3-2
Block Control Byte (BCB)	3-4
Function Control Sequence (FCS)	3-7
String Control Byte (SCB)	3-8
Record Control Byte (RCB)	3-9
Subrecord Control Byte (SRCB)	3-11
SNA Buffer Format	3-13
Transmission Buffer Size	3-13
Compressing and Compacting Data for SNA Transmissions	3-13
SNA Transmission Buffer	3-13
String Control Byte (SCB)	3-14
Compression SCB	3-15
Compaction SCB	3-15
Record Identifier (RID)	3-15
Data Records	3-17
Format of Data Records Containing Unspanned Data	3-17
Format of Data Records Containing Spanned Data	3-18
Trailing Blank Truncation	3-19

NJE packages data in NJE records within transmission buffers. It uses one buffer format for non-SNA transmissions (BSC and CTC) and another for SNA transmissions. This chapter describes the formats of those buffers in detail. It further describes the data within those buffers.

Non-SNA Buffer Format

Transmission Buffer Size

The minimum size of a transmission buffer is 300 bytes. This size allows the transmission buffer to contain a complete 256-byte record, compression bytes, and BSC control bytes. The maximum size of a transmission buffer is 32K. A BSC connection can use a smaller buffer size than a CTC connection because the time to retransmit and the probability of errors on a BSC connection are greater than on a CTC connection.

Compressing Data for BSC and CTC Transmissions

For BSC and CTC NJE transmission, data compression is always used; compression reduces the length of records for transmission by removing blanks and duplicate characters. String control bytes (SCBs) are used to compress data before placing each record into the transmission buffer. The SCB is described in detail later in this section.

Note: BSC and CTC NJE transmissions use compression only; compaction (reducing the length of records by representing certain 8-bit characters with only 4 bits) is **not** used.

BSC and CTC Buffers

A variable number of NJE records can be combined into a transmission buffer. Each record in the buffer is comprised of a series of character strings and prefixed by a string control byte (SCB). Each NJE 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 buffer. Finally, a block control byte (BCB) is added as the first character of each transmission buffer for error detection and correction.

Following is the layout of a typical transmission buffer for BSC and CTC communications.

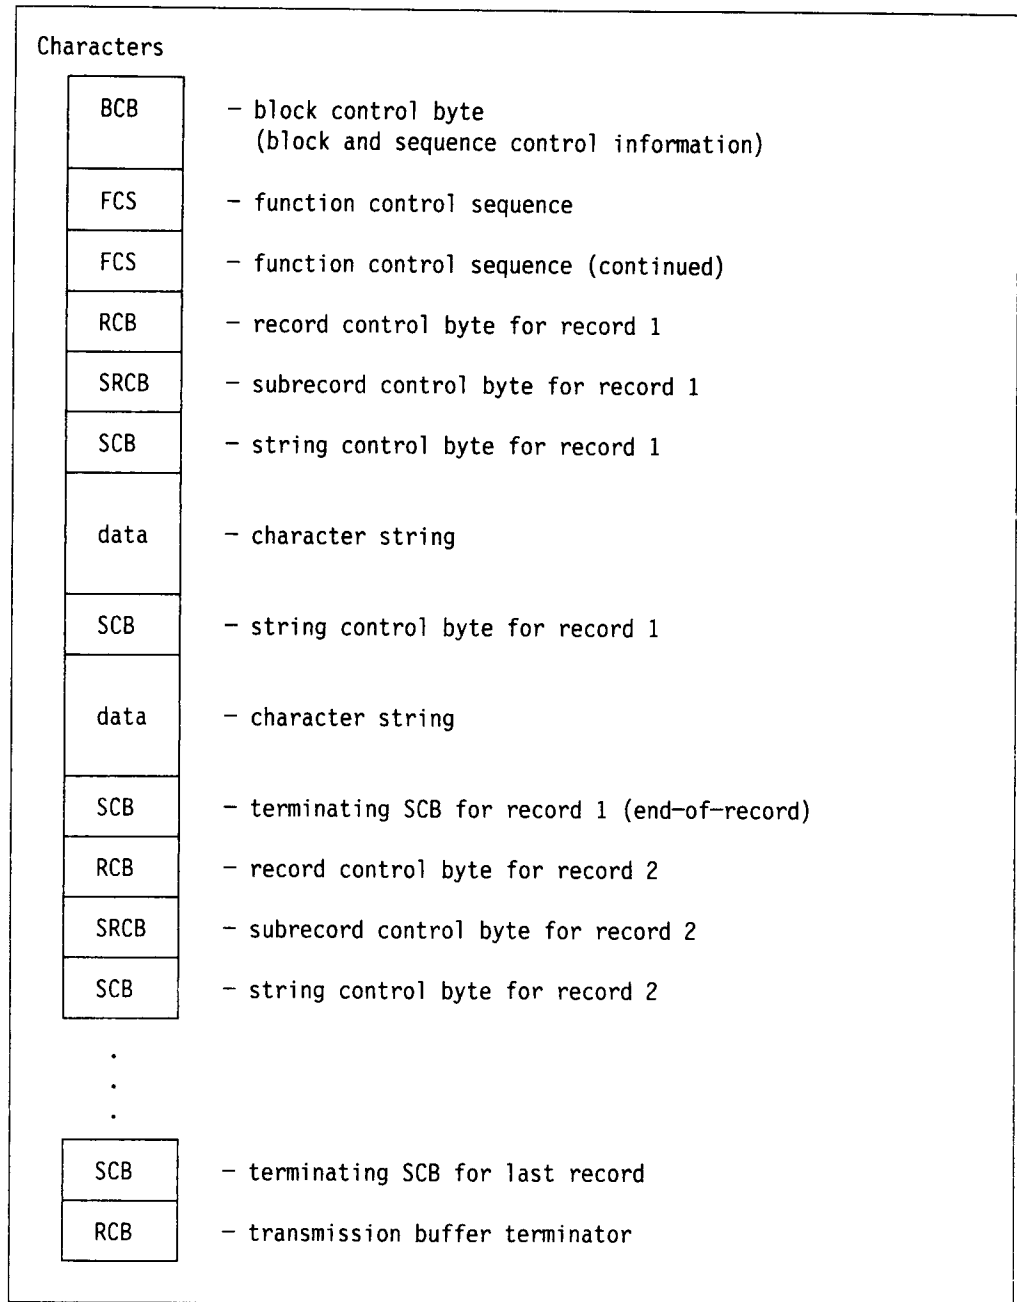


Figure 3-1. Format of a Transmission Buffer for BSC and CTC

Block Control Byte (BCB)

Every BSC buffer begins with a block control byte (BCB) which contains inbound and outbound buffer sequence counters. These counters are used to synchronize inbound and outbound transmissions and to detect and correct sequence errors. Figure 3-2 shows the BCB bit definitions.

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 IBM's use
11xx	Reserved for IBM's use
1... cccc	Modulo 16 sequence counter

Figure 3-2. BCB Bit Definitions

After a BSC line completes initialization, each directly-attached node initializes the inbound and outbound BCB counters to 0000 and maintains it in modulo sixteen. The outbound BCB count is incremented by one each time an acknowledgement (ACK) is received. The inbound BCB count is also updated to reflect the number of transmission buffers that the node has received. When a node receives a BCB the outbound count should be one more than the last BCB it received (the inbound count). If so, data transmission is normal and no data has been lost.

If the count in the BCB is not what is expected, the receiving node must indicate an error by sending an RCB indicating a BCB sequence error. When the sequence error is received, the node receiving it terminates the connection because error recovery is not possible.

If a duplicate BCB is received, the system receiving it assumes the last transmission buffer it sent has been lost and its last transmission buffer that did not receive an acknowledgement must be re-sent. The duplicate buffer is discarded. This allows the NJE network to recover from errors rather than terminating the connection by indicating a sequence error.

Because the BCB is also used to detect lost blocks, a null block containing a BCB should always be sent as an acknowledgement rather than using a DLE ACK0 as the acknowledgement. Figure 3-3 and Figure 3-4 clarify the importance of using the DLE as an acknowledgement.

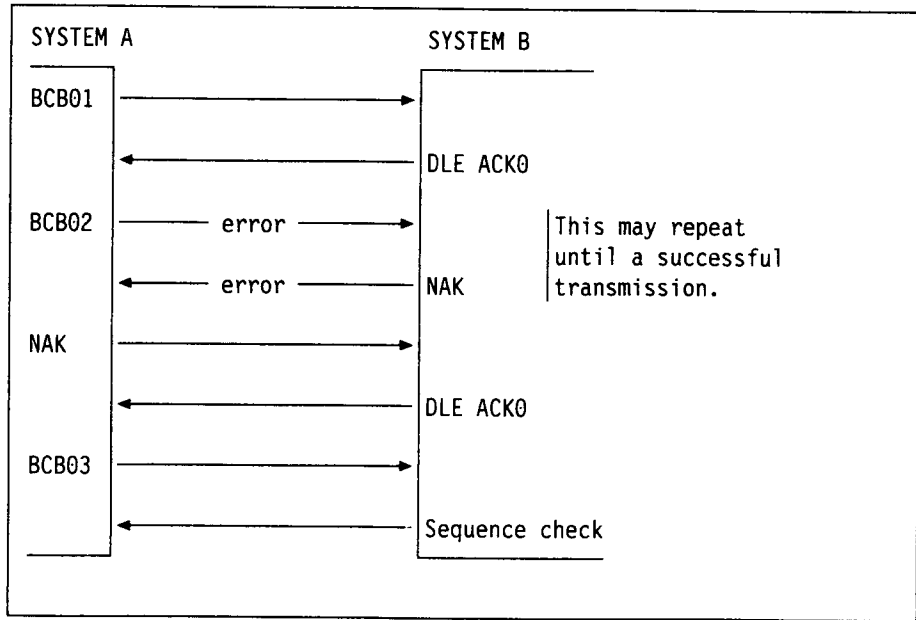


Figure 3-3. Results of Not Transmitting Null Records.

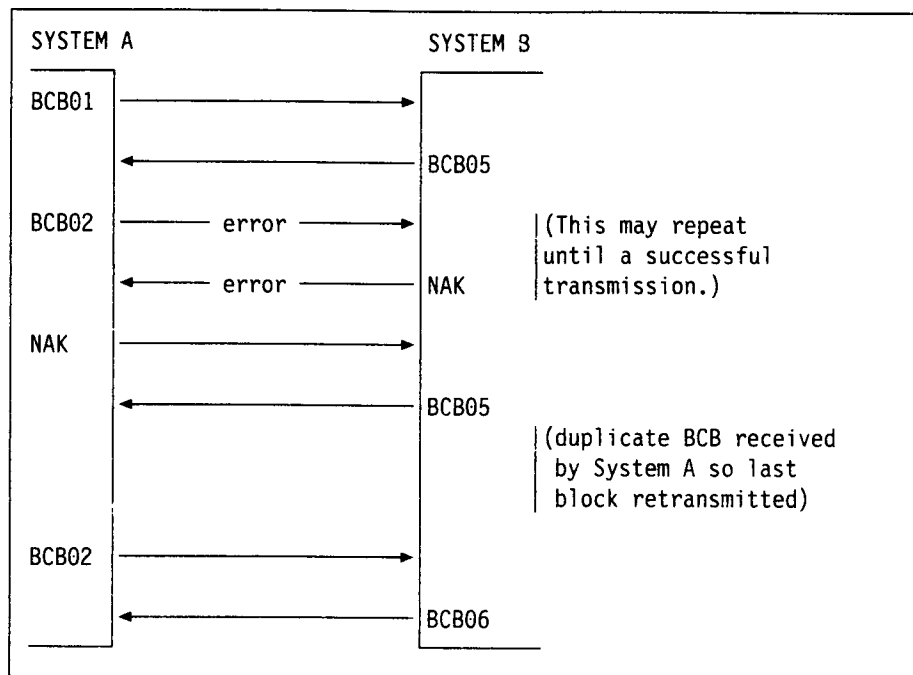


Figure 3-4. Correct Recovery with Null Records

Use Figure 3-5 to determine the state a node is placed in and the action it takes after receiving a transmission block.

Figure 3-5. Transmission Block Handling State Table

Current State of Node	Event	Action Node Takes	New State of Node After Action is Taken
Waiting for a response	Receives a data block		Checks the BCB
	Receives DLE ACK0	Transmits last non-NAK block of data	Waits for a response
	Receives a NAK	Retransmits last non-NAK	Waits for a response
	Receives an error (not CE-DE status)	Increments error count	Checks error count
Checking a BCB	Corrects the BCB	Transmits last non-NAK block of data	Waits for a response
	Detects a duplicate BCB	Retransmits the last non-NAK	Waits for a response
	Detects an incorrect BCB	Transmits a sequence error	Waits for a response
Recovering from an error and waiting for a response	Receives a data block		Checks the BCB
	Receives DLE ACK0	Transmits last non-NAK block of data	Waits for a response
	Receives a NAK	Retransmits last non-NAK	Waits for a response
	Receives an error (not CE-DE status)	Increments error count	Checks error count
Checking error count	Error count is not exceeded	Transmit a NAK	Recovers from the error and waits for a response

Function Control Sequence (FCS)

Following the BCB are two bytes used by each receiver to control inbound flow. These bytes are called the function control sequence (FCS). There is one bit to hold (off) or enable (on) *each* of the eight inbound streams and one bit to hold/enable *all* inbound streams. The latter bit is referred to as the "wait-a-bit". Figure 3-6 defines the binary representation and the meaning of bits in the FCS.

Binary	Meaning
1... .. 1... ..	Must be 1... .. 1... ..
.0..	Normal state
.1..	Suspend all stream transmission (Wait-a-bit)
..rrrr ..	Reserved for IBM's use
.... 1... ..	Function stream identifier for SYSIN stream 1
.... .1..	Function stream identifier for SYSIN stream 2 or SYSOUT stream 7
.... ..1.	Function stream identifier for SYSIN stream 3 or SYSOUT stream 6
.... ...1	Function stream identifier for SYSIN stream 4 or SYSOUT stream 5
....1..	Function stream identifier for SYSIN stream 5 or SYSOUT stream 4
....1.	Function stream identifier for SYSIN stream 6 or SYSOUT stream 3
....1.	Function stream identifier for SYSIN stream 7 or SYSOUT stream 2
....1	Function stream identifier for SYSOUT stream 1

Figure 3-6. FCS Stream Bit Definitions

The FCS bit assignments cause the BSC protocol restriction of a maximum of eight SYSIN and SYSOUT streams concurrently in each direction.

The compressed transmission records are placed in the buffer after the FCS bytes. Each transmission buffer contains transmission records that belong to the same NJE job. NJE protocols do not allow transmission records from different streams to be placed in the same transmission buffer.

The system receiving an FCS containing Wait-a-bit that is on, should delay responding. The delay can be in the range of one or two seconds if the Wait-a-bit was sent with a null record. A delay of greater than two seconds is not recommended because a timeout will occur if the response is not received within three seconds. If the data was sent with the Wait-a-bit, 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 no data to transmit, but cannot receive data.

The node receiving the FCS containing a Wait-a-bit that is on should respond with either a:

- DLE ACK0
- Null buffer.

IBM recommends responding with null buffers. If a DLE ACK0 is sent, the node receiving it responds to the DLE ACK0 based on the last FCS that was sent. (This is the normal response to a DLE ACK0). NJE networking does not allow a node to transmit a data record in response to a Wait-a-bit record.

String Control Byte (SCB)

In BSC and CTC NJE transmissions, the record segments are first compressed. This compression results in string control bytes (SCBs) interspersed in the stream data with duplicate character strings compressed. The resulting format always begins with an initial SCB and may have additional SCBs interspersed.

The SCB defines how many blanks or repeating non-blank characters are to be inserted by the receiver (up to 31) or the length of strings without any duplicate characters (up to 63 characters). An SCB **must** appear at least once every 64 bytes, and describe up to 63 bytes of following data. A "stand-alone" SCB with a value of zero, sometimes referred to as an "EOR SCB", is always placed at the end of each NJE record.

The SCB bit settings specify the compression parameters. SCBs are always used, whether or not the data is actually compressed; that is, even non-compressed data must be interspersed with SCBs that have the two high-order bits set (see Figure 3-7).

<i>Figure 3-7. SCB Definition</i>		
2 high-order bits	Remaining bits	Meaning
00	00 0000	Indicates the end of the NJE record.
01	00 0000	Indicates the transmission of the stream should be terminated.
10	0b bbbb	Indicates that 'bbbb' blanks should be inserted after the SCB.
10	1d dddd	Indicates that the single character following the SCB should be duplicated 'dddd' times.
11	cc cccc	Indicates that 'ccccc' uncompressed characters (maximum of 63) follow the SCB. (The record is compressed; these characters are not.)

Record Control Byte (RCB)

Each logical record begins with a record control byte (RCB), which contains the stream identifier (type and number). 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 uncompressed records (signon, signoff, and path manager records).

A transmission buffer cannot contain records with different RCBs. The connection must be terminated (all streams) if a transmission buffer is received that contains an unexpected, unrecognized, or incorrect RCB. This includes:

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

When the next compressed transmission record will not fit in the buffer, a special "stand-alone" RCB of X'00' (EOB) is placed after the last record and the buffer is truncated at this point for transmission.

Note that segments are never spanned across transmission buffers, however, uncompressed records may be.

Valid RCBs received out of sequence will also cause a connection to terminate. For example, a X'B0' receiver cancel with an unstarted stream referenced in the SRCB constitutes an error situation.

A request to start a stream that has already been started is handled differently. In this case, the request is rejected with the X'B0' (permission denied) RCB. The connection is not terminated; instead, the transmitting system can terminate the stream upon receipt of the X'B0' RCB. Figure 3-8 on page 3-10 depicts possible binary and hexadecimal values and their meanings in a record control byte.

Binary	Hex	Meaning	
0000	0000	00	End-of-block (BSC)
<i>rrrr</i>	<i>rrrr</i>	01-8F	Reserved
1001	0000	90	Request to initiate stream (SRCB = RCB of stream to be initiated)
1010	0000	A0	Permission to initiate stream (SRCB = RCB of stream to be initiated)
1011	0000	B0	Negative permission or receiver cancel (SRCB = RCB of stream to be denied)
1100	0000	C0	Acknowledge transmission complete (SRCB = RCB of stream received)
1101	0000	D0	Ready to receive stream (SRCB = RCB of stream to be received)
1110	0000	E0	BCB sequence error
1111	0000	F0	General control record
1001	0001	91	Reserved
1 <i>rrr</i>	0001	A1-F1	Reserved
1001	0010	92	Reserved
1 <i>rrr</i>	0010	A2-F2	Reserved
1 <i>iii</i>	0011	93-F3	Reserved
1 <i>iii</i>	0100	94-F4	Reserved
1 <i>iii</i>	0101	95-F5	Reserved
1 <i>iii</i>	0110	96-F6	Reserved
1 <i>rrr</i>	0111	97-F7	Reserved
1 <i>iii</i>	1000	98-F8	SYSIN record
1 <i>iii</i>	1001	99-F9	SYSOUT record
1001	1010	9A	Operator command/console message
1 <i>rrr</i>	1010	AA-FA	Reserved
1 <i>rrr</i>	1011	9B-FB	Reserved
1 <i>rrr</i>	1100	9C-FC	Reserved
1 <i>rrr</i>	1101	9D-FD	Reserved
1 <i>rrr</i>	1110	9E-FE	Reserved
1 <i>rrr</i>	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 3-8. RCB Definition

For complete system-dependent and product-specific information, see Appendix A, "System-Dependent Considerations."

Subrecord Control Byte (SRCB)

The subrecord control byte (SRCB) is interpreted according to the RCB value, as shown in the following figures. The SRCB defines the record type (job header, data set header, job trailer, or user data record) plus SYSOUT carriage control type and user data spanning indicators (first, last, or middle segment). Refer to "Format of Data Records Containing Spanned Data" on page 3-18 for additional information on records containing spanned data.

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 completed stream
D0	RCB of ready stream receiver
E0	Expected count - BCB sequence error (received count is in BCB)
F0	An identification character as follows: <ul style="list-style-type: none">A = ReservedB = Network SIGNOFFC-H = ReservedI = Initial network SIGNON<ul style="list-style-type: none">Must be only record in transmission bufferJ = Response to initial network SIGNON<ul style="list-style-type: none">Must be only record in transmission bufferK = Reset network SIGNONL = Accept (concurrency) network SIGNONM = Add network connectionN = Delete network connectionO-R = Reserved for IBM's useS-Z = Unused
91-94	Reserved

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

Figure 3-9. SRCB Definition - Part 1

RCB	SRCB
95-F5	Reserved
96-F6	Undefined
97-F7	Undefined
98-F8	NJE SYSIN control information as follows: 1000 0000 - Standard record 1100 0000 - Job header 1110 0000 - Data set header 1101 0000 - Job trailer 1111 0000 - Reserved 1111 0000 - Reserved for IBM's use
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) 10cc ss00 - Spanned record control as follows: ----- 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 - Reserved for IBM's use
9A	Operator Command/Message (NMR) 1000 0000 (X'80')

Figure 3-10. SRCB Definition - Part 2

For system-dependent and product-specific information, see Appendix A, "System-Dependent Considerations."

SNA Buffer Format

The SNA data buffer is called a request/response unit (RU). The RU carries control information and data between logical units. A **control RU** contains a request or acknowledgement, while a **data RU** contains function management headers, or SCBs and data.

Transmission Buffer Size

A data RU may contain as many job networking records as can fit into the RU. The size of a data RU is determined by the BIND parameters specified at the start of the session. The maximum size of an RU can be 65535 bytes. The maximum length for the job networking record is 259 bytes (256 bytes of user data plus a 3-byte record identifier).

Compressing and Compacting Data for SNA Transmissions

SNA data transmission requires the use of compression. Compression reduces the length of records for transmission by removing blanks and duplicate characters. Compaction (reducing the length of records by representing certain 8-bit characters with only 4 bits) is optional and is controlled by NJE and SNA protocols. (See "Session Initialization" in Chapter 6 for additional information on compaction and compression.)

The SNA string control byte (SCB) defines the beginning and end of compressed and compacted data. An SCB must begin each RU following the exchange of FM headers. Refer to "String Control Byte (SCB)" on page 3-8 for additional information for the SCB.

SNA Transmission Buffer

SNA transmission buffers contain SCBs, RIDs, and data records. Each data record is prefixed with a 3-byte header, called a record identifier (RID), which identifies the type of record that follows. The purpose of the RID is similar to that of the RCB and SRCB in non-SNA NJE transmissions. The three bytes of the RID are as follows:

RIDRCB	Contains the address of the RCB.
RIDSRCB	Contains the address of the SRCB.
RIDLEN	Contains the length of the data. NJE networking uses RIDLEN to locate the beginning of the next data record in the transmission buffer.

Figure 3-11 illustrates the format of a SNA transmission buffer.

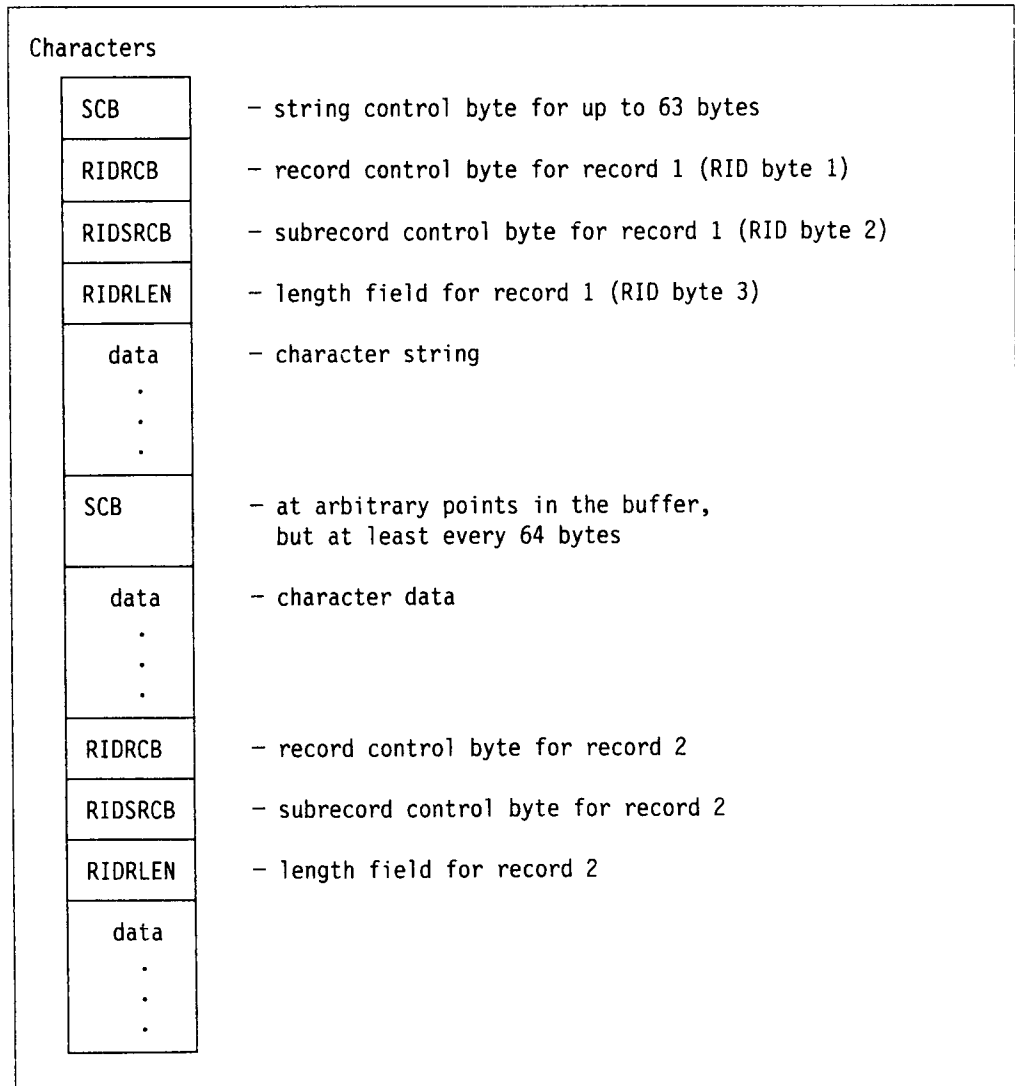


Figure 3-11. Format of a Transmission Buffer for SNA

String Control Byte (SCB)

In contrast with non-SNA use of SCBs, for SNA buffers, SCBs are used to compress/compact the entire buffer; that is, the entire buffer is treated as data as far as the SNA transport subsystem is concerned.

The SCB bit settings specify the compression and compaction parameters, but the encodings are different for SNA than for non-SNA transmissions. SCBs are always used, whether or not the data is actually compressed; that is, even uncompressed data must be interspersed with SCBs that have the two high-order bits set.

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 3-12.

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

Figure 3-12. String Control Byte (SCB) Format for SNA

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

Compression SCB

Compression may be optionally indicated in the BIND for LU TYPE 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 SCB

Compaction allows certain 8-bit character sequences to be represented in network transmission as 4 bits. The compacted characters 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 in effect. All other characters are considered non-compactable and will be transmitted in their 8-bit form. Obviously, the non-compactable characters should be those that are least frequently used in the data stream.

Master and non-master characters are transmitted via FMH3 at session initialization.

Record Identifier (RID)

A record identifier (RID) is a 3-byte header that precedes every NJE SNA record. The RID identifies the type of record that follows. 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 following figures depict the format of the three bytes of the RID:

1. RIDRCB - Figure 3-13
2. RIDSRCB - (depends on value of RIDRCB) Figure 3-14 on page 3-16
3. RIDRLN - (depends on value of RIDRCB) Figure 3-15 on page 3-17.

Note: Byte 2 of the RID, RIDSRCB, contains the same values as the SRCB for non-SNA records.

Value	Record type	Meaning
X'90'	stream	Request to allocate SYSIN/SYSOUT stream
X'A0'	control stream	Permission to allocate stream granted
X'B0'	control stream	Permission to allocate stream denied or receiver cancel
X'C0'	control stream	Acknowledge end of transmission
X'D0'	control stream	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 3-13. RIDRCB: Byte 1 of the RID

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

The *iii* in the value field for RIDRCB identifies the particular SYSIN or SYSOUT stream. Values of 1-7 binary are permitted; a value of zero is not allowed.

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 the initiated SYSIN/SYSOUT receiver
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 canceled by transmitter
X'99'-X'F9'	NJE SYSIN control information Transmission end information X'00' - standard SYSOUT stream end X'40' - SYSOUT stream canceled by transmitter
X'9A'	NJE SYSOUT control information X'80' (JES2 sends an SRCB of X'00' with NMRs.)

Figure 3-14. RIDSRCB: Byte 2 of the RID

RIDRCB	RIDRLEN
X'90'	N/A (set to zero)
X'A0'	N/A (set to zero)
X'B0'	N/A (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 of data record - 1 ¹
X'99'-X'F9'	Length of data record - 1 ¹
X'9A'	Length of nodal message record - 1 ¹

Figure 3-15. RIDRLEN: Byte 3 of the RID

Data Records

This section describes unspanned/spanned data format and trailing blank truncation. This information applies to both SNA and non-SNA NJE transmissions as described in the previous sections.

Format of Data Records Containing Unspanned Data

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.




Figure 3-16. Data Record Without Carriage Control

If machine carriage control is used, any CCW operation code (opcode) may be sent in that record. Store-and-forward systems must forward all CCW opcodes intact. The destination may discard any records that it cannot process.




Figure 3-17. Data Record With Carriage Control

For NJE SYSIN jobs, the default record length and format for each record is a 80 bytes fixed. SYSIN data records do not contain carriage control.

¹

For the last three values shown in the figure, a value of zero in RIDRLEN indicates that no data follows. The values X'01' - X'FF' represent record lengths of 2-256. Because NJE records begin with a one-byte length, a record length of 1 need not be represented.

Format of Data Records Containing Spanned Data

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 32760 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 are also not included in the SEGL. A sample spanned record transmission is shown in Figure 3-18.

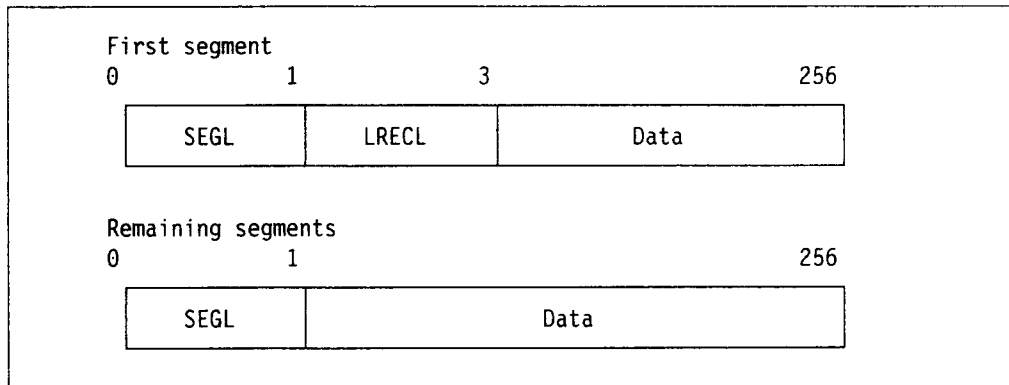


Figure 3-18. Spanned Data Record

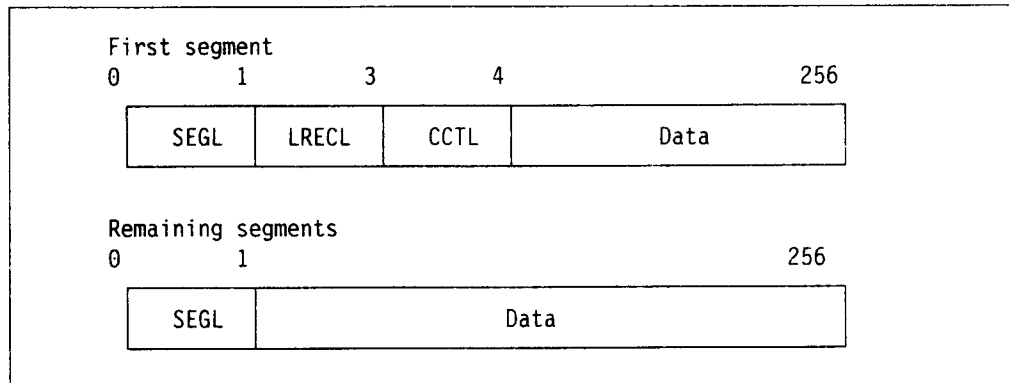


Figure 3-19. Spanned Data Record With Carriage Control

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

If machine carriage control is used, any CCW opcode may be sent in that record. Store-and-forward systems must forward all CCW opcodes intact. The destination may discard any records that it cannot process. For further information, see the appropriate product-specific sections in Appendix A.

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.

Compression and compaction may also be applied for transmission, but these have reconstruction mechanisms independent of the original LRECL.

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 4. Stream Control

Initiating Stream Transmission	4-2
BSC and CTC Protocols for Stream Control	4-3
SNA Protocols for Stream Control	4-3
Abnormal Stream Termination	4-4

This chapter describes the multi-leaving functions of NJE. After NJE records are placed into transmission buffers, the buffers are sent in a stream to a remote node. Multi-leaving describes the capability to concurrently transfer multiple streams on the same BSC connection or SNA session. NJE protocols also allow for full duplexing. Full duplexing refers to the ability to send a stream in one direction while receiving an unrelated stream on the same line in the other direction.

Any differences in multi-leaving between the protocols are noted.

Initiating Stream Transmission

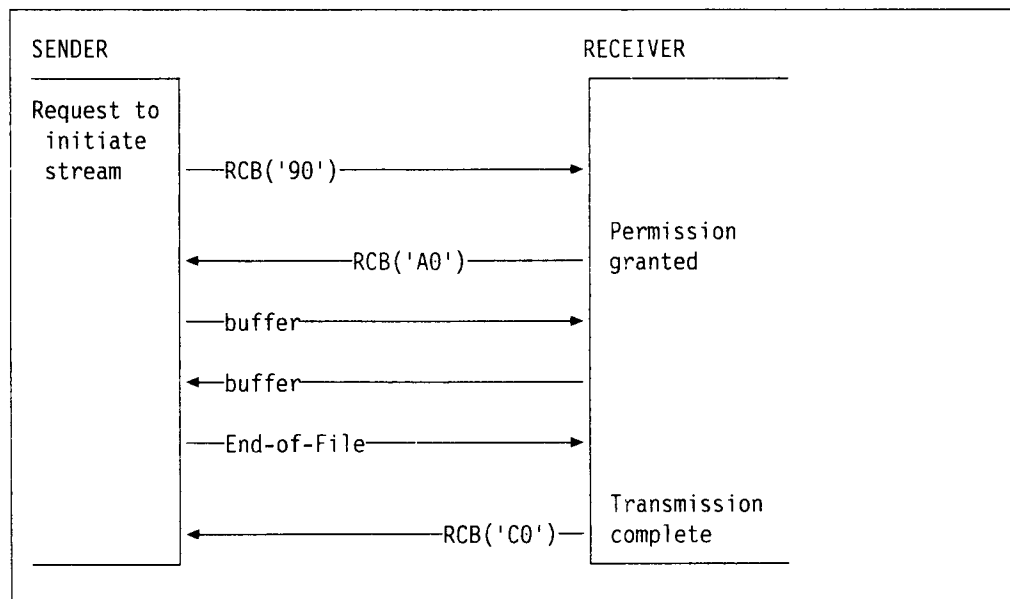


Figure 4-1. Protocols Used to Transmit a Stream

When a node is ready to transmit a stream, it prepares a control record to initiate the stream. An RCB with a value of X'90' is the control record that indicates a request has been made to initiate a stream. The receiving node receives the request and prepares a control record to indicate if it can receive the stream. If the node can receive the stream, the receiving node responds by preparing an RCB with a value of X'A0'.

Otherwise, the receiving node prepares an RCB with a value of X'B0' (negative permission control record) to indicate it cannot accept the stream. When a negative permission control record is returned, the sending node should either wait for an interval of time or until it receives a receiver-initiated control record (RCB X'D0') for that stream before sending another request to initiate the stream.

Note: If JES2 or RSCS are the subsystems at the sending node, the job is held if a negative permission control record is received. POWER and JES3 drain the NJE jobs that are waiting to be sent and requeue the NJE jobs so that the job can be resent.

The sending node cannot send any transmission buffers until the proper control records are exchanged. The sender may only send a transmission buffer after the receiver has responded with a permission to initiate a stream (RCB X'A0'). Only one job may be sent per request to initiate a stream. If the receiving node receives a request to initiate a stream that is already active, the receiving node should return a negative permission (RCB X'B0').

The sending node sends an end-of-file control record when transmission for the stream is complete.

When transmission complete (RCB X'C0') is received, the sender may purge all copies of the job on his own system. It may not do so before RCB X'C0' is received. At this point the stream is logically closed.

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.

BSC and CTC Protocols for Stream Control

BSC protocols allows:

- A total of 8 NJE SYSIN or NJE SYSOUT jobs in either direction
- An NMR
- A control stream.

BSC and CTC connections perform "pseudo" full duplexing because each node must flip-flop between the send and receive states. Both nodes are synchronized.

The sender should not activate one stream at a time that uses a given FCS stream suspension bit. The receiver should respond with a negative permission control record for any invalid requests.

Note: The restrictions on the number of streams that may be concurrently active are documented in "Function Control Sequence (FCS)." Some products implement the same restrictions for SNA stream control, although NJE does not impose them.

SNA Protocols for Stream Control

SNA protocols allow:

- the concurrent flow of 28 NJE SYSIN and/or SYSOUT jobs in either direction
- An NMR
- A control stream.

All NJE records flow on logical streams, which correspond to transmitter and receiver pairs on the network link. NJE records are controlled by destination indications in the record identifier (RID). For information about the contents of the RID, refer to "Record Identifier (RID)" on page 3-15.

NJE logical streams must be initiated before any data can be sent. An SNA stream control record must be the only record in the request/response unit (RU).

Abnormal Stream Termination

Either the sender or the receiver may terminate the transmission of a job before the end-of-file or transmission complete control record is sent. This abnormal termination may be caused by any problem either the sender or receiver has with a particular job *after* the initiating protocol has completed (that is, the receiving node has sent a permission granted RCB X'A0' control record).

Note: When the sending node uses SNA protocols, a RIDSRCB of X'40' rather than an SCB of X'40' indicates the sending node initiated stream termination. (The SCB is described in "String Control Byte (SCB)" on page 3-8.)

The abnormal termination protocol when the sending node initiates terminating the session for the receiving node is as follows:

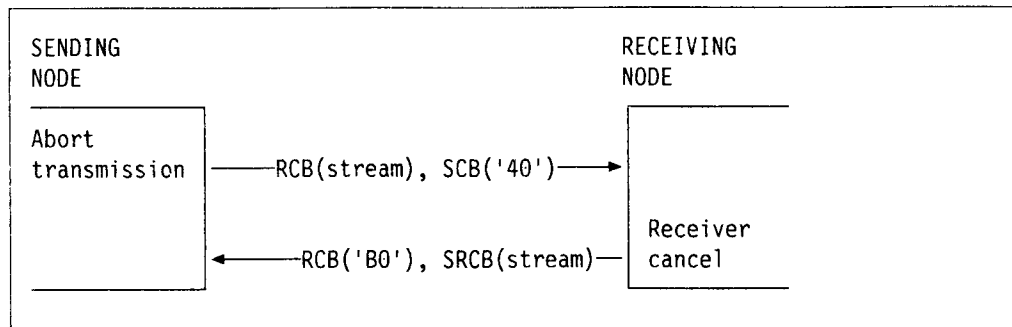


Figure 4-2. Sending Node Initiates Stream Termination

No additional data may be sent on this stream until the receiver cancel control record 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 termination was caused by system problems, rather than a user or operator cancel request, a copy of the terminated job should be kept in the sender's spool space for transmission at a subsequent time.

Abnormal termination protocol is as follows:

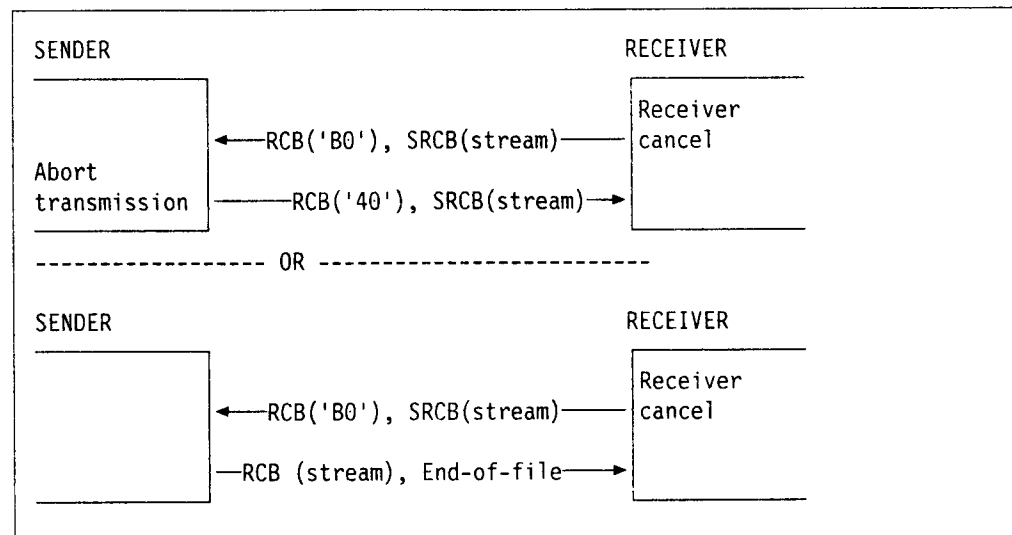


Figure 4-3. Receiving Node Initiates Stream Termination

When the sender receives the receiver cancel control record, the following steps must be performed at the sending node:

1. Stop sending the job for the stream specified in the receiver cancel message.
2. Send a terminate transmission control record to the receiver to show that the transmission of the specified job has been stopped, or send an end-of-file control record 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 terminated stream. A request for permission to transmit stream control records must be sent before beginning 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 terminates the job once, it will be terminated during subsequent attempts to retransmit the job.

Chapter 5. NJE Connection Control

Signon Concurrence Feature	5-2
External Interface	5-2
Implementation Example	5-2
Protocol Description	5-3
Establishing a Connection Without NPM	5-3
Establishing a Connection or Session With NPM	5-4
Signon Connection Protocol Between Two Nodes With NPM	5-4
Full Primary Trunk Protocol (Low End Initiation)	5-5
Full Primary Trunk Protocol (High End Initiation)	5-5
Full Secondary Trunk Protocol	5-6
Full Reset Trunk Protocol	5-6
Signon Protocols Between A Node With NPM and A Node Without NPM	5-7
Connection Status Information	5-7
Disconnections	5-7
Terminating a Session or Connection	5-8

Depending on the networking facility and the type of protocol a node uses to transmit a transmission buffer, a node can establish a connection with another node in the network.

If one of the nodes attempting to establish a connection has JES2 installed, the network path manager (NPM) influences the protocols used to establish the connection. JES2 installations can include the the CONNECT initialization statement in their initialization streams to specify a node does not use NPM. Refer to “Establishing a Connection or Session With NPM” on page 5-4 for the protocols used if you are establishing a connection with a node that has JES2 installed. Refer to “Establishing a Connection Without NPM” on page 5-3 for the protocols used if you are establishing a connection with nodes that *do not* have JES2 installed. The Signon Concurrence Feature is an option your installation can implement. “Signon Concurrence Feature” describes the feature.

Signon Concurrence Feature

The signon concurrence feature 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.

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 to be used if the system at the other end concurs.

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 corresponding extended NJE feature is used between them.

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 the SSFW into NCCIFEAT before transmitting the I record.
- Secondary
 1. Copy NCCIFEAT from the I record into the AFW.
 2. AND the AFW with the SSFW, storing result in the AFW.
 3. Copy the AFW into NCCIFEAT of the J record.
- Primary
 1. Copy NCCIFEAT from the J record into the AFW.
 2. AND the AFW with the SSFW and store the result in the AFW.
 3. Compare the AFW with NCCIFEAT. If the fields are different, terminate the line or use the features indicated by NCCIFEAT.

The previous 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 an I or J record must then assume that none of the optional features are to be used on the session.

Protocol Description

The system transmitting the I record (primary) must set the features bits in the 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 an I record is received, the secondary node must look at the extended features bits and for any features that it is also willing to use and must transmit the corresponding feature bit in the J record. The secondary node should then use the features. Upon receiving the J record, the primary node should then examine the features bits and use any features that are specified. See "Initial Signon Record" on page 2-67 for a description of the actual bit definitions.

Establishing a Connection Without NPM

The nodes first attempt to establish a connection by exchanging a set of control records. The control records vary depending on the type of protocols the nodes will use to transmit the stream.

The NJE primary node always sends the initial signon record. If BSC or CTC protocols are being used, the NJE primary node is the node that receives the DLE ACK0 BSC control record. If SNA protocols are being used, the NJE primary node is the node with the higher node name. The SRCB that precedes the initial signon record contains an identifier of C'I'. "Initial Signon Record" on page 2-67 describes the format of the initial signon record.

The node receiving the initial signon record, called the NJE secondary node, responds with a response signon record. The SRCB that precedes the response signon record contains an identifier of C'J'. The J record contains a connection event sequence of minus one (X'FFFFFF'). Figure 5-1 on page 5-3 describes the format of the response signon record.

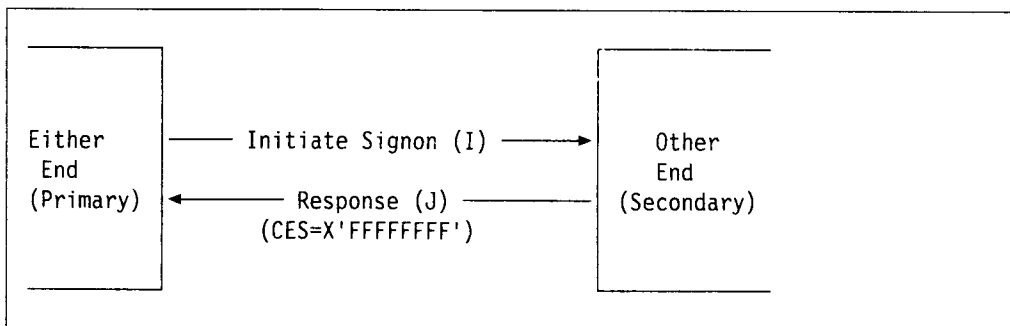


Figure 5-1. Connection Protocols used by the Non-Path Manager Nodes

Establishing a Connection or Session With NPM

A node that has JES2 installed uses the Network Path Manager (NPM) to establish a connection or session between nodes. NPM is also responsible for:

- Promulgating connection information to the other members.
- Maintaining information about which lines should be used to reach a given node.
- Informing other subcomponents which nodes should be reached over a given line.
- Providing routing information for jobs and messages.
- Processing NJE signon and connection/disconnection records from other Network Path Managers.
- Selecting a line based on specified resistances.

Nodes that use NPM may have code parameters on initialization statements that prohibit connect and disconnect records from being sent for some connections. The following are examples of connections where the installation may prevent connect and disconnect records from being sent:

- **Private connections.** Connection records are created for the directly connected node. For private connections no other nodes are notified of the nodes existence.
- **Non-path manager connections** Connection records are not created when a node using NPM signs onto a node not using NPM.

Signon Connection Protocol Between Two Nodes With NPM

The protocol for directly connecting two NJE nodes depends upon the capabilities of the two Network Path Managers, the number of lines between the two systems, and the installation-supplied names of the nodes. 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 always increasing, NPMs can decide what information is the most recent, discarding any old connection information. When CES values are assigned, the path manager ensures 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 depend 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 consistent throughout the network, it is necessary that all systems using the path manager use Greenwich Mean Time (GMT) settings in their TOD clocks, and that these clocks are properly synchronized. This also affects the accuracy of the reader start time and other SMF information.

Full Primary Trunk Protocol (Low End Initiation)

This form of signon protocol is used on BSC connections only. Figure 5-2 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.

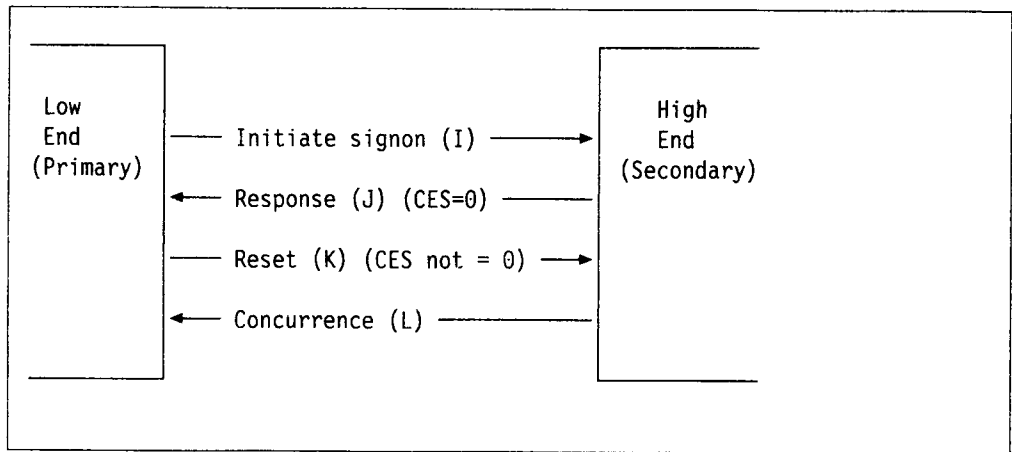


Figure 5-2. Normal Connection (Low Initiates)

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 5-3 illustrates high end-initiated full connection protocol.

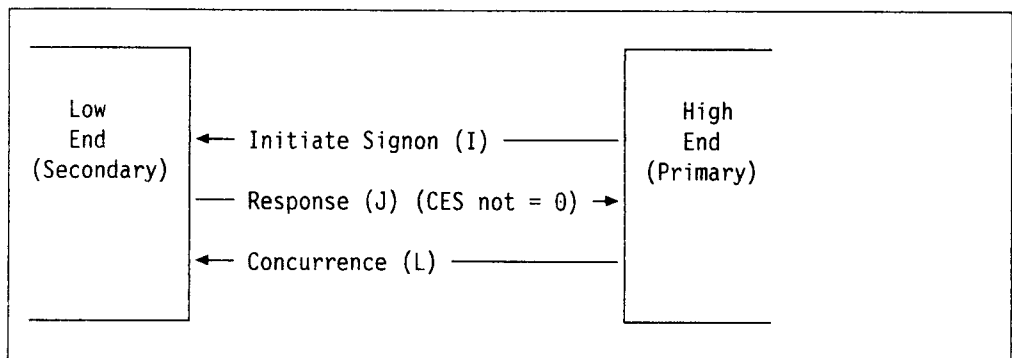


Figure 5-3. Normal Connection (High Initiates)

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 than the original. Because this does not represent a new connection, no CES is assigned and no distinction is made between low end and 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 5-4 illustrates secondary full trunk connection protocol.

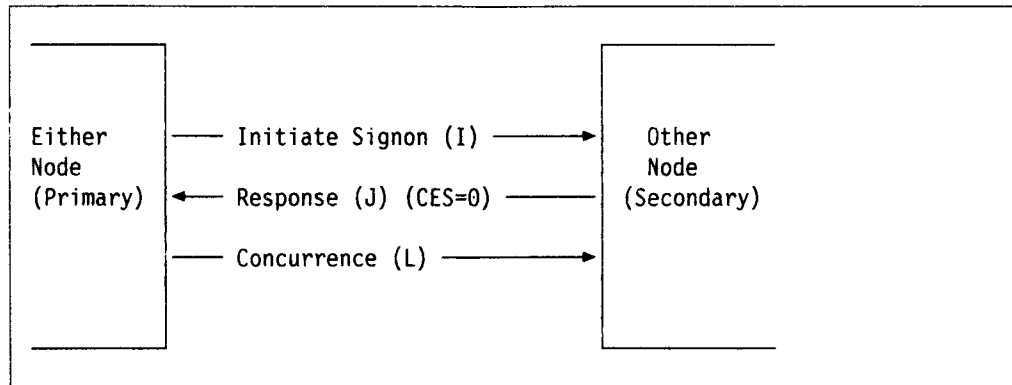


Figure 5-4. Secondary Trunk Connection

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 5-5 illustrates primary assignment reset protocol.

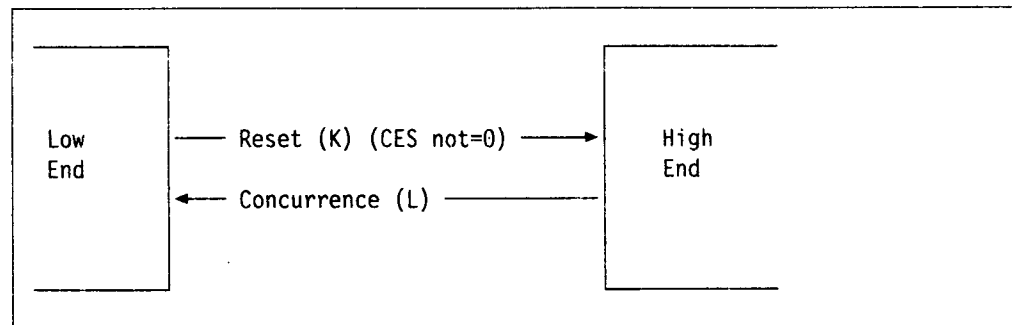


Figure 5-5. Trunk Reset Protocol

Signon Protocols Between A Node With NPM and A Node Without NPM

The protocols used to signon between a node using NPM and a node that doesn't use NPM are the same as those used to signon between nodes that do not use NPM. Refer to "Establishing a Connection Without NPM" on page 5-3. The node not using NPM discards all NPM records other than B, I or J.

All nodes that do not use the NPM must be defined to each node in the network. Figure 5-6 illustrates the protocols used to connect nodes using the nodal path manager.

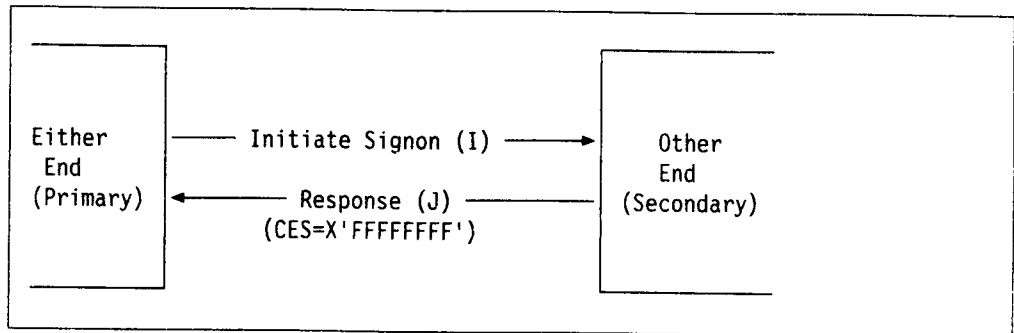


Figure 5-6. NPM Connection Protocols

Connection Status Information

Whenever a dynamic connection is agreed upon, each path manager involved sends an "add connection" control record to systems not involved in the connection over all other NJE lines. The add connection control record is used by receiving path managers to determine best paths to nodes within the network. Each NPM node forwards the add connection control records. If a connection is already known (CES indicates that the 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. Disconnections 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 subtract connection control records.

Add and subtract connection control records may be blocked in the buffer with reset and concurrence control records. This is 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 "Networking Connection Control Records (NCCR)" on page 2-67 for the format of the control records.

Disconnections

When an NJE line has disconnected, the path manager then clears its own reachable nodes in its tables, validates the queues, and notifies attached nodes of the disconnection.

Terminating a Session or Connection

Normal disconnection of an NJE node (using BSC or CTC) occurs when the initiating system sends a final signoff 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 occurs when the initiating system prepares 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 "Signoff" on page A-19 for RSCS-specific processing.

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

Chapter 6. Link Protocols

Bisynchronous Communications Lines	6-2
Initialization	6-2
Initialization Error Recovery	6-4
Error Recovery	6-5
Termination	6-6
Normal Sequences	6-6
V.27 Modem Contention Resolution	6-6
V.27 Signon by the Primary Node	6-7
V.27 Signon by the Secondary Node	6-7
Channel-to-Channel Adapter	6-8
Initialization	6-8
Error Recovery	6-9
Termination	6-10
Normal Sequences	6-10
PREPARE Mode Option: Suspend I/O (BSC and CTC Adapter)	6-11
Normal Sequences	6-11
Requesting I/O Suspension	6-12
Suspension and Resumption of I/O	6-12
BSC I/O Suspension	6-12
CTC Adapter I/O Suspension	6-14
Protocols Common to BSC and CTC Adapters	6-16
BSC Error Protocols	6-16
CTC Adapter Error Protocols	6-21
Wait-a-Bit and PREPARE	6-24
Systems Network Architecture - LU Type 0	6-24
Overview	6-24
Session Initialization	6-25
BIND RU Format	6-26
Function Management Headers	6-27
Compaction Table Format	6-31
Session Termination	6-33
Normal Termination with Quiesce	6-34
Normal Termination - Immediate	6-34
Error Recovery	6-35
ACF/VTAM Considerations	6-36
Application Exits	6-36
Data Flow Commands	6-37

Bisynchronous Communications Lines

Initialization

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. If the nodes cannot agree on which node will be the primary and which will be the secondary, then the nodes are in contention.

BSC connections are initialized in one of the following two ways:

1. Figure 6-1 illustrates when there is not any contention between two nodes when signing on.

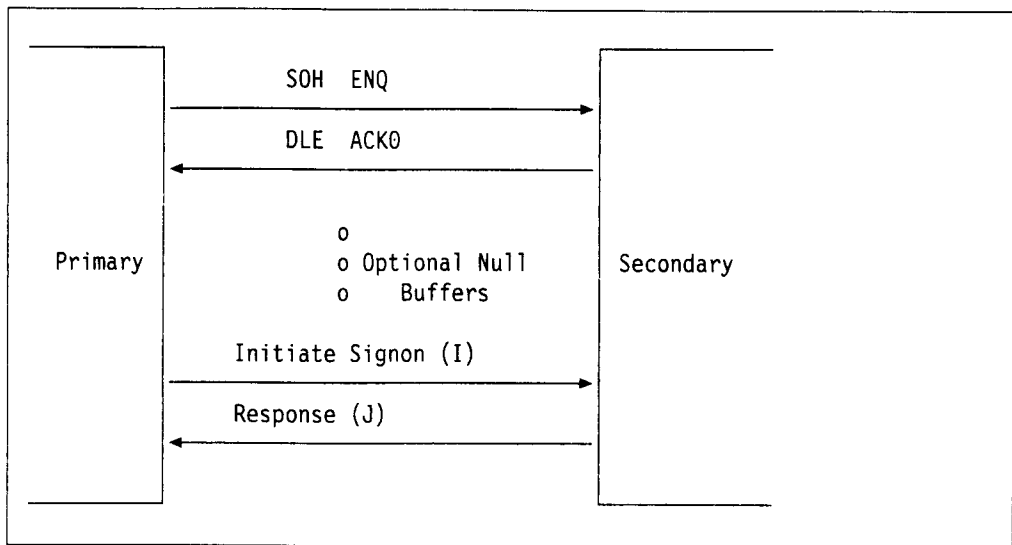


Figure 6-1. BSC Initialization

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 provide the facility to be pre-defined as a secondary node.

2. 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 ENQs in one of the two directly-connected communication controllers (CCs).

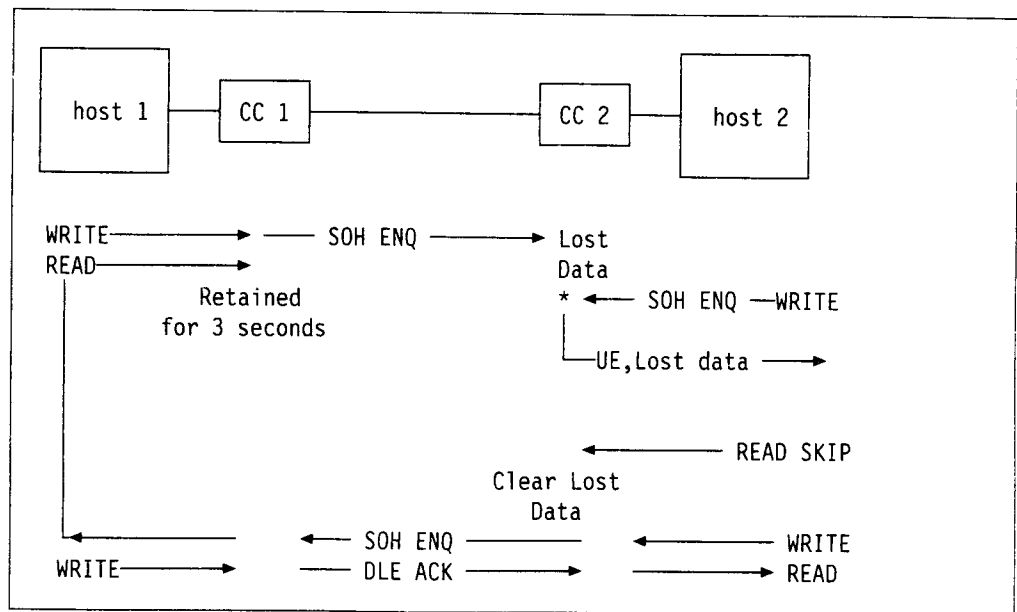


Figure 6-2. BSC Signon Contention

If inbound data arrives at a communication controller (CC) and a READ is not active, 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 that 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 ENQs 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 pre-defined.

Figure 6-3 shows the channel commands used in BSC initialization.

<i>Figure 6-3. BSC Channel Commands</i>		
Opcode	Definition	Comments
X'2F'	Disable	This CCW causes the adapter to be reset. If this is a switched communications line, the circuit will be terminated.
X'23'	Set Mode	Conditions for this 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 normal 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	Reads data from adapter.
X'06'	Prepare	Causes the adapter to wait for synchronization.
X'29'	Dial	Enables the adapter and automatically issues a call to the receiver.

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 successful 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 6-4.

<i>Figure 6-4. BSC Initialization Error Recovery</i>		
Unrecognized data		Rewrite SOH ENQ
Device Status	CCW	Recovery Action
Unit	Write	Read Skip - Rewrite SOH ENQ
Exception	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

Error Recovery

Post-initialization error recovery 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 6-5.

SOH ENQ		Restart line
NAK		Retransmit last non-NAK transmission
Unrecognized data		Write NAK
Device Status	CCW	Recovery Action
Unit	Write	Read Skip - Write NAK
Exception	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 6-5. BSC Error Recovery

Termination

Normal termination occurs when a signoff record is sent. (See "Terminating a Session or Connection" on page 5-8) There is no special BSC character sequence sent for termination and no response is expected to the signoff record. If a read is chained to the write of the signoff, the read will time out. After the signoff record is written, a disable should be issued to disable the communications adapter.

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.

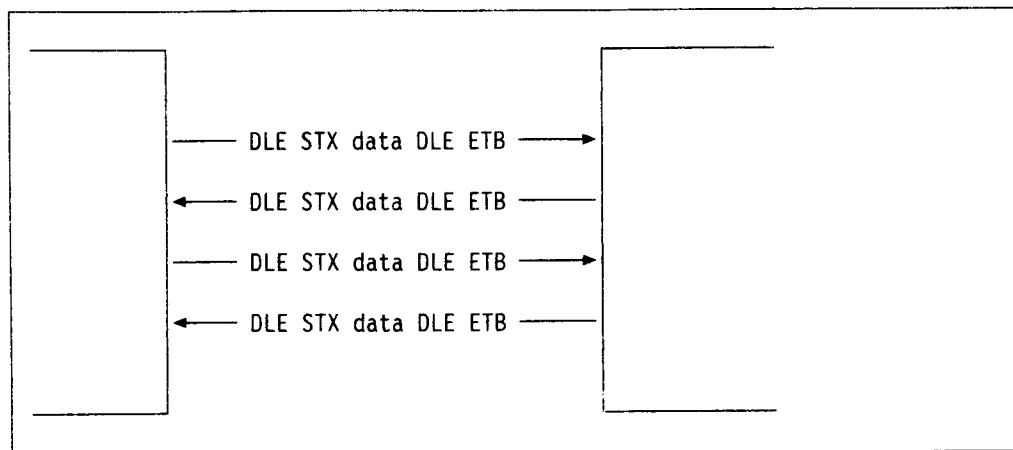


Figure 6-6. BSC Normal Communication

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

For V.27 signon, a given node 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.

V.27 Signon by the Primary Node

This scenario assumes that the node attempting signon has been identified as the primary node.

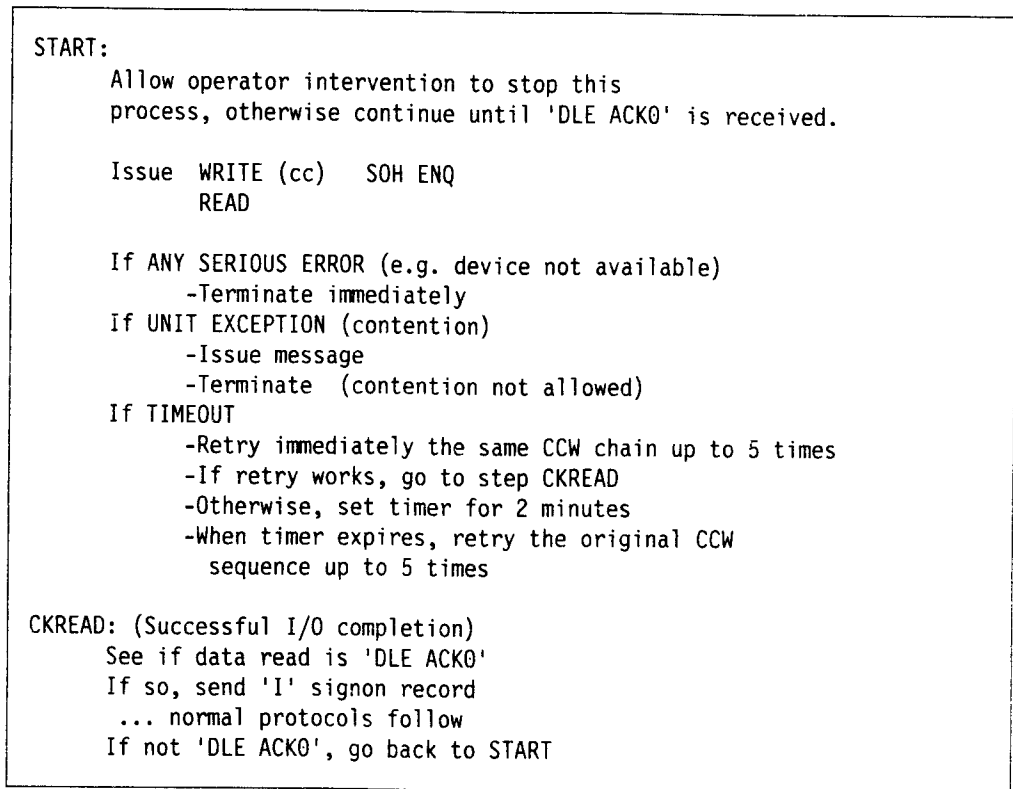


Figure 6-7. V.27 Modem Protocols: Primary Node Signon

V.27 Signon by the Secondary Node

This scenario assumes that the node attempting signon has been identified as the secondary node.

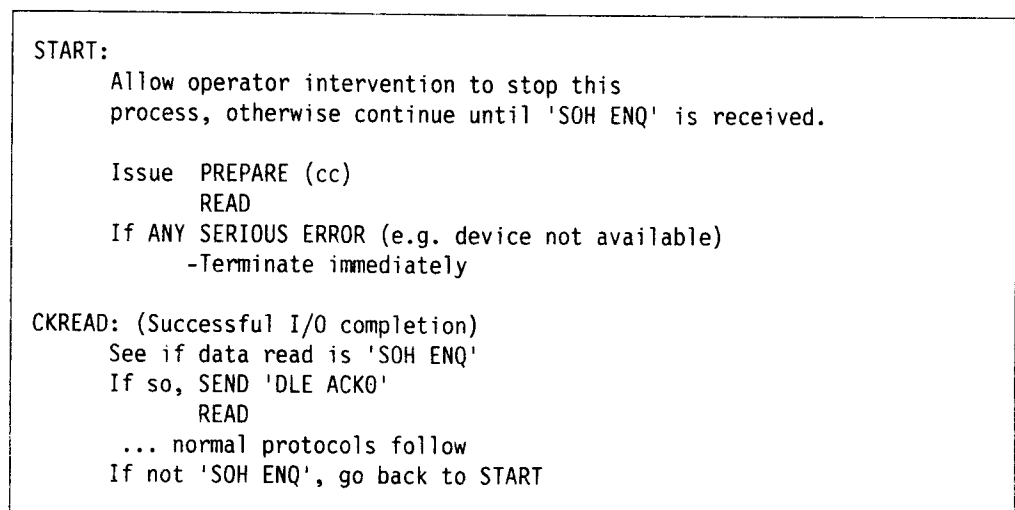


Figure 6-8. V.27 Modem Protocols: Secondary Node Signon

Channel-to-Channel Adapter

NJE also supports communication by a channel-to-channel (CTC) adapter as a communication line. The support is for 360-mode operations over block multiplexor channels.

Note: POWER supports only virtual channel-to-channel adapters.

To properly operate as a NJE device, the UCWs for the CTC adapter should be plugged non-shared. (This is a hardware option on block multiplexor channels.) The UCW is used to store the address for performing the channel reconnection. If the UCW for the CTC adapter is not plugged non-shared, the CTC adapter 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.

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 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 command 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 CTC adapter to respond.

When side B starts, it issues a sense device command, detects the control state (X'07') on side A, and writes a SYN NAK (X'323D'). This allows side A to become the primary, which 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 "CTC Adapter Initialization" on page A-16 for JES3-specific processing; see "CTC Adapter Initialization" on page A-28 for RSCS-specific processing.)

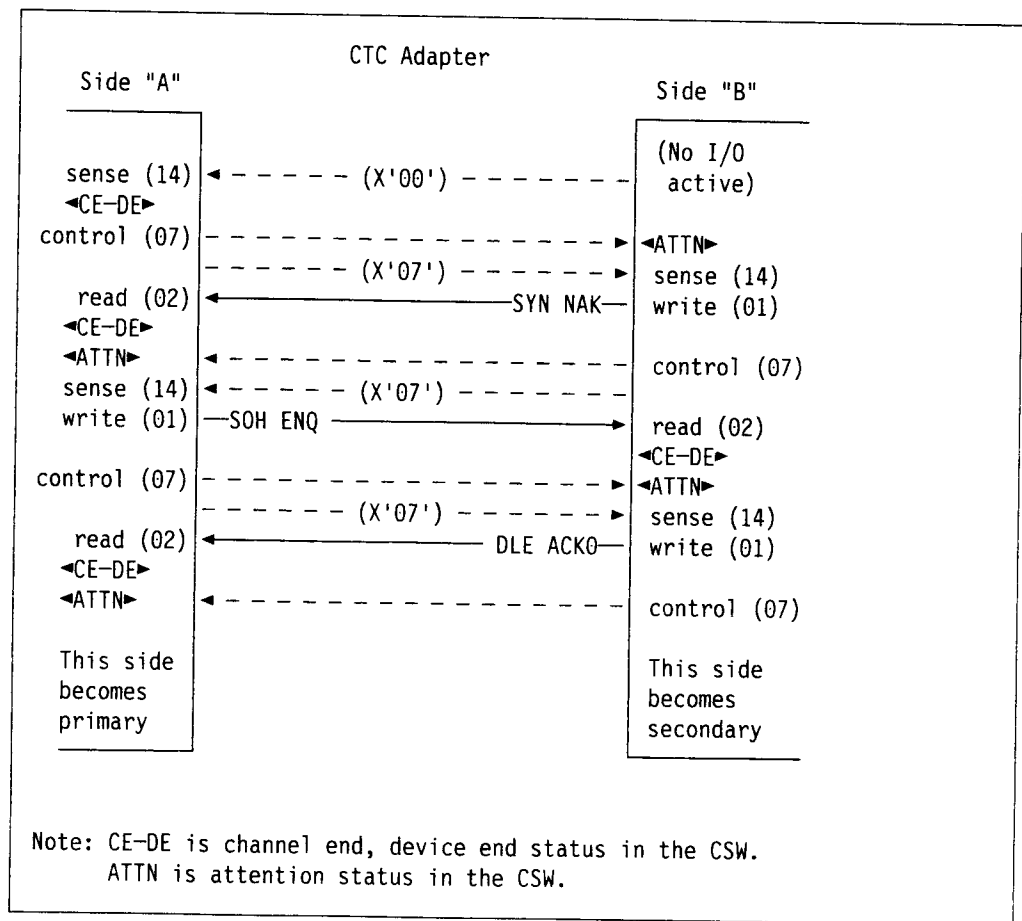


Figure 6-9. CTC Adapter Initialization

Error Recovery

No error recovery is used on the CTC adapter. When running in compatibility mode, the CTC adapter 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.

Prior to the advent of channels that buffer status, attention and busy could only occur if the two systems had lost synchronization. A non-shared 30xx subchannel may cause an attention and/or busy status. This is because the subchannel has already buffered the attention from the other side's I/O request and intends to present it. In channels that do not buffer status, the I/O would have been sent to the CTC adapter, which 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 "CTC Adapter Attention and Busy Status" on page A-8 and "CTC Adapter Attention and Busy Status" on page A-16 for system-specific handling of this condition.)

Termination

Normal termination occurs when a signoff record is sent. (see "Terminating a Session or Connection" on page 5-8) There is no special character sequence sent for termination and no response is expected after a signoff record is sent. In this case, the standard CTC program should terminate with the write CCW. (See "CTC Adapter Termination Deviations" on page A-16 for JES3 deviations.)

Normal Sequences

Instead of using CCWs designed for BSC line transmissions, the channel-to-channel adapter sequences are as those shown in Figure 6-10.

Opcode	Definition	Comments
X'14'	Sense	This CCW will cause other end of adapter control 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 the other end of the adapter issues a SENSE. A block multiplex channel is freed for other programs.
X'02'	Read	Reads data from the adapter.

Figure 6-10. CTC Adapter Program

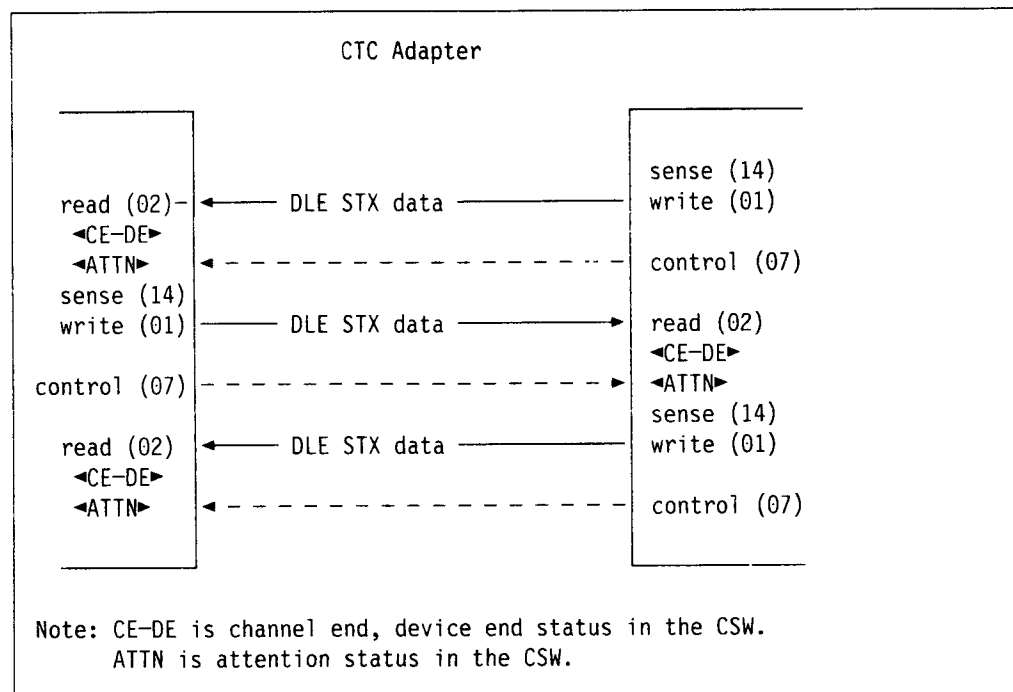


Figure 6-11. CTC Adapter Normal Sequences

PREPARE Mode Option: Suspend I/O (BSC and CTC Adapter)

This option (implemented only by RSCS Version 2) 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 CTC adapters, 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 "Signon Concurrence Feature" on page 5-2 for the description of signon concurrence.) This causes no problems for back-level systems because this has already been 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 mode 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.

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 6-12. (The CTC adapter normal CCW chain is similar, except that the read CCW is preceded by a control command and the write by a sense.)

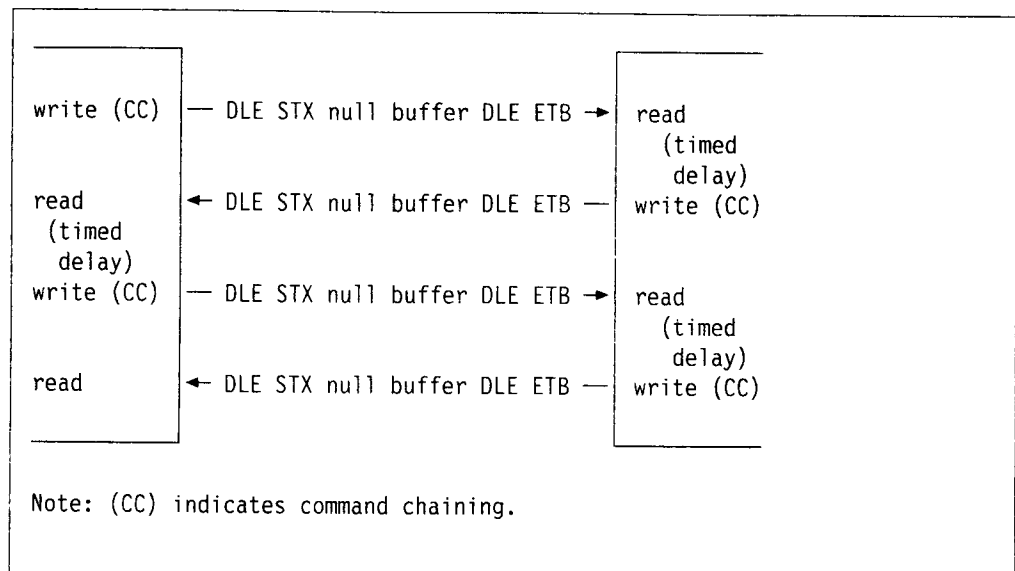


Figure 6-12. BSC Normal Communication

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 (side A) may decide to suspend I/O activity. It does this by writing a null buffer to the other side (side 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 CTC adapters or BSC. Thus, the write is command chained either directly (BSC) or indirectly (through a control for CTC adapters) to a read. If side 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 (CTC adapter). If it does not want to suspend I/O, side 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: The two null buffers with DLE ETX must be exchanged without an intervening buffer with DLE ETB in order for PREPARE mode 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.

Suspension and Resumption of I/O

How communication is actually suspended and resumed with PREPARE mode differs for BSC and CTC adapters. This is described in the following separate sections.

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 halt device (HDV) or halt I/O (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 resume normal processing once data has been received. This sequence is shown in Figure 6-13 on page 6-13.

If both systems attempt to initiate transmission at the same time, one of the DLE ENQs 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 "BSC Error Protocols" on page 6-16) to force resynchronization.

Note: During resumption, all null buffers are terminated with the normal DLE ETB sequence.

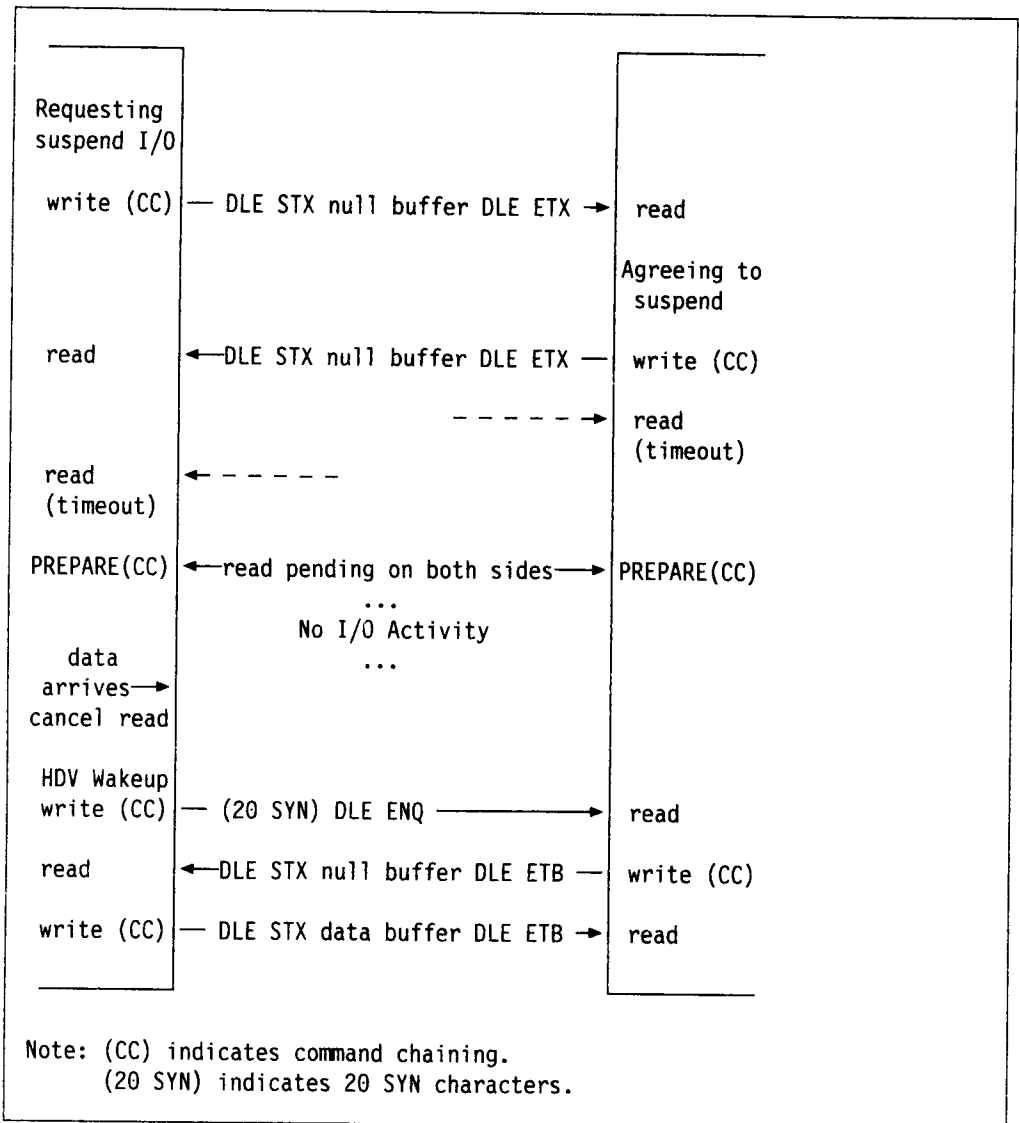


Figure 6-13. BSC Communication in PREPARE Mode

CTC Adapter I/O Suspension

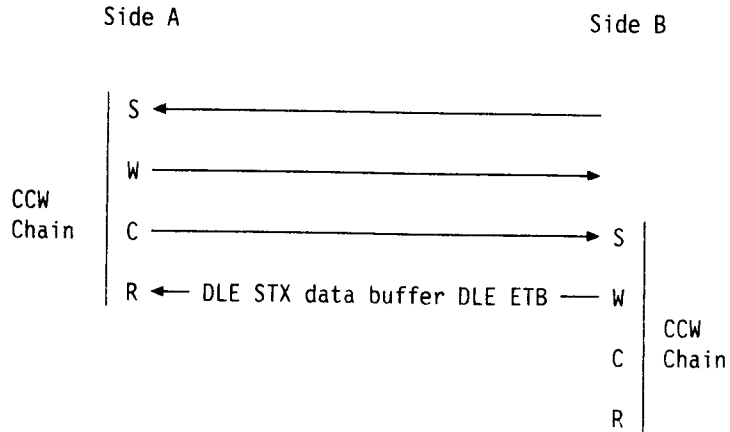
During I/O suspension on a CTC adapter, no I/O sequences are issued so that the adapter is left with no I/O in progress during idle periods. A control CCW, which signals attention, indicates data transmission is to be resumed. The procedure, illustrated in Figure 6-14 on page 6-16, is as follows:

- I/O suspension is requested and agreed upon as described in "Requesting I/O Suspension" on page 6-12.
- Both sides then wait without any I/O outstanding.
- If either side wants to resume, it must first issue a control CCW which will signal an attention to the other side (B, for this example). The control will be command chained to a read.
- When the attention interrupt is detected by side B, it will issue its normal CTC adapter CCW chain (sense, write (null buffer), control, read).
- The sense causes the control issued by side A to complete, and the null buffer is written by side B and read by side A.
- Side A, which had data to send, can send the data by issuing the normal CCW chain as soon as the first "wake up" 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 down during the PREPARE, the "wake up" 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, an HDV would be issued and the link taken down.

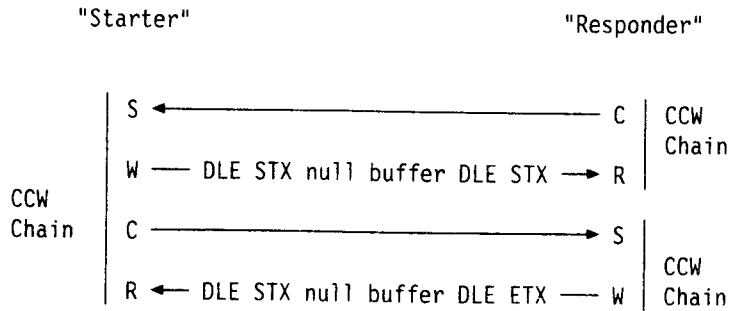
The following flows show the procedure graphically. In the diagrams:

S = Sense CCW
C = Control CCW
R = Read CCW
W = Write CCW

Normal Sequence (for CTC adapter)

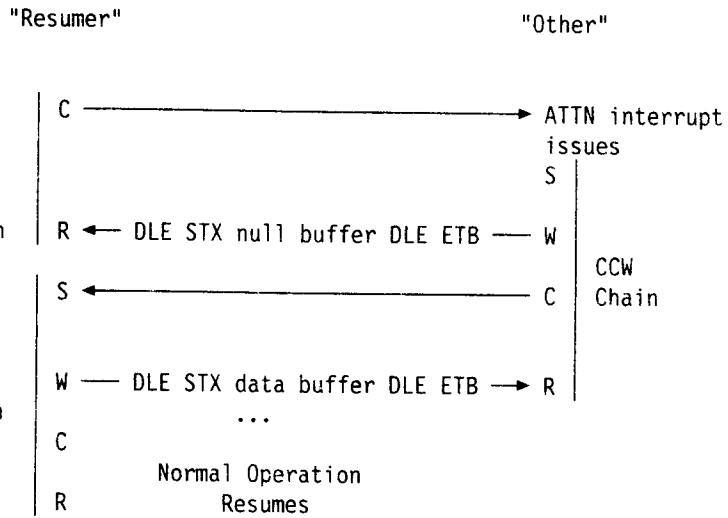


PREPARE Initiation (for CTC adapter)



No I/O Active

I/O Resumption



Protocols Common to BSC and CTC Adapters

During the I/O suspension period, either side may optionally set a timer for some reasonable interval (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, the link is taken down by normal error recovery procedures.

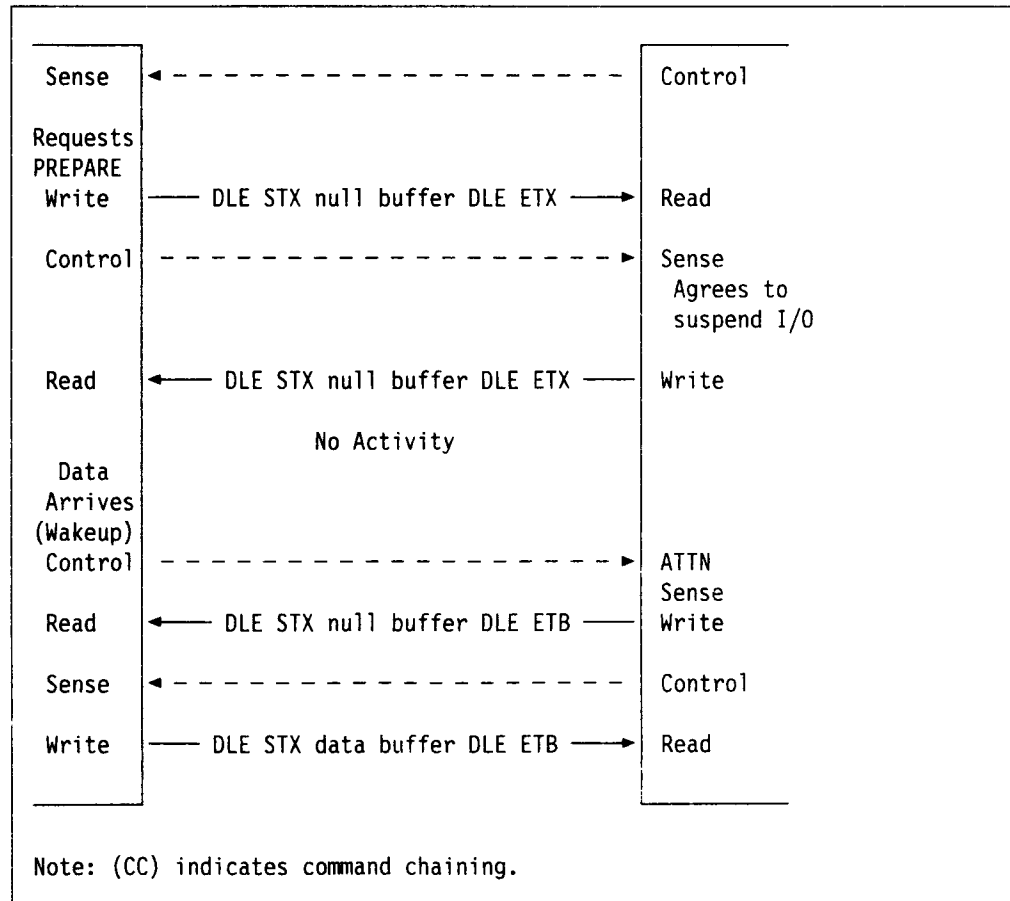


Figure 6-14. CTC Adapter Communication in PREPARE Mode

BSC Error Protocols

The symbols used in the flows that follow 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 receives a negative acknowledge, 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.

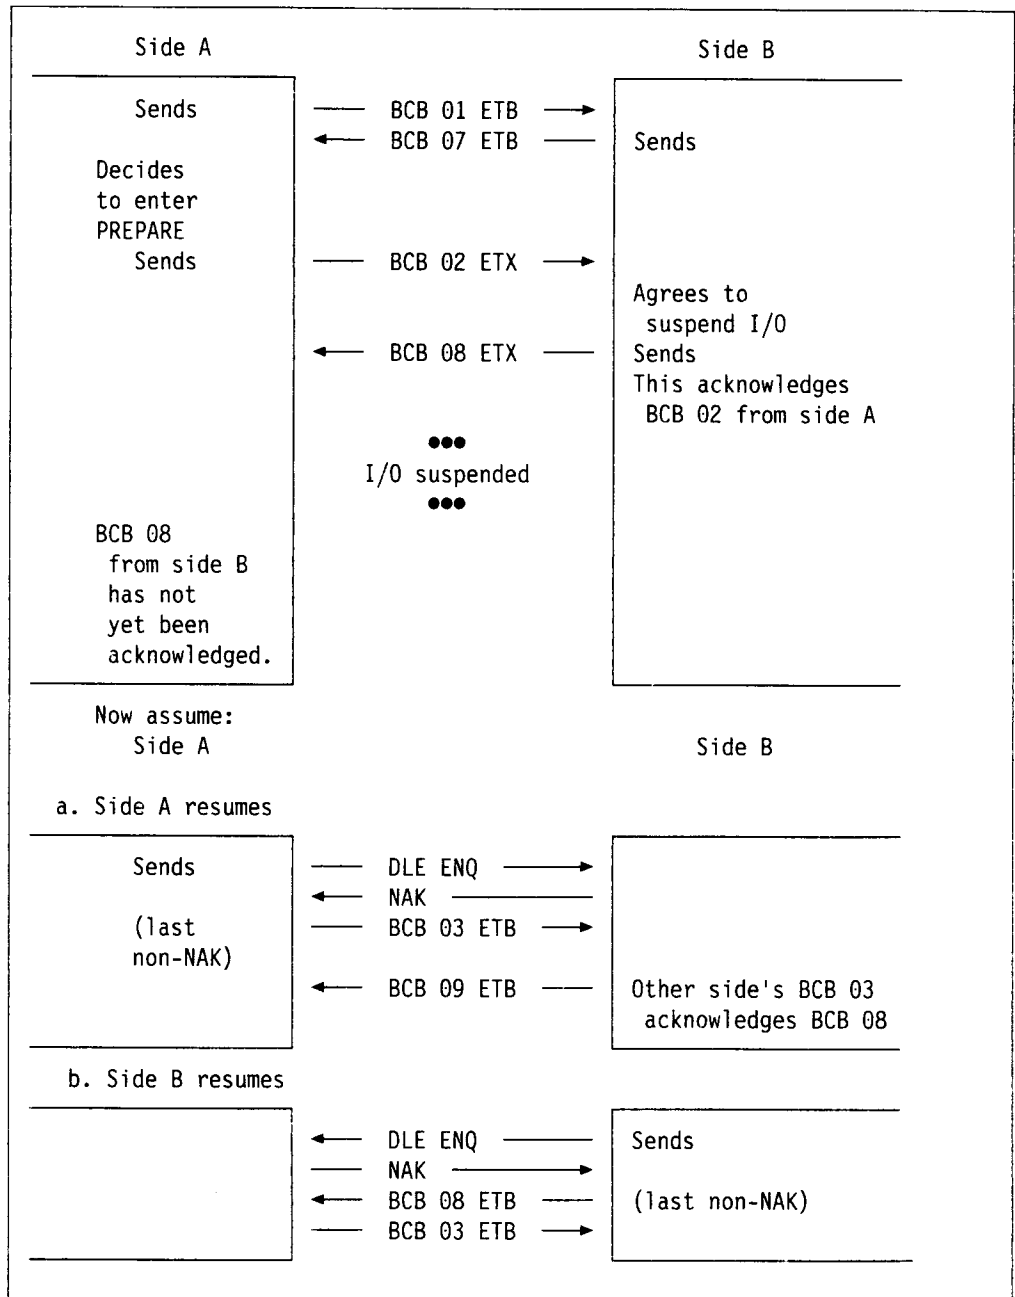


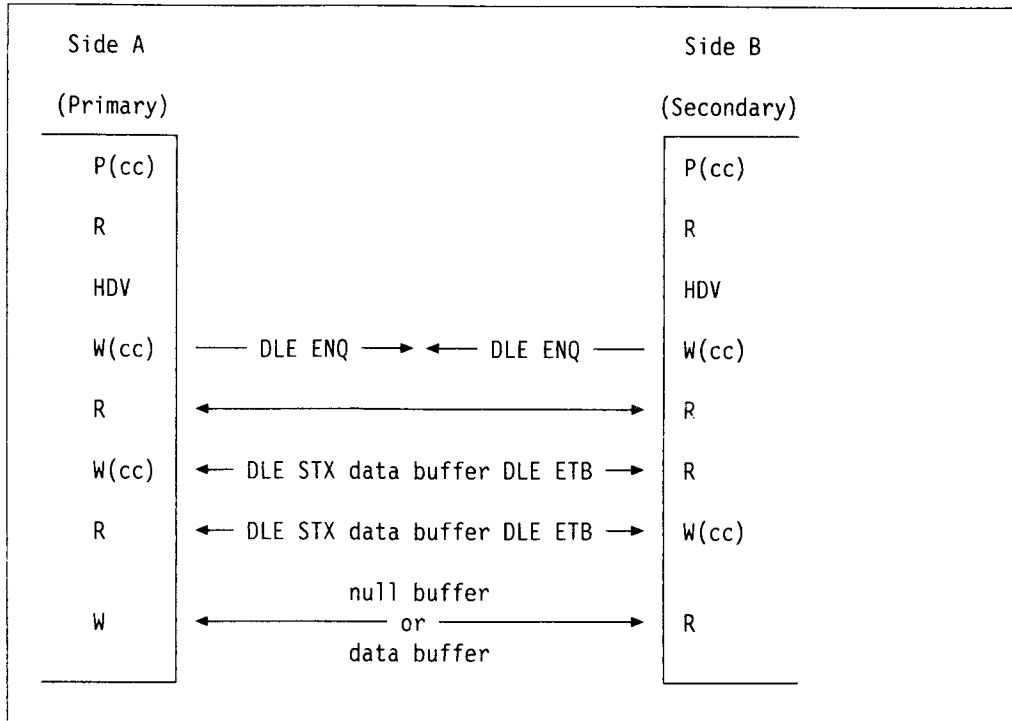
Figure 6-15. BSC Error Protocols

2. In the case of a wake up 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.

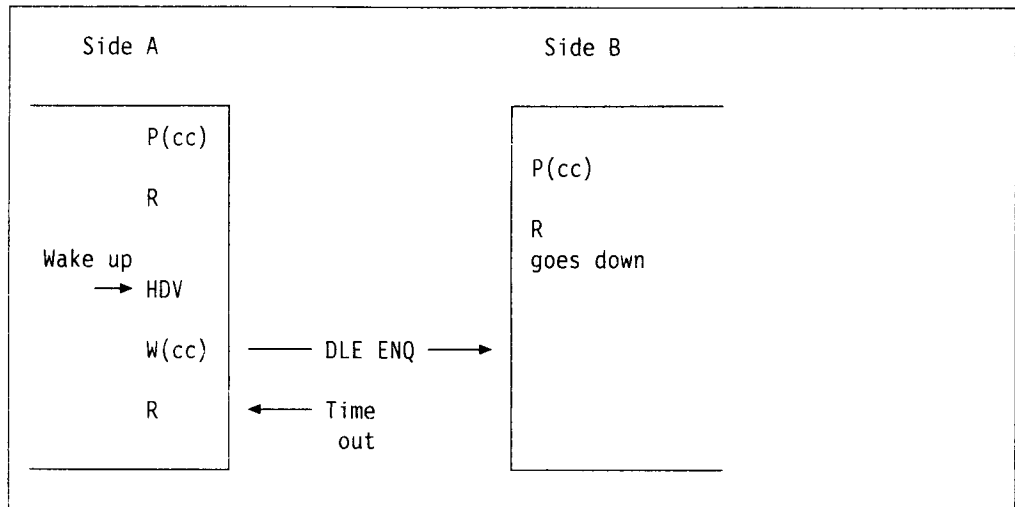
If one side receives a unit exception when it tries to write the DLE ENQ, that side uses the primary/secondary relationship in the same way as above after issuing a read skip to clear the contention.

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:

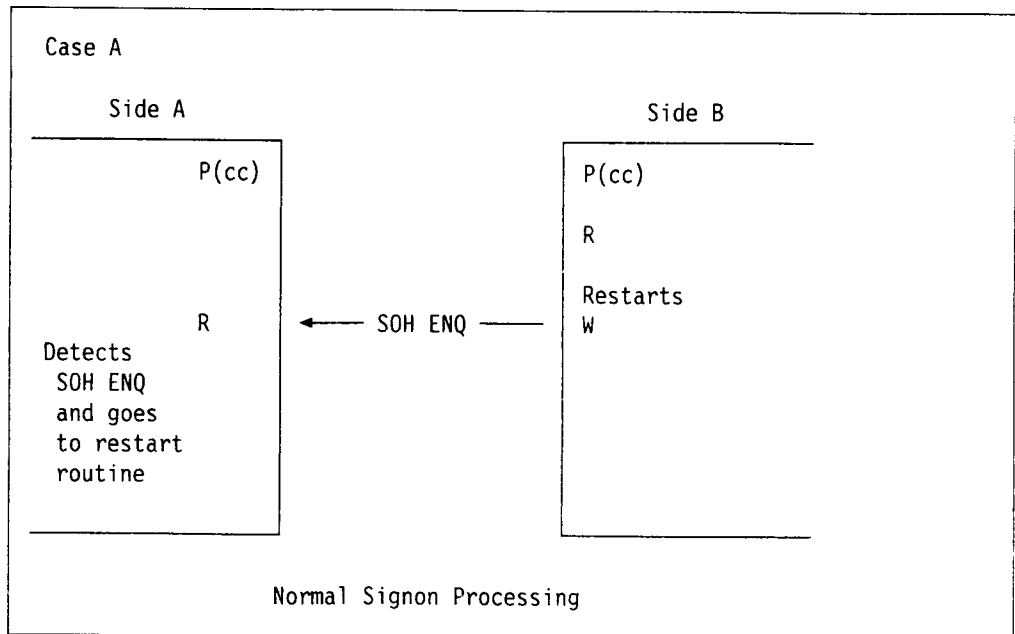


3. If one side goes down while I/O is suspended, the other side will detect it at resume time by a timeout as follows:

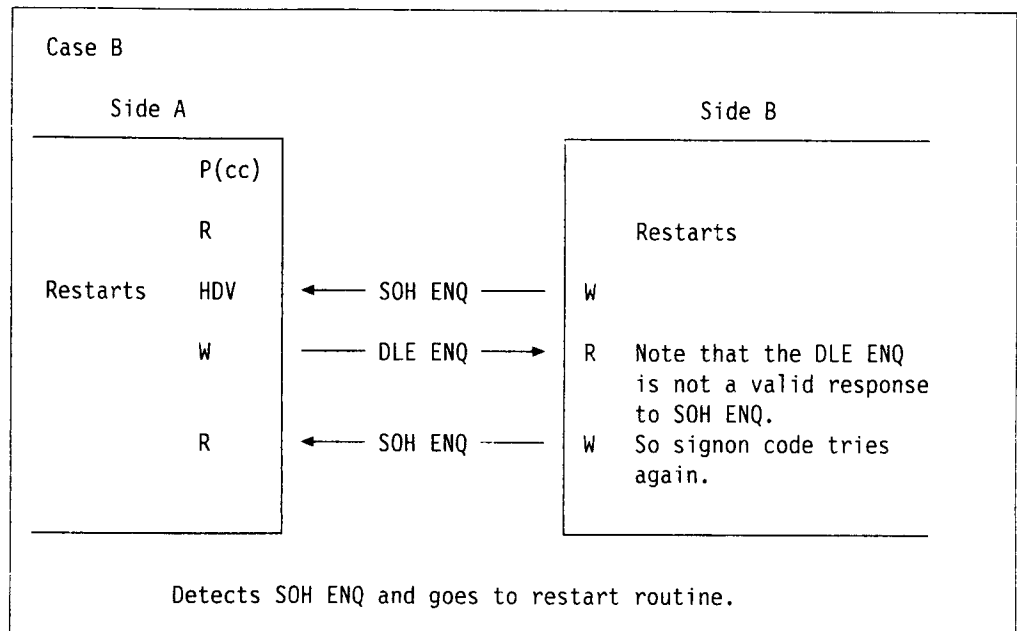


When the read times out in this case, normal time-out processing should be used. This includes sending a NAK.

4. If one side restarts while I/O is suspended, the following scenario will occur:

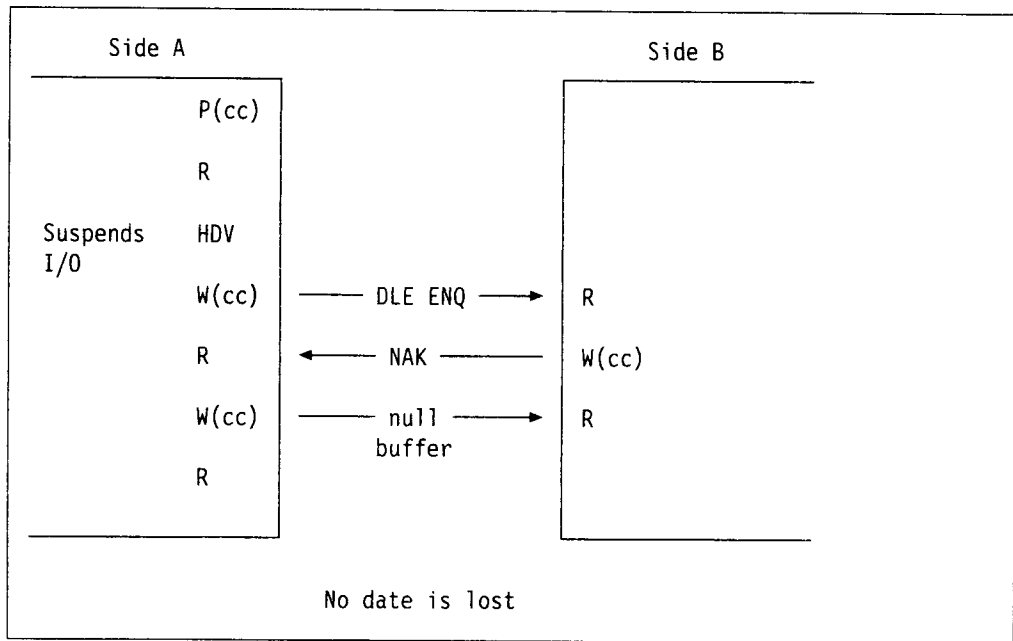


If the restart occurs at the same time as a resume, the following can happen:

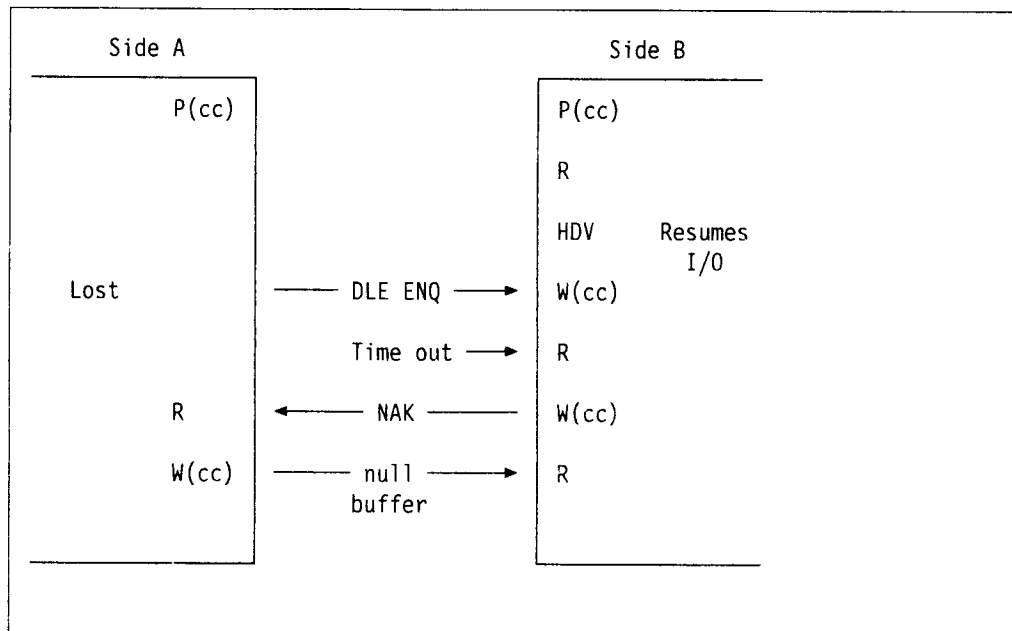


Note: If both sides try to restart while I/O is suspended, the signon contention protocol will resolve contention problems.

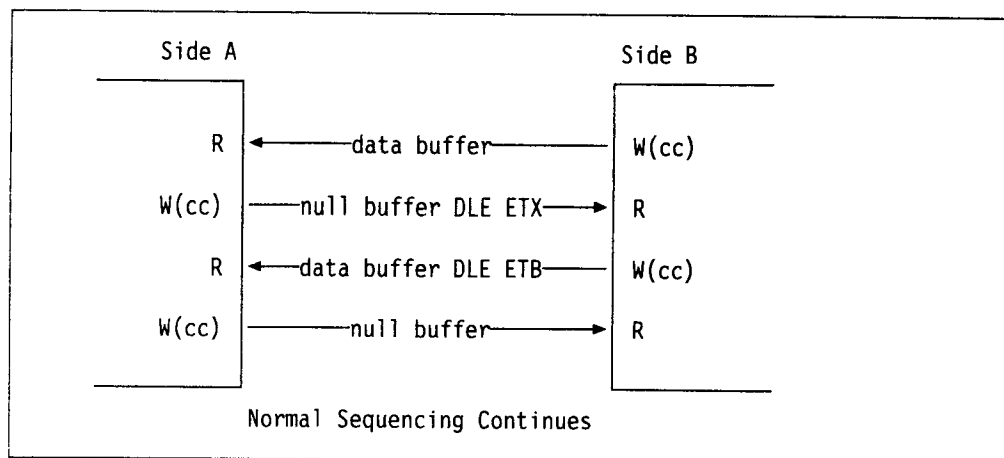
5. A system that has not agreed to use PREPARE at signon time should never receive DLE ENQ. However, if such a sequence is ever received unexpectedly, the proper procedure is to respond to it with a NAK.



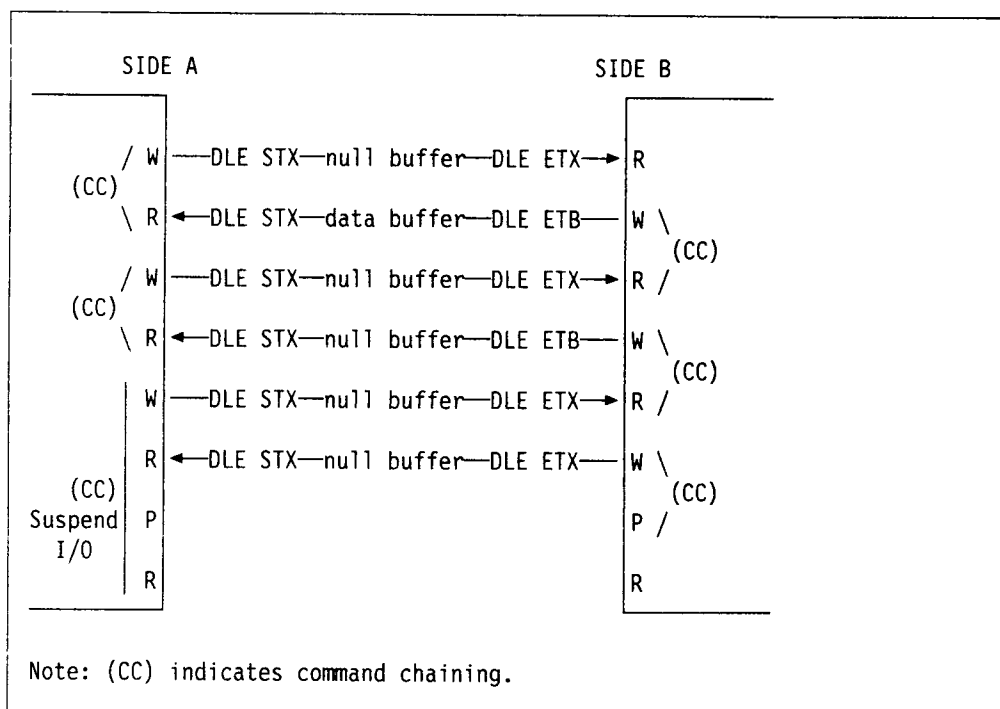
6. If the resuming side sends DLE ENQ and the sequence is lost (due to a hardware error), a timeout will occur and that side should send a NAK (following normal timeout recovery).



7. Note that the only valid sequences read by the READ 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 following flow shows what happens when side A requests I/O suspension and side B does not want to enter this state because it has something to send.



9. Either side may continue to request suspension, even if the other side does not agree to it.



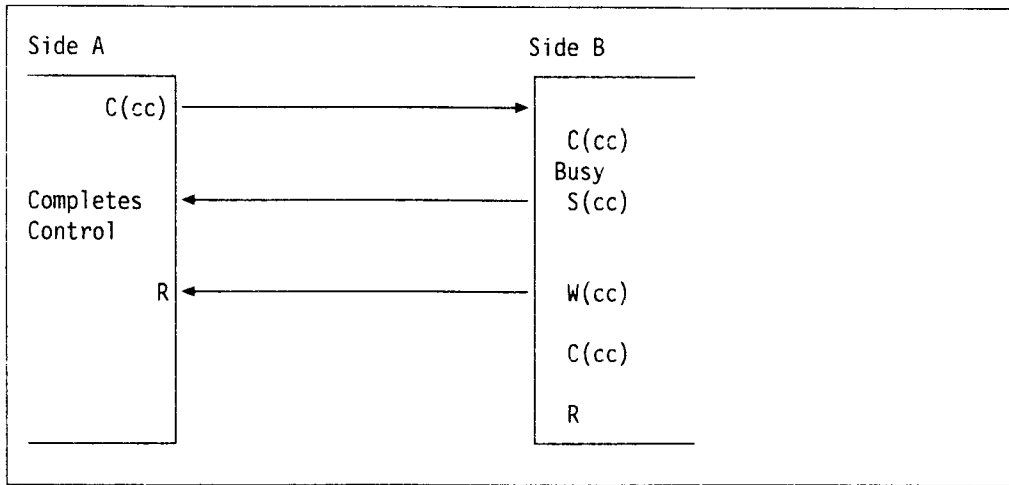
CTC Adapter Error Protocols

The symbols used in the following flows are:


- R = Read
- P = PREPARE
- W = Write
- S = Sense
- C = Control
- TO = Time Out
- cc = Command Chain

1. A NAK is never sent on a CTC adapter; if a NAK is ever received, the link should be taken down.

2. The hardware prevents any contention during I/O resumption.



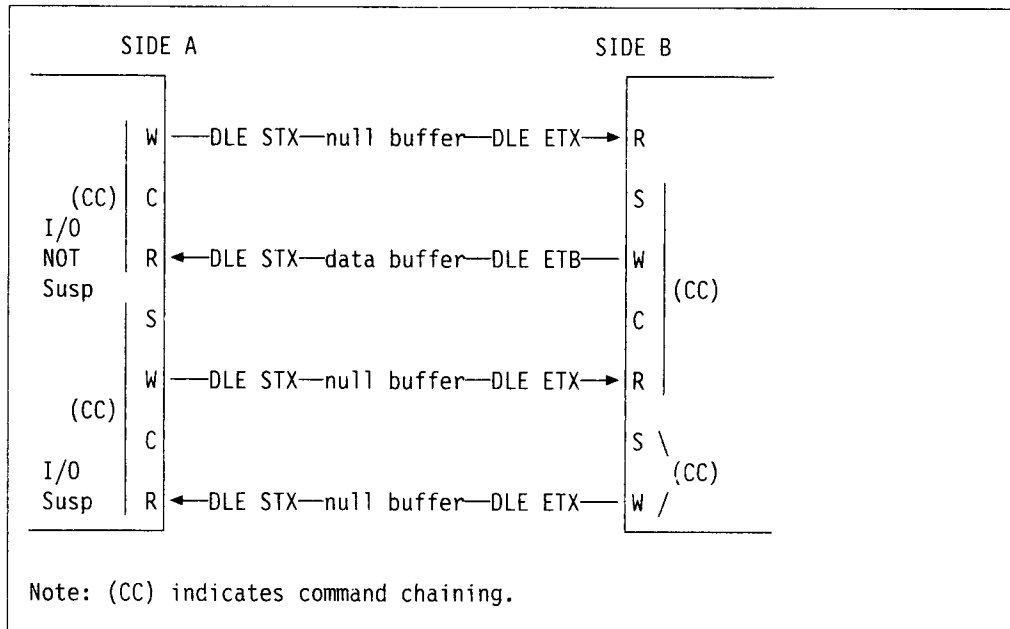
3. If one side goes down while I/O is suspended, the other must detect this during the resume as follows:



4. Restart

If side B responds with null-buffer DLE ETX immediately, suspension will then take place. However, if side B responds with a normal buffer, normal I/O continues.

- Side A may go on writing null buffer DLE ETX if it still wants to suspend the I/O. Suspension will not take place until side B responds immediately with the null buffer DLE ETX.



- To prevent the possible loss of data in unexpected conditions, all reads should be for a full buffer. 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.

Wait-a-Bit and PREPARE

- PREPARE mode can never be used during any wait-a-bit condition (by definition) because no files are active when a PREPARE is initiated.
- PREPARE mode 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.

Systems Network Architecture - LU Type 0

Overview

NJE defines its own type of LU, LUTYPE 0, for all its application-to-application communication in an SNA environment. LUTYPE 0 is a full duplex protocol that 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.

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
		B'bbbbbb'	Reserved
	1		Reserved
	2-7		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
		B'0'	Reserved
	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	0	B'0'	PS profile 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'0000000'	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 "SNA Support" on page A-7 for additional JES2 BIND information.
- See "Data Flow Control" on page A-36 for POWER data flow control.
- See "Bind Parameters" on page A-27 for additional RSCS BIND information.

Function Management Headers

Function management headers (FMHs) are used to control the data presentation for the session. FMHs are sent as only-in-chain elements, with a definite response required. They are not compressed or 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 or not *and* whether or not the sender elects to use compaction.

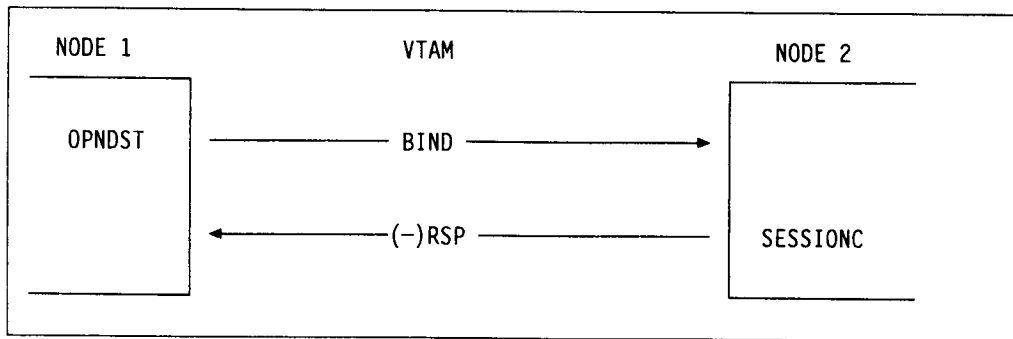


Figure 6-17. SNA Session Initialization Error Recovery

See "SESSIONNC Usage" on page A-7 for JES2-specific parameters.

BIND RU Format

The BIND area is sent from the primary LU to the secondary LU to activate a SNA session. The BIND contains the parameters that 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
		B'bbbbbb'	Reserved
	1		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	0		PS profile
			Basic format
		B'0000000'	LUTYPE 0
15-25			No protocols specified
26	0-1	B'00'	Reserved-encryption not used by NJE
	2-7	B'0000000'	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 "SNA Support" on page A-7 for additional JES2 BIND information.
- See "Data Flow Control" on page A-36 for POWER data flow control.
- See "Bind Parameters" on page A-27 for additional RSCS BIND information.

Function Management Headers

Function management headers (FMHs) are used to control the data presentation for the session. FMHs are sent as only-in-chain elements, with a definite response required. They are not compressed or 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 or not *and* whether or not 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 is not supported as a receiver, then side B must not send an FMH3.
2. Similarly (and independently), if side B indicates compaction is 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 bit must not be set, and need not be checked. See "Functional Management Headers" on page A-36 for POWER-specific information; see "Function Management Headers" on page A-27 for RSCS-specific information.

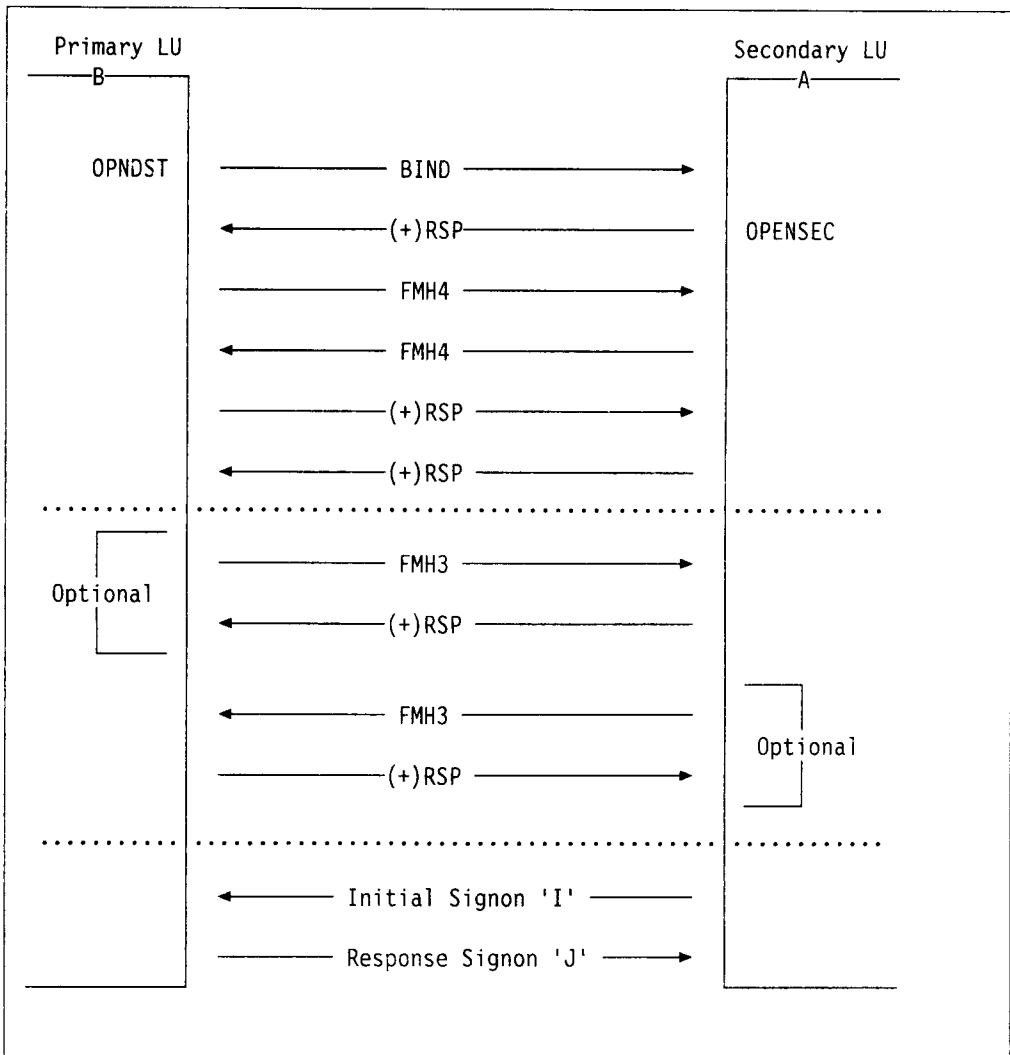


Figure 6-18. SNA FMH and Signon Flows

Figure 6-18 shows 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 the signon I record). In fact, the figure shows an example where the roles are reversed.

FMH4

FMH4 is a non-architected SNA header 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 "Networking Connection Control Records (NCCR)" on page 2-67 for the preferred method of determining the transmission buffer size using a field in the signon record.

FMH4 indicates whether a compaction table and signon records can be received by the node that 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 follows.

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. The FMH4 format is shown in Figure 6-19.

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' to 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
7	3-7	B'00000'	Reserved
	0	B'1'	RID format 1
	1-7	B'0000000'	Reserved

Figure 6-19. FMH4 Format

See "Function Management Headers" on page A-8 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 6-31 for a description of compaction characters.) The format for FMH3 is shown in Figure 6-20.

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 6-20. FMH3 Format

See "Function Management Headers" on page A-8 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 (greater than 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 6-31 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 (length less than 300, bits set wrong, or compaction table format wrong, etc.)
Action: Send negative RSP and UNBIND.
2. Receiving an FMH4 or FMH3 when not expected or not allowed by NJE.
Action: Send negative RSP and UNBIND.
3. Multiple FMHs in buffer.
Action: Product choice. Either ignore the extra data or send a negative 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 6-21 on page 6-32. 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'
'abcdefghijklmnopqrstuvwxyz'
```

Space in the upper left ($m \times m$) 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 16×16 table) is 256 minus $(m \times (m + 1))$. The smallest value of m for that 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 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 decompact, 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 Figure 6-21 on page 6-32 is:

BYTE	VALUE	MEANING
----	-----	-----
0	X'5B'	Length of header: $4 + 256 - (m \times m)$, 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
													v	w	y
													s	t	u
													o	p	r
													l	m	n
													h	i	k
													e	f	g
													b	c	d
													8	9	a
													5	6	7
													2	3	4
													Z	0	1
													W	X	Y
													Q	04	V
?	\	:	#	@	'	=	"	B	C	F	H	J	K	M	P
(+		&	1E	\$	*)	;	0C	-	/	,	6C	_	▶
	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 6-21. 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 uncompact form. Figure 6-22 on page 6-33 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
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	93	Bits 2-7 ('010101'B) indicate 21 bytes								
		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 decompact 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 6-22. 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.

Session Termination

The following describes two types of termination:

- Normal termination with quiesce.
- Normal termination - immediate.

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 primary LU (PLU) to secondary LU (SLU). Rather, it is simulated to both the PLU and the SLU and the NS exits at both ends are driven.

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 6-23. SNA Session Termination - Initiated by Primary Node

If the secondary LU wishes to terminate the session, it must issue RSHUTD. After all data activity has stopped, the primary LU 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.

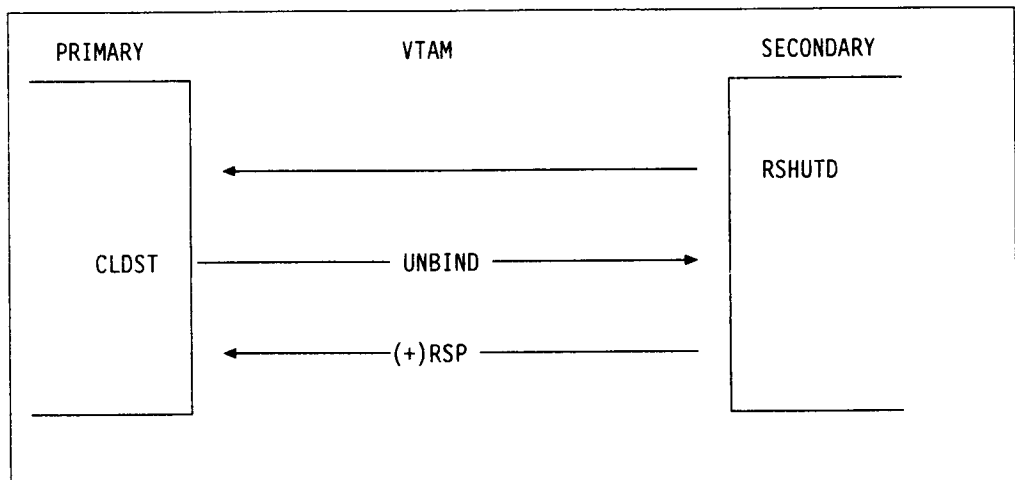


Figure 6-24. SNA Session Termination - Initiated by Secondary Node

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 an 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 LU initiates normal session termination.

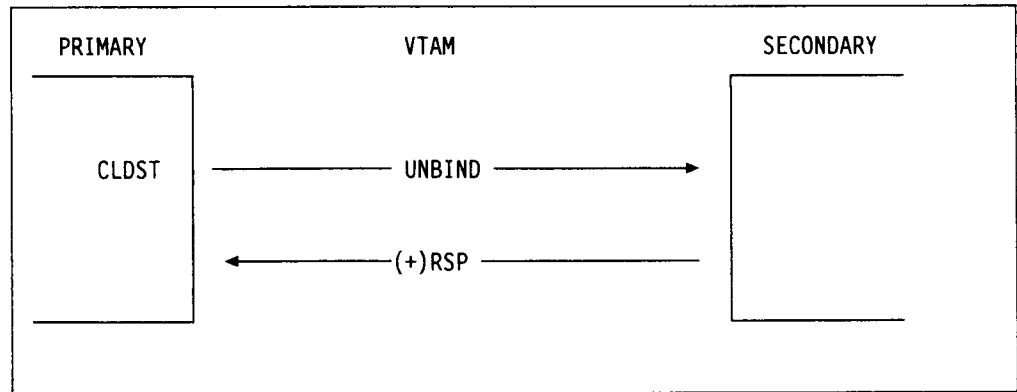


Figure 6-25. SNA Immediate Session Termination - Initiated by Primary Node

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.

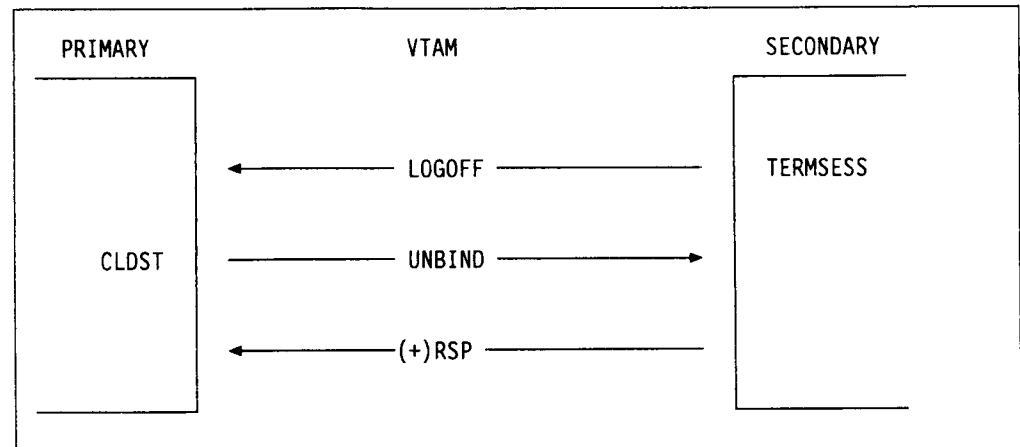


Figure 6-26. SNA Immediate Session Termination - Initiated by Secondary Node

See "Termination" on page A-8 for JES2 specific information.

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 that 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.

ACF/VTAM Considerations

Application exits and data flow commands are defined to VTAM and can be used as described below.

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 TERMSSESS or CLSDST).

NSEXIT

VTAM calls this exit with a CLEANUP RU if the session has been lost.

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 A. System-Dependent Considerations

This section contains system-dependent considerations and deviations from the protocols for each of the networking facilities.

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.

Figure A-1. Various NJE Features Supported by Networking Facilities

Feature	JES2	JES3	RSCS	POWER	BDT
BSC Communications	Yes	Yes	Yes	No	(J3)
SNA Communications	Yes	No(B)	Yes(*)	Yes	Yes
CTC Communications	Yes	Yes	Yes	No	No(J3)
Network Path Manager	Yes	No	No	No	No
Formatted Commands	Yes	(A)	(A)	No	N/A (J3)
Data Compaction	Yes	No	No	(A)	No
Output Fanout	Yes	(A)	Yes(*)	(A)	N/A (J3)
Spanned Headers	Yes	Yes	Yes	Yes	Yes
Multiple Streams	7	2(**)	Yes(*)	Yes	28
Wait-A-Bit All Streams	Yes	(A)	Yes	Yes	N/A
Null Buffers vs. DLE ACK0	(A)	(A)	Yes(*)	Yes	N/A
PREPARE Mode - Suspend I/O	No	No	Yes(*)	No	N/A
V.27 Contention Resolution	No	No	Yes(*)	No	N/A
Enhanced Termination - Quiesce	No	No	No	No	Yes
Signon Concurrence	(A)	(A)	Yes(*)	Yes	Yes

Key:

- YES = Supports the Feature
- NO = Does Not Support The Feature
- N/A = Not Applicable
- (A) = Accepts, but does not Send.
- (B) = Function Supported by BDT
- (J3) = Function Supported by JES3
- (*) = Provided with RSCS Networking Version 2.
- (**) = JES3 supports 1 SYSIN and 1 SYSOUT stream in parallel.

JES2

Network Control

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 respective 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.

Network Addressing, Topology & Routing

Naming Conventions

JES2 supports up to 9999 remotes and 9999 nodes. Remotes and nodes are defined during JES2 initialization. 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.

Commands and Messages (NMRs)

JES2 sends an SRCB of X'00' on NMRs.

Command Transmission

JES2 builds unformatted nodal message records (NMRs) as a result of \$N operator commands.

Formatted (Global) Commands

JES2 builds formatted NMRs for "global" commands. The following global commands are supported by JES2:

Display job - Using the \$G D operator command

Hold job - Using the \$G H operator command

Release job - Using the \$G A operator command

Cancel job - Using the \$G C operator command

Route job (SYSIN or SYSOUT) - Using the \$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 job name is specified and JES2 finds more than one job with that name. A message is issued to the local console if the global command is rejected.

Command Authorization

JES2 supports four levels of command authority, which may be specified on a node basis. The authority levels are as follows:

Network - The node has the same authority as local consoles; this includes device, job and system authority.

Device - The node has the authority to issue device-related commands.

Job - The node has the authority to issue job-related commands.

System - The node has the authority to issue certain system commands.

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.

SYSIN (Job Input) Transmission

Store-and-Forward Transparency

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

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 job name 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 job name in the job header. All relevant job header information (such as priority and accounting information) is taken from the preceding JOB statement.

Job Header

The job header and trailer are created when a job is read into the JES2 system and 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.

Note: JES2 segments control records in lengths of 253 bytes instead of the preferred 255.

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. JES2 appends this section to the job header even at intermediate nodes.

Job Scheduling Section

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 (pre-SP1.3.3) require no additional code to store-and-forward this section.

User Sections

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).

Job Trailer

The actual number of lines, cards, pages, 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 contain 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.

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 maximum length of an NJE SYSIN record is 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.

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

If an error is encountered while processing JECL statements, JES2 input processing skips the remaining records in the job and queues the job (with error messages) for output. Error messages are returned to the originating node and user.

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 job header contains a userid and the job is transmitted from its originating 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 at the originating node. If the originating node and execution node are the same, this results in a TSO SEND command specifying the userid. Otherwise, the notification is sent in a nodal message record to the originating node, and that node issues the SEND command to the userid. There is also notification when each SYSOUT data set group reaches the destination node. (See "Notification" on page A-6 for details.)

SYSIN Job Routing

Use of System Qualifier

There is no facility to route a job to a particular system in a JES2 MAS configuration. 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 jobs 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).

Jobid Assignment

The JES2 job identifier (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 a receiving system (as may occur for spin data sets). The newly-assigned jobid is not transmitted in the job header.

SYSOUT (Job Output) Transmission

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 does not add the accounting section to the job trailer.

JESNEWS

In certain situations, JES2 attempts 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 from an intermediate node. JESNEWS is appended only if both of the following two conditions are met:

1. JESNEWS exists on the transmitting system.
2. The job log contains variable length records.

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 also provides an option to preserve trailing blanks on spool for specified SYSOUT classes.

Job Header

As an intermediate node, JES2 adds the JES2 and job scheduling sections to the end of the job header if that header does not already contain those sections.

Job Copies

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 calculate the total number of copies for a data set.

The \$T O operator command is used to update the characteristics of data sets at a JES2 node. The 3800 characteristics is updated only if the existing data set header contains a 3800 section. JES2 does not update any data sets that have multiple "clone" copies (that is, /*JOBPARM copies or those created by a \$N operator command).

Job Trailer

Accounting Section

As an intermediate node, JES2 will not add the accounting section to the end of the job trailer for SYSOUT.

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 adds 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) that precedes the buffer).

3800 Section

On the execution node, JES2 does not create a 3800 section if the data set has only the default 3800 characteristics specified.

Multiple Data Set Headers

JES2 is able to send and receive multiple data set headers for a data set with more than one destination.

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 originating node. See "Notification" on page A-4 for a description of the end-of-execution message.

Job Output Routing

Default Output Routing

Unless otherwise specified, JES2 routes job output back to the originating node and remote workstation unless specifically routed with a /*ROUTE or other JCL/IECL statement.

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.

Interactive Data Transmission Facility

A file sent by the TSO/E TRANSMIT command 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 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 an installation 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".

Stream Support and Control

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.

Stream Initiation and Suspension

Receiver Cancel

JES2 responds to a "receiver cancel" with a 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.

SNA Support

Session 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 node.

Bind Parameters

JES2 creates the BIND from three sources: MODTAB, JES2 initialization parameters, and an internal table that forces certain parameters such as TS and FM profile types. JES2 allows only the maximum primary and secondary RU sizes to be variable and requires the network topology and output dispersal flags to be on in FMH type 4.

SESSIONC Usage

JES2 sends SESSIONC for the following cases:

- Invalid application name in the JES2 application table.
- No logical SNA line available for this session.
- Invalid parameters in the BIND.
- This node is already in session or another OPNSEC is pending.

Function Management 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 and the receiving node has indicated compaction is accepted via the flag bit in FMH4.

RU Composition

RU Size Determination

The NJE RU size is determined by the TPDEF BUFSIZE 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.

Termination

A JES2 NJE session using SNA is normally terminated by a \$P LINE command. Normal session termination may be initiated by either the primary or secondary LU node. The purpose of abnormal termination is to clear the session as quickly as possible. A session is abnormally terminated when the line is restarted (\$E LINE). TERMSESS is sent by JES2 for the following reasons:

1. CLEANUP RU received in NS exit
2. Logic error - no OPNDST AUTH = (ACQ).

BSC Line and CTC Adapter Support

CTC Adapter Initialization

JES2 issues the SYN NAK as specified in "Initialization" on page 6-8.

CTC Adapter Attention and Busy Status

JES2 performs one retry in this condition, which resets the attention 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.

Error Recovery

JES2 follows the actions in Figure 6-5 on page 6-5, and terminates after 10 errors.

Use of Null Buffers

JES2 uses DLE ACK0 instead of null buffers for positive acknowledgement when there is no data to send.

Accounting

Accounting Records

JES2 uses SMF type 26 (job purge) records to record all successful SYSIN job transmissions; for all successful SYSOUT transmissions, JES2 uses SMF type 57 records. Because 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, which can present problems 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 job executes at the originating node.

Type 57 Records

The following standard job header information is missing from the type 57 record cut by JES2:

- Job name
- Time and date on the reader at original node
- User identification from the common exit parameter area.

NJE Network Management Records

JES2 records the following information reflecting network events:

- 55 Network signon
- 56 Network integrity (invalid password)
- 58 Network signoff.

Other Records

The following information is missing from the type 6 record cut by JES2 and other records (such as, types 4 and 5):

- Original job number
- Original node name.

Network Account Number

JES2 uses the following scheme to derive the network account number in the job header:

1. Obtains it from the /*NETACCT JECL statement, if specified by the user.
2. If not explicitly specified in the JECL, then it may be converted from a local account number to a network account number through a local-to-network account table set up by the NETACCT JES2 initialization statements.
3. Otherwise, defaults it to the local account number.

Miscellaneous Considerations

Installation 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 TRANSMIT command. 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 following exits in JES2 input processing could be used to screen jobs received by 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).

These 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.

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.

JES3

Network Control

Network Connection Control Records

Path Manager Records

JES3 does not include a path manager; JES3 supports only SRCB record types I, J, and B.

Network Addressing, Topology, and Routing

Naming Conventions (Remote vs. Userid)

To the JES3 user and operator, there are only two levels of qualification for the specification of destinations and origins:

1. Node
2. Either the remote identifier 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 identifier. The networking code does not know what this secondary destination is, and usually considers it to be a remote identifier. If the secondary destination must be placed in a field which, by definition, could contain either a userid or a remote identifier, and a flag set to indicate which it is, the flag that is set will indicate a remote identifier.

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. Also, the networking code runs only on the global processor. 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 system qualifier in the NJE header indicates to which system in the complex an individual user belongs. 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 qualifier value (1-8) corresponds to the relative position of the main processor control (MPC) table 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 that are not submitted from TSO, NJHGORGQ will be zero.

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). See "Stream Support and Control" on page A-14 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.)

Commands and Messages (NMRs)

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 indexed in a table called the console status table (CST). 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 supports only 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.

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 to 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 NMRROUT field, which comprises the second two bytes of the 8-byte destination field.

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. Formatted commands are translated into equivalent JES3 commands if such exist.

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.

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 and 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 an exit (IATUX35) that will accept other commands, or place further restrictions on the commands listed above. This exit replaces JES3 standard validation/authorization of the command. (In addition, exit IATUX18 may be used for additional authorization checking.)

SYSIN (Job Input) Transmission

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.

Data Set Header (RCCS)

JES3 does not support receipt of the record characteristics change section (RCCS) that accompanies SYSIN data greater than 80 bytes. If received, this section is ignored. It will not be store-and-forwarded.

Acceptable Job Streams

Requirement For Two Job Statements

To transmit a job from a JES3 system for execution at another node, the following JCL must be submitted:

```
//abc JOB xxxxx first job statement
//*NETACCT .... (optional)
//*ROUTE XEQ nodename
//xyz JOB yyyy 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.
- Submitted jobs 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.

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 "Notification" on page A-14 for notification during SYSOUT processing.

SYSIN Job Routing

Use of System Qualifier

See "Network Addressing, Topology, and Routing" on page A-10 for details.

Use of Userid

See "Network Control" on page A-10.

Operator Re-routing

JES3 does not permit operator re-routing of SYSIN jobs.

Jobid Assignment

JES3 tries to assign the original job number to SYSIN and SYSOUT jobs received over the network. If the original job number is not available, JES3 assigns the next available number.

SYSOUT (Job Output) Transmission

Job Output Routing

In SP 1.3.4, JES3 uses the submitting node as the default destination for output routing. Prior releases used the execution node as the default routing node.

Store-and-Forward Transparency

Spanned Record Support

JES3 supports spanned records.

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 do not exist in jobs sent by JES3.

Notification

JES3 sends notification messages for job output:

- 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 (destination) node.

See "Notification" on page A-13 for notification during SYSIN processing.

Job Output Routing

Use of System Qualifier

See "Network Addressing, Topology, and Routing" on page A-10 for details.

Use of Userid

See "Network Control" on page A-10.

Default Output Routing

JES3 sets 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.

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 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 (jobs or SYSOUT).

Multiple Streams

JES3 supports one stream for job transmission and one stream for SYSOUT transmission. In each case, the stream supported is stream 1 (RCB 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 sends a receiver initiated sequence (RCB X'D0') to NODE2.
8. NODE2 then varies all the senders to NODE1 online. Jobs in specialized reschedule status that are waiting for the senders to become available will now be available for scheduling.
9. NODE2 can reissue the request permission to initiate a 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 to handle both jobs and SYSOUT, "receiver on" cannot 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.

SNA Support

JES3 does not support SNA as a communication medium.

BSC Line and CTC Adapter Support

CTC Adapter Initialization

JES3 checks the data returned from the sense CCW and if it is X'00', a control-read sequence is issued; if it is X'01', a read-sense-write-control-read is issued; otherwise, a write-control-read is issued.

CTC Adapter Termination Deviations

For SP 1.3.4, JES3 does not send a SIGNOFF record to terminate a CTC adapter connection. Instead, a negative acknowledge (NAK) is transmitted and the connection is terminated. For SP 2.1.5, JES3 sends a signoff to terminate a CTC adapter connection.

CTC Adapter Attention and Busy Status

Some systems (RSCS, for example) place a TIC command between the standard CONTROL and READ commands for the CTC adapter. 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 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.

CTC Adapter Zero Byte Operation

It is possible for a READ or WRITE operation on a CTC adapter to terminate with successful 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 without sending a negative acknowledge. (In JES3 1.3.4, NAK is not sent; in 2.1.5, a signoff is not sent.)

Error Recovery

BSC Initialization Error Recovery

JES3 differs from Figure 6-4 on page 6-4 ("BSC Initialization Error Recovery") 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 6-5 on page 6-5 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.

BCB Handling

JES3 has the following deviations from the correct BCB recovery procedures in Figure 6-15 on page 6-17:

1. When in state S1 or S3 and a NAK is received (E3), if the line is a CTC adapter, 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 re-transmitted.
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.

Null Buffers

JES3 does not send null buffers as a response. DLE ACK0 is always sent as an acknowledgement.

Signon Deviations

JES3 does not negotiate the buffer size in the SIGNON records. If a mismatch occurs, the line is terminated.

Accounting

Use of Job Header and Trailer Fields

There are several fields in the NJE job header that relate to the recording of SMF information. For example, NJHGETIM, which is the estimated job execution time, and NJHGELIN, the estimated output print lines. The job trailer is also 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 that is recorded by JES3 is culled from other sources, such as internal control blocks.

Accounting Records (SMF)

System management facilities (SMF) 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 (types 26 and 57) are recorded by JES3 that contain network-related information. 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 that record information of a network management nature.

Network Account Number

Users submitting jobs from JES3 can specify certain accounting information they wish to accompany their job. This is accomplished by using the `/**NETACCT` control statement in the jobs JCL stream. Refer to "Acceptable Job Streams" on page A-13 for information on the placement of this statement in the JCL stream. The information that 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 built for the job.

Miscellaneous Considerations

Use of Utility Jobs in JES3

JES3 uses "utility jobs" to process network traffic. A utility job is created when:

- An executing job produces output destined for another node.
- A job is submitted for execution on another node.
- A job or SYSOUT data is received for either local processing or store-and-forwarding.

A "job" in the JES3 sense is a set of well-defined, schedulable elements that 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.

Installation Exits

JES3 networking code contains several installation exits. Some 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 exits follows.

- IATUX35 - allows the installation to perform special command validation/authorization for commands that execute locally.
- IATUX36 - allows the user the pull accounting information from the first segment of the job header for received jobs and SYSOUT.
- IATUX37 - allows the modification of the first segment of the data set header for received SYSOUT data sets.
- IATUX38 - allows the user to perform special processing of SYSOUT classes for received data sets.
- IATUX39 - 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.
- IATUX40 - allows modification of the job header that is built for locally-submitted jobs.

- IATUX43 - allows access to the entire job header for SYSOUT data sets from locally-executed jobs that are to be transmitted. The exit is called at intermediate nodes

RSCS

Network Control

Signon

RSCS Version 1 will only act as the primary for signon. RSCS sends the signon record with reset BCB specified. In Version 2, this restriction is eliminated.

Network Path Management

RSCS does not support a path manager. Any connection between RSCS and JES2 must be pre-defined by a JES2 CONNECT statement. RSCS will not forward Network Path Manager (NPM) topology records.

Network Connection Control Records

RSCS supports only (SRCB) type I, J, and B records.

Signoff

RSCS will not automatically restart the line after a signoff is received.

Network Addressing, Topology, and Routing

Remotes and Userids

VM makes extensive use of user identifications (userids) to identify virtual machines. Therefore, RSCS has record of the userids of originators and final receivers of network data. It also handles data for remote workstations. However, a "remote", as defined for JES2 and other NJE systems, is not used in exactly the same way in a VM environment.

The VM 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 identifiers for workstations that are used by other NJE systems with no restrictions on names. However, a remote and userid cannot 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 for userids, fields defined for remotes.

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.

Commands and Messages (NMRs)

General NMR Processing

RSCS Version 1 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 identifiers, logical routing information, and node qualifiers, this information is lost if the NMR record is sent through an RSCS node. The fields that are not supported are NMRTQUL, NMRFMQUL, NMRUCM, NMRDESC, and most other redefinitions of NMROUT except for userid or remote identifier.

In Version 2, this 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 identifiers in the command NMR. The console command response is treated like any other message originating from the node upon which RSCS is running.


Commands

Formatted Commands


RSCS networking does not support formatted commands that originate at its node. It will pass on formatted commands (received from another NJE node) correctly. It also will translate formatted commands into equivalent RSCS commands according to the setting of the NMRFOP field as follows:

1. NMRFOPD (display job) - Query file spool identifier
2. NMRFOPC (cancel job) - Purge spool identifier
3. NMRFOPA (release job) - Change spool identifier nohold
4. NMRFOPH (hold job) - Change spool identifier hold
5. NMRFOPR (route job) - Transfer spool identifier to local userid.

With Version 1, if a formatted command passes through a VM to VM connection, the command will not be understood. The reason for this is that the connection between two RSCS nodes must use 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 that 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 that implement full NJE protocols.



(In Version 2, this 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 Version 1, certain NMRTYPE bits are not store-and-forwarded under various conditions; however, all bits are store-and-forwarded

in Version 2 all bits are store-and-forwarded and the NMRFLAGT field is always turned on when a message contains a destination userid.

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.)

SYSIN (Job Input) Transmission

Store-and-Forward Transparency

Segmented Headers

RSCS supports segmented job headers, data set headers (for both SYSIN and SYSOUT), and segmented job trailers up to 4092 bytes.

Trailing Blank Truncation

For any type of file (SYSIN, print, or punch SYSOUT), any records with trailing blanks have the blanks truncated when they are forwarded by RSCS.

In Version 2 (with VM/SP Release 4), 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.

Data Set Header (Record Characteristics Change Section)

RSCS does nothing with input data set headers. It does not look at them at all, but merely puts them in the same file with the rest of the SYSIN job in which they are contained.

All SYSIN data that is store-and-forwarded through RSCS must be 80 bytes or less. RSCS rejects (with receiver cancel) any SYSIN job that contains records longer than 80 bytes. (In Version 2, this restriction is eliminated.)

Notification

RSCS sends a notification message (DMTxxx147I) when the file is successfully transmitted on a link. For Version 2, DMTAXM104I is sent when the file is spooled to the user.

Spool File Identifier Assignment

CP assigns a new spool file identifier when it receives a file. No attempt made to re-assign the same number it had on the originating node (in NJHGJID).

SYSOUT (Job Output) Transmission

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 forwards them with a machine carriage control character of X'09' added. The NDHGRCFM and NDHGLREC fields in the data set header are changed to reflect this addition.

In Version 2, this 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.

In Version 2.3, the NDHGRCFM field has no bearing.

2. When records with ASA carriage control are store-and-forwarded, RSCS converts the ASA carriage control to machine carriage control. The NDHGRCFM header field is changed accordingly. Records sent with machine carriage control character are store-and-forwarded without change.
3. Records (with or without carriage control of any type) sent as spanned records are store-and-forwarded without change.
4. RSCS truncates any characters in print SYSOUT records that exceed the number allowed for the printer whose type is used to store the file on CP spool. (The method in which this printer type is determined is described in detail in section RSCS/CP Spool Interface Considerations.

In Version 2, this restriction is eliminated for store-and-forward files. At the destination node, records are truncated as described above.

Punch File Transparency

RSCS handles incoming SYSOUT punch files as follows:

1. In Version 1, 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 a receiver cancel.

This restriction is eliminated in Version 2.

2. If NDHGRCFM indicates that a punch file contains ASA or machine carriage control, RSCS replaces the carriage control in each record with a machine punch operation code (opcode). 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 with no carriage control).

Trailing Blank Truncation

For any type of file (SYSIN or print or punch SYSOUT), any records with trailing blanks have the blanks truncated when they are forwarded by RSCS.

In Version 2 (with VM/SP Release 4), blanks are 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 CP component of VM 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 that RSCS uses are 1403, 3211, and 3800. (Note that CP supports more virtual print types than RSCS.) 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 Version 1 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 will modify any records that conflict (truncate them) so they no longer conflict. This is done only at the destination node.

Figure A-2 is a table of virtual device types that are defined for different incoming NJE file types:

NJE data	CP device type
SYSIN	Punch - limited to 80 characters
SYSOUT:	
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 below in greater detail.)

Figure A-2. RSCS Virtual Device Types

Notes:

1. The table is accurate only for RSCS Version 1 and above running on a VM/SP Level 2 and above system and assumes all data set headers contain the print/punch flag.
2. When a data set header contains an RSCS subsection, RSCS uses the originating device type in this section to determine the VM device type. Therefore, the SYSOUT part of this table only applies to files which do not originate on a VM system.

Once RSCS has determined the device type 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 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. **One result of this processing is that RSCS can not store-and-forward print records greater than 204 unless they are sent as spanned records.**

Version 2 defines files in the same way as shown in Figure A-2. However, it is able to store records that are longer than the normal length specified for the device type, without truncating data.

3800 Files

Version 1, Release 3 of 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)
- Virtual 3800 files that are currently only understood by a VM system and thus are only originated by RSCS.

Virtual 3800 files are also sent with a 3800 section in their data set 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 opcodes as well as machine carriage control characters and, secondly, they are the *only* type of file with a 3800 subsection that are originated by RSCS.

When virtual 3800 files are created by an RSCS node, they are sent with both an RSCS subsection and a 3800 subsection in the data set 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.

RSCS creates a virtual 3800 file out of any 3800 print file that 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 at the destination node and inserts select CCWs that 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 node, the printout should be as the originator intended. If the same file is sent back to another non-VM 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 types of print files (see "Store-and-Forward Transparency" on page A-21).

Version 2, Release 1 3800 support is the same as it is on Version 1, Release 3 except that

- Virtual 3800 records sent without carriage control are not modified on forwarding. They are forwarded without carriage control. However, records sent with ASA carriage control are changed to machine carriage control. NDHDFCFM is changed to reflect this.
- Spanned 3800 records are not unspanned when they pass through an RSCS node. This is only done when the files are printed on the node RSCS is running on.

Support of SRCB X'B0'

RSCS does not support the AFP SRCB (X'B0'). It cannot initiate files with this SRCB. RSCS Version 1 rejects, with a X'B0' RCB, any file containing this SRCB (page mode records).

In Version 2, 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 that contains them.

Spanned Records

Version 1 handles spanned records with machine carriage control and no carriage control when acting as an end node, provided the maximum length field in the data set header (NDHGLREC) is filled in correctly. ASA carriage control is not supported in Version 1 for spanned records if RSCS is the end node. (In Version 2, this restriction is eliminated.)

Data Set Header

Use of CP TAG Command

RSCS has no direct JCL equivalent to allow users to specify parameters for jobs or data sets. However, RSCS allows a file originator a limited method of specifying certain fields in the data set 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 to be included for fields RSCS does not set. The CP TAG command, described in *RSCS Operation and Use Guide*, is used to specify these parameters. The TAG command can only be used for fields in the data set header.

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 data set or many data sets may be grouped in the job because they are all output from its execution. In the second case, each data set is preceded by a data set header. RSCS handles these cases as follows:

1. When a single data set is received by RSCS with multiple destinations, the data set is store-and-forwarded as separate files (one for each destination). RSCS will only make separate files if the data sets involved are going out on different links.
2. When any data set within a job has different characteristics from another in the same job, each data set is store-and-forwarded as a separate file.
3. Spin data sets (flag NDHGF1SP on) are always stored as separate files.

Note: When data sets are store-and-forwarded as separate files, each file has the original job header and job trailer.

Version 2 uses RSCS uses the LIST processor to generate multiple data set headers.

Use of HOLD in the Data Set Header

RSCS Version 1 does not respond to flag NDHGF1HD in the data set header. This flag, if on, indicates that the data set is to be held at its final destination, not at the receiving node. Because RSCS does not process headers on all links (see section "VM-Specific Line Drivers" on page A-29), 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, this restriction is eliminated.)

Notification

RSCS sends a notification message when the file is successfully transmitted on a link. (DMTxxx147!)

Special Processing

Print File Processing

If RSCS is the destination, it will default CCWs it does not understand.

Punch File Processing

Punch files originated by RSCS Version 1 are sent as fixed 80-byte records without carriage control. In Version 2, punch files may be sent either 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 using the PUNCC option.

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 NJE system goes to the system printer or punch and not to the reader of the user who originated the job. (In Version 2, this restriction is removed.)

Fanout (Optimized) Support

RSCS Version 2, using the LIST processor, will not send multiple data set headers with a file.

Operator Rerouting

The TRANSFER command can be used by the RSCS operator (and by the file originator, in Version 2) to re-route files to other nodes.

In Version 1, this command only changes the tag block; it does not change the network header fields. Therefore, if the file is sent to another node over an NJE link, the updated routing information will be lost. (This restriction is removed in Version 2.)

Stream Support and Control

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. (In Version 2, this restriction is eliminated.)

RSCS networking supports only RCBs of X'98' for input and X'99' for output. RSCS rejects any other streams that are sent. (In Version 2, this restriction is eliminated.)

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, this restriction is eliminated.)

Abnormal Termination Protocol

Version 1.3

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 that 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.

Stream Initiation and Suspension

Stream Suspension

Because RSCS Version 1 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 sends wait-a-bit for a stream it is receiving when it only has one buffer left to hold the next input. (In Version 2, this restriction is eliminated.)

SNA Support

RSCS Version 1 does not support SNA sessions for NJE. Version 2 of RSCS provides native SNA support.

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.

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 to be searched must be specified when defining the remote system to its local VTAM.

Function Management Headers

Compaction

RSCS does not support compaction. If an FMH3 is received, RSCS will terminate the session.

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 that it is 300 or greater. The actual transmission block size is the smaller of the two values specified in the signon records.

BSC Line and CTC Adapter Support

CTC Adapter Initialization

RSCS Version 1 does not follow the procedure outlined in "Initialization" on page 6-8. 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 the SOH ENQ. If the other side sends an I signon record, RSCS will ignore it and still send the I itself. (In Version 2, this restriction is eliminated.)

Wait-a-bit Processing

RSCS Version 1 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 in response time, 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 RSCS accepts and processes received buffers that have wait-a-bit set. It responds with a null buffer without any response time difference (that is, no special delay).

Error Handling

Initialization Error Recovery

RSCS does not follow the procedure outlined in Figure 6-4 on page 6-4. It retries 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 Version 1 follows the actions in Figure 6-5 on page 6-5, with the following exception. RSCS does not specifically diagnose the cause of the error for command reject, bus out check, equipment check, data check, data overrun, or lost data. If this happens on a write, the data is resent. If it happens on a read, a NAK is sent.

Version 2.1 RSCS follows the actions in Figure 6-5 on page 6-5 except for data check, data overrun, or lost data. If these errors happen on a write, the data is re-transmitted. If they happen on a read, a negative acknowledge is sent.

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, this restriction is eliminated.)

Accounting

Accounting Records

For information about RSCS accounting records, see *RSCS Exit Customization*, LY24-5240.

Miscellaneous Considerations

VM-Specific Line Drivers

RSCS Version 1 was designed so that there is different code used to communicate with non-VM NJE systems from that used to communicate with other RSCS systems. The result of this structure is that certain information about jobs and NMRs that is carried in the non-VM NJE code is lost when sent thru an RSCS node. The headers and trailers described in this document are *not* sent on a link that 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 that receives a file on an RSCS-RSCS link.

In Version 2, this restriction is eliminated. RSCS uses the same code for all communications.

Installation Exits

There are installation exits in RSCS to examine headers and trailers, but not for transmitted files.

POWER

Network Control

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 pre-defined by a JES2 CONNECT statement.

Network Addressing, Topology, and Routing

Naming Conventions

POWER supports up to 200 remotes (RJE workstations) and 250 in Version 2.2; both versions support an unlimited number of nodes. Remotes are defined during 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. Members 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 and used to dynamically alter the network routing and topology.

Network Routing

POWER does not support the JES2 Network Path Manager. All routing information must be pre-defined 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 primary route is not active, but the secondary route is active, then the secondary route will be chosen. Both routes cannot be used together. As soon as the primary route is available, all new transmissions will be sent over the primary route and the secondary route will cease to be used as soon as the current transmission finishes.

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 that 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 re-transmission from the beginning.

Transmission may be discontinued voluntarily or involuntarily:

- Voluntarily, by either the transmitter or receiver.
- Involuntarily, in the case of a line disconnection.

In either case, the receiving system discards the job or SYSOUT data that it has partially received. In the case of voluntary termination of transmission, the transmitting system will always requeue the job or SYSOUT data for re-transmission.

Commands and Messages (NMRs)

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.

Store-and-Forward Transparency

Commands and messages are transmitted to their final destination only if there is a path available. If the primary 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 regarding the inability to deliver the command. If the route is not available, messages are discarded without informing the originator.

Formatted (Global) Commands

Formatted commands are not supported by POWER.

Command Authorization

POWER supports three types of command authorization, which must be specified on a node basis in the node definition table:

- Network - The node has the authority to control many 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 - The node can only manipulate jobs or output that either originated at or is destined for that node.
- Nojob - The node is not allowed to manipulate anything within the network. Users can still display queues on other nodes.

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.

SYSIN (Job Input) Transmission

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 store-and-forward node can change the execution node of the job and this will be updated in the job header.

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 (such as 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.

Job Trailer

The job trailer is built as soon as the job enters the POWER system. Because POWER doesn't use a lot of the counts that JES2 and JES3 use for scheduling purposes, the majority of the fields in this record are not used and are set to defaults.

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, the record characteristics change 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.

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.

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 JCL JOB statement. Therefore, if a node other than the originating node is to be notified, then the originating node identifier is destroyed by the notify node name.

If the job header contains a notify userid and the job is transmitted from its originating node for execution on another node, the job transmitter of each intermediate node 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 by the originating node field.

SYSIN Job Routing

Use of the 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 on the * \$\$ JOB statement and is contained in the POWER section.

Use of the Userid

The userid field on the XDEST parameter can be used to route a job to a VM system for execution. If POWER is the target execution node, then the NJHGXEQU field 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.

Jobid Assignment

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 identifier.

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 SYSOUT. The new and original job numbers are displayed in the messages that are issued to acknowledge transmission of jobs or SYSOUT to the next node. The job header always contains the original (input system's) jobid. The jobid assignment procedure outlined above is followed by both job receivers and SYSOUT receivers.

If the SYSOUT has been segmented, the job name will remain the same as the job and each segment will retain the original job number.

SYSOUT (Job Output) Transmission

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 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 SYSOUT types (print and punch data), it will be split.

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 add a user section.

Spanned Records with Carriage Control

POWER supports spanned records with carriage control, both as intermediate node and as destination node.

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 2.1 converts all ASA carriage control to machine carriage control.

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 because their information has already been copied to the data set headers.

Job Trailer

The majority of the fields in the job trailer are not used by POWER and are set to default values.

Data Set Header

The data set header contains the information necessary to handle SYSOUT 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 as soon as any job has executed on a POWER node or SYSOUT 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 the POWER spool, even if the data set is not intended for transmission. The data set header record is flagged as an internal record, so it is ignored by local print/punch processors and external writers at the destination node. 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 SYSOUT data types in one transmission stream. If there is print and punch output to be sent, then this will be transmitted as two distinct transmissions, each having its own job header and job trailer.

3800 section

At execution time, POWER creates a 3800 section for all output for which a 3800 device had been specified. The values in the section are taken from the values specified in the * \$\$ LST statement or from the SETPRT defaults.

In POWER, there are two sets of defaults for the 3800 -- the hardware defaults and the software defaults. A value of X'00' denotes the use of hardware defaults and a value of X'40' denotes the use of software defaults for POWER; in JES2 and JES3, a value of X'00' denotes the use of software defaults (hardware defaults are unknown). This means that when POWER is the end destination, it must determine if the file originated at a POWER node. When that is not the case, the defaults are converted.

Output Processing Section

This section is built by POWER (starting with Version 2.3) when a user-defined output parameter is coded on the * \$\$ LST/PRT statement or when an OPTB structure is passed via the SAS interface at the beginning of output spooling.

POWER supports OPTBs up to a maximum length of 32,670. No additional code is required to have down-level POWER systems (prior to Version 2.3) store-and-forward this section.

Multiple Data Set Headers

POWER does not create multiple destination data sets. The receipt of multiple headers with no SYSOUT data in between indicates to the SYSOUT receiver that the data set has multiple destinations; the receiver builds 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, therefore, never generates multiple data set headers.

Notification

When a job's system output (or any part of it) reaches its ultimate destination, the SYSOUT receiver at the destination node examines the userid in the job header. If "NTFY=" was specified on the JCL JOB statement, the SYSOUT receiver issues a message to the specified user indicating that the job's output was received.

Job Output Routing

Default Output Routing

Unless otherwise specified on control statements, the default print and punch nodes are set to the originating 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.

Fanout (Optimized) Support

POWER does not support fanout. Any input stream containing multiple data set headers is broken down into streams consisting of single data sets using the following criteria:

POWER keeps all data sets belonging to a transmission together, provided the following important criteria do not change:

- 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 changes within one unit of transmission, then a new spool entry is built together with the job header and data set header. This is necessary because the operator can change the destination of any job or SYSOUT in the spool file and make its destination the local node. This means that all entries on spool are capable of being processed locally.

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.

Stream Initiation and Suspension

Operator Stream and Line Control

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 SYSOUT transmitter 1 are set in active status. If the POWER 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 1-7. The specification of JOB or OUT is necessary to set the correct RCB and FCS bits and inhibit the starting of invalid combinations of JOB and SYSOUT transmitters (for example, JOB transmitter 5 and SYSOUT transmitter 4).

In POWER, receiver tasks live only for the time span of one transmission (for everything between the job header and job trailer). After the transmission is complete (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 a job or SYSOUT is placed in the transmission queue for this node.

SNA Support

Data Flow Control

RSHUTD is the only data flow control (DFC) function supported by POWER. If any other function is received, POWER will terminate the session with a CLSDST or a TERMSESS.

Bind Parameters

POWER does not negotiate any bind parameters. They are “hard-coded” and most of the parameters received by the other node from POWER are ignored.

Functional Management Headers

The session will be terminated if the FMH received by POWER is not acceptable.

Compaction

POWER accepts and handles an FMH3 and decompacted inbound data. POWER does not compact outbound data.

RU Composition

RU Size Determination

The transmission block size is defined in the network definition table. This value is used in the signon record 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 RUs containing mixed record types.

BSC Line and CTC Adapter Support

BSC

Buffer Size

For BSC, the maximum transmission buffer size allowed by POWER Version 2.2 is 1800.

Error Recovery

Initialization Error Recovery

POWER will retry on Bus Out Check and Equipment Check before terminating.

CTC Adapter

The channel-to-channel adapter is supported by POWER Version 2.3 under control of VM/SP.

Accounting

For accounting information, see Chapter 4 in *VSE/POWER Networking User's Guide*, SC33-6140.

Miscellaneous Considerations

Installation Exits

For information about installation exits, see Chapter 4 in *VSE/POWER Networking User's Guide*, SC33-6140.

Appendix B. Abbreviations

- ACF/VTAM.** Advanced Communications Facility/Virtual Terminal Access Method
- AFP.** Advanced Function Printer
- ASA.** American Standards Association
- ASP.** Attached Support Processor
- BCB.** Block Control Byte
- BSC.** Bisynchronous Communication
- CCW.** Channel Command Word
- CES.** Connection Event Sequence
- CPDS.** Composed Page Data Stream
- CTC.** Channel to Channel (adapter)
- CTCA.** Channel to Channel Adapter
- EOF.** End of File
- FMH.** Function Management Header
- FCS.** Function Control Sequence
- HASP.** Houston Automated Spooling Priority
- JCL.** Job Control Language
- JECL.** Job Entry Control Language
- JES2.** Job Entry Subsystem 2
- JES3.** Job Entry Subsystem 3
- LRECL.** Logical Record Length
- LU.** Logical Unit
- ML.** Multi-Leaving
- MVS.** Multiple Virtual Storage
- MVS/SP.** Multiple Virtual Storage/System Product
- NCC.** Network Connection Control
- NDH.** Network Data Set Header

NJH. Network Job Header

NJT. Network Job Trailer

NJE. Network Job Entry

NJI. Network Job Interface

NMR. Nodal Message Record

NPM. Network Path Manager

OPTB. Output Parameter Text Block

POWER. Priority Output Writers, Execution Processors and Input Readers

RCB. Record Control Byte

RID. Record Identifier

RJE. Remote Job Entry

RSCS. Remote Spooling Communications Subsystem

RU. Request/Response Unit

SAF. Store and Forward

SCB. String Control Byte

SCIP. Set Control Interval Processing

SEGL. Segment Length

SNA. System Network Architecture

SRCB. Subrecord Control Byte

SWB. Scheduler Work Block

TSO. Time Sharing Option

TSO/E. Time Sharing Option/Extensions

TSS. Time Sharing System

UCW. Unit Control Word

VM. Virtual Machine

VM/370. Virtual Machine/370

VSE/POWER. Virtual System Extended/Priority Output Writers, Execution Processors and Input Readers

Glossary

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.

adapter. Is a mechanism for attaching parts. The mechanism can be either:

- 37xx line adapter
- Integrated channel adapter (ICA)
- PC communications card

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

AFPDS. Advanced function print data stream

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 signals generated at the sending and receiving stations.

C

CCTL. Carriage Control. The first character of an output record; it indicates the placement of that record in the print or punch medium.

CCW. Channel Command Word

CES. See "connection event sequence".

channel-to-channel (CTC) adapter. 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.

compaction. A method of reducing the length of records for transmission by representing certain 8-bit characters with 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). The CES value prevents 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.

D

destination node. The node that receives and processes an NJE job or an NMR.

destination user. A user or device (printer or punch) that is notified when a NJE SYSOUT job has been processed at the destination node.

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 SOH or STX.

end-of-transmission (EOT). A BSC control character indicating the end of a message transmission.

execution node. The node where an NJE SYSIN job or NMR command is executed.

- SYSOUT or messages generated from the execution are forwarded to a destination node.

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 on 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, a specialized control format to select a destination and control the way data is sent to the destination.

H

HASP. Houston Automated Spooling Priority - the predecessor to JES2.

I

intermediate node. A node that lies in the path of either the:

- Originating node and execution node
- Execution node and the destination node.

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. It controls all input and output, ensures complex-wide data integrity, and 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. See "NJE job".

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, and read by the job entry subsystem (JES2, JES3, or VSE/POWER).

job header. 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. The NJE control record that 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

LU0. 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 record. See "network connection control".

NDH. network data set header - see "data set header".

network. A group of one or more nodes that communicate by using direct lines, CTC links, or telecommunication lines.

network connection control (NCC) record. An NJE control record used to establish or break a connection. NCC records are 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/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 SNA, but is an application layer that uses SNA, BSC, and CTC transmission facilities.)

(3) The JES2 program product implementation of the NJE protocol.

NJE control records. The control records used by a network to define a unit of work. The possible NJE control records are:

- Job header
- Data set header
- Job trailer
- NMR header

NJE job. A unit of work that begins with a job header, is followed by data, and ends with a job trailer. The data between the job header and job trailer can be either SYSIN data or SYSOUT data.

NJE SYSIN job. An NJE job that contains JCL for a job and one or more SYSIN data sets.

NJE SYSOUT job. An NJE job that contains one or more SYSOUT data sets. Each SYSOUT data set is separated by a data set header.

NJE transfer unit. A unit of work that the network needs to process (an NJE job or an NMR).

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 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 complex that is part of a network using NJE protocols. The attributes for the node are defined in the initialization stream for the complex.

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

originating node. The node where a user submitted an NJE transfer unit for the network to process.

originating user. The user that submits an NJE transfer unit at the originating node. The originating user submits the NJE transfer unit at an operator console, terminal, or an RJE workstation. An NJE transfer unit may also originate from a file.

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. See "permission rejected".

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). In SNA transmissions, the RID consists of a 3-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 network.

RU. request/response unit

S

store-and-forward node. See "intermediate node".

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 "system network architecture". (2) (adjective) Adhering to the System Network Architecture definitions. SNA is used to describe links, I/O devices, telecommunication controllers, etc.

SOH. See "start-of-heading".

spool. simultaneous peripheral operation online (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 the 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' sent by the receiver after the transmitter sends an EOF. 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 that 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 may specify its further processing.

system network architecture. A formal definition of the format, protocol, and sequencing of information 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' 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 CTC adapter 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

accounting
 user
 definition of 1-5
ACK0 6-6
addressing nodes 1-6
Add/Subtract Connection Record 5-7

B

BIN 2-44
BIND 3-15
block control bytes (in BSC buffer) 3-4
BSC A-1
buffer fields
 SNA 3-14
buffer format
 BSC and CTC 3-2
 non-SNA 3-2

C

CCITT 6-6
CES 5-4
CICSDATA 2-46
CKPTLINE 2-42
CKPTPAGE 2-42
CKPTSEC 2-42
CLSDST 5-8
COMPACT 2-42
Compaction A-1
compressing
 data 1-11
connection
 establishing 5-3
 introduction 1-14
 of sessions 1-14
 terminating 5-8
connection control records
 format 2-67
CONTROL 2-42
CTC 6-8, A-1

D

data
 compressing 1-11
Data compression
 SNA data 3-13
data format (spanned and unspanned records) 3-17
data set header
 fields 2-21
DATAACK 2-46

destination
 node
 definition of 1-3
 user
 definition of 1-5
destination node, error handling 2-51
DLE ACK0 6-4, 6-6, 6-8, A-1
DLE ETB 6-6
DLE STX 6-6, 6-10
DPAGELBL 2-44
DUMP 2-46
DUPLEX 2-44

E

ETB 6-6
execution
 node
 definition of 1-3

F

Fanout 6-29, A-1
FILEID 2-44
FORMDEF 2-43
function control sequence 3-7
function control sequence (FCS)
 introduction 1-11

I

INDEX keyword 2-42
initial signon record
 format 2-67
interpreting fields for
 an NMR
 using fields NMRFLAGC 2-58
 using fields NMRTYPEF 2-58
introduction
 to NJE 1-1

J

job header
 fields 2-3
 introduction 1-8
job trailer
 fields 2-54
 introduction 1-9

L

LINDEX keyword 2-42
LU 6-24

LUTYPE 0 6-24
LU0 6-24

M

MESSAGES 2-43

N

NAK 6-5
NCCADL 2-70
NCCAEVNT 2-70
NCCANODA 2-70
NCCANODB 2-70
NCCAQULA 2-70
NCCAQULB 2-70
NCCAREST 2-70
NCCCDL 2-69
NCCCEVNT 2-69
NCCCRB 2-69
NCCCREST 2-69
NCCIBFSZ 2-67
NCCIBUFSZ 2-68
NCCIDL 2-67, 2-68
NCCIEVNT 2-67, 2-68
NCCIFEAT 2-67
NCCIFLG 2-67, 2-68
NCCIFLGM 2-68
NCCILPAS 2-67, 2-68
NCCINODE 2-67, 2-68
NCCINPAS 2-67, 2-68
NCCIPREP 2-67
NCCIQUAL 2-67, 2-68
NCCIREST 2-67, 2-68, 2-69
NCCITRM 2-67
NCCRCB 2-67, 2-69, 2-70
NCCSRCB 2-67, 2-68, 2-69, 2-70
NDHACPYG 2-35
NDHAFLCT 2-33
NDHAFLG1 2-32
NDHAFLSH 2-34
NDHAF1BN 2-32
NDHAF1BR 2-32
NDHAF1J 2-32
NDHALEN 2-32
NDHAMOD 2-32
NDHAMODF 2-35
NDHATAB1 2-33
NDHATAB2 2-34
NDHATAB3 2-34
NDHATAB4 2-34
NDHATREF 2-33
NDHATYPE 2-32
NDHCFLG1 2-38
NDHCLEN 2-38
NDHCLREC 2-38
NDHCMOD 2-38
NDHCRCFM 2-38
NDHCTYPE 2-38
NDHFLAGS 2-22
NDHFLG2 2-30
NDHGCLAS 2-23
NDHGDD 2-23
NDHGDSCT 2-27
NDHGDSNO 2-23
NDHGFCB 2-28
NDHGFCBI 2-27
NDHGFLG1 2-24
NDHGFORM 2-28
NDHGF1DF 2-25
NDHGF1HD 2-24
NDHGF1IN 2-25
NDHGF1LC 2-25
NDHGF1LG 2-24
NDHGF1OV 2-24
NDHGF1SP 2-24
NDHGF1ST 2-25
NDHGF2PR 2-30
NDHGF2PU 2-30
NDHGLEN 2-22
NDHGLNCT 2-28
NDHGLREC 2-26
NDHGMOD 2-22
NDHGNAME 2-30
NDHGNODE 2-22
NDHGNREC 2-23
NDHGPMD 2-31
NDHGPROC 2-22
NDHGRCFM 2-25
NDHGRMT 2-22
NDHGSTEP 2-22
NDHGTYPE 2-22
NDHGUCS 2-29
NDHGUCSD 2-31
NDHGUCSF 2-31
NDHGUCSO 2-30
NDHGXWTR 2-29
NDHLEN 2-22
NDHPCOMP 2-37
NDHPDISP 2-37
NDHPFLG1 2-37
NDHPIDEV 2-37
NDHPJBSF 2-37
NDHPJNUM 2-37
NDHPLEN 2-37
NDHPMOD 2-37
NDHPNSEP 2-37
NDHPOPTN 2-37
NDHPPART 2-37
NDHPPASS 2-37
NDHPPRIO 2-37
NDHPRCFM 2-37
NDHPSETP 2-37
NDHPSTRT 2-37

NDHPSYID 2-37
 NDHPTYPE 2-37
 NDHPUSER 2-37
 NDHSCPDS 2-39
 NDHSDLEN 2-40
 NDHSEQ 2-22
 NDHSFLEN 2-39
 NDHSFLG1 2-39
 NDHSFLG2 2-40
 NDHSGPID 2-40
 NDHSLEN 2-39
 NDHSMOD 2-39
 NDHSNSTR 2-39
 NDHSOPTB 2-40
 NDHSPLEN 2-40
 NDHSPRID 2-40
 NDHSTYPE 2-39
 NDHSVERB 2-40
 NDHSVERS 2-40
 NDHTCNOD 2-53
 NDHTFEN 2-52
 NDHTFLG0 2-52
 NDHTFLG1 2-52
 NDHTFLG2 2-52
 NDHTF0JB 2-52
 NDHTLEN 2-52
 NDHTLENP 2-52
 NDHTLENT 2-52
 NDHTMOD 2-52
 NDHTOGRP 2-53
 NDHTOUSR 2-53
 NDHTPOEN 2-53
 NDHTPOEX 2-53
 NDHTSECL 2-53
 NDHTSGRP 2-53
 NDHTSNOD 2-53
 NDHTSTYP 2-52
 NDHTSUSR 2-53
 NDHTTYPE 2-52
 NDHTVERS 2-52
 NDHT1EXT 2-52
 NDHT2DFT 2-52
 NDHT2MLO 2-52
 NDHT2RMT 2-52
 NDHT2SHI 2-52
 NDHT2SUS 2-52
 NDHT2TRS 2-52
 NDHULEN 2-53
 NDHUMOD 2-53
 NDHUTYPE 2-53
 NDHVCLAS 2-36
 NDHVDIST 2-36
 NDHVFLG1 2-36
 NDHVFNAM 2-36
 NDHVFTYP 2-36
 NDHVIDEV 2-36
 NDHVLEN 2-36
 NDHVMOD 2-36
 NDHVPGLE 2-36
 NDHVPRIO 2-36
 NDHVRELN 2-36
 NDHVTAGR 2-36
 NDHVTYPE 2-36
 NDHVVRSN 2-36
 network facility 1-2
 NJE job
 contents of 1-8
 data 1-8
 introduction 1-8
 format of 1-8
 introduction 1-8
 job header 1-8
 job trailer 1-8
 unit of work 1-2
 NJE protocols
 introduction 1-7
 NJE SYSIN job
 format of 1-8
 NJE SYSOUT job
 format of 1-8
 NJE transfer unit
 definition of 1-10
 NJE transfer units
 transmission buffers 1-11
 NJHEBYTE 2-18
 NJHELEN 2-18
 NJHEMOD 2-18
 NJHEPAGE 2-18
 NJHETYPE 2-18
 NJHFLAGS 2-4
 NJHGACCT 2-9
 NJHGBLDG 2-15
 NJHGDEPT 2-15
 NJHGECD 2-14
 NJHGELIN 2-14
 NJHGETIM 2-14
 NJHGETS 2-10
 NJHGFLG1 2-5
 NJHGFORM 2-13
 NJHGF1PR 2-5
 NJHGHOPS 2-8
 NJHGICRD 2-14
 NJHGJCLS 2-4
 NJHGJCPY 2-8
 NJHGJID 2-4
 NJHGJNAM 2-9
 NJHGLN 2-4
 NJHGLNCT 2-8
 NJHGMCLS 2-4
 NJHGMOD 2-4
 NJHGNPAS 2-10
 NJHGNREC 2-15
 NJHGNTYN 2-15
 notify node 2-16

NJHGORGN	2-10	NJHTSGRP	2-20
NJHGORGQ	2-7	NJHTSNOD	2-20
NJHGORGR	2-10	NJHTSTYP	2-19
NJHGYPASS	2-10	NJHTSUSR	2-20
NJHGPRGN	2-14	NJHTTYPE	2-19
NJHGPRIO	2-6	NJHTVERS	2-19
NJHGPRTN	2-11	NJHT1EXT	2-19
NJHGPRTR	2-12	NJHT2DFT	2-19
NJHGPUNN	2-12	NJHT2MLO	2-19
NJHGPUNR	2-13	NJHT2RMT	2-19
NJHGRROOM	2-14	NJHT2SHI	2-19
NJHGTYPE	2-4	NJHT2SUS	2-19
NJHGUSID	2-9	NJHT2TRS	2-19
NJHGXEQN	2-11	NJHULEN	2-20
NJHGXEQU	2-11	NJHUMOD	2-20
NJHLEN	2-4	NJHUTYPE	2-20
NJHPDCY	2-17	NJH2ACCT	2-16
NJHPDD	2-17	NJH2FJOB	2-16
NJHPDDM	2-17	NJH2FLG1	2-16
NJHPDGP1	2-17	NJH2FSTC	2-16
NJHPDGP2	2-17	NJH2FTSU	2-16
NJHPDG1D	2-17	NJH2GRP	2-16
NJHPDG1R	2-17	NJH2LEN	2-16
NJHPDG1T	2-17	NJH2MOD	2-16
NJHPDG1W	2-17	NJH2SGRP	2-16
NJHPDG2X	2-17	NJH2SUSR	2-16
NJHPDISP	2-17	NJH2TYPE	2-16
NJHPDMD	2-17	NJH2USE	2-16
NJHPDN	2-17	NJH2USID	2-16
NJHPDND	2-18	NJH2USR	2-16
NJHPDNM	2-18	NJTFLAGS	2-54
NJHPDNTH	2-18	NJTGACRD	2-55
NJHPDNTM	2-18	NJTGALIN	2-55
NJHPDNY	2-17	NJTGAXPR	2-56
NJHPDSKT	2-17	NJTGFLG1	2-54
NJHPFLG1	2-17	NJTGIOPR	2-56
NJHPLEN	2-17	NJTGIXPR	2-55
NJHPMOD	2-17	NJTGLEN	2-54
NJHPSYID	2-17	NJTGMOD	2-54
NJHPTYPE	2-17	NJTGSTOP	2-55
NJHPUSER	2-17	NJTGSTRT	2-55
NJHSEQ	2-4	NJTGTTYPE	2-54
NJHTCNOD	2-20	NJTGXCLS	2-54
NJHTFEN	2-19	NJTLEN	2-54
NJHTFLG0	2-19	NJTS	2-57
NJHTFLG1	2-19	NJTSABYT	2-57
NJHTFLG2	2-19	NJTSAPAG	2-57
NJHTF0JB	2-19	NJTSEQ	2-54
NJHTLEN	2-19	NJTSLEN	2-57
NJHTLENP	2-19	NJTSMOD	2-57
NJHTLENT	2-19	NJTSTYPE	2-57
NJHTMOD	2-19	NJTULEN	2-57
NJHTOGRP	2-20	NJTUMOD	2-57
NJHTOUSR	2-20	NJTUTYPE	2-57
NJHTPOEN	2-20	NMR	
NJHTPOEX	2-19	command	
NJHTSECL	2-20	format of	1-9
		format of	1-9

NMR (continued)
 introduction 1-9
 message
 format of 1-9
 unit of work 1-2

NMRCONID
 message 2-65

NMRDESC
 message 2-65

NMRECSID
 message 2-65

NMRFD
 formatted command 2-62

NMRFFLG
 formatted command 2-62

NMRFFLGD
 formatted command 2-62

NMRFFLGO
 formatted command 2-62

NMRFJID
 formatted command 2-62

NMRFJNAM
 formatted command 2-62

NMRFLAG 2-58
 formatted command 2-61
 message 2-63
 unformatted command 2-59

NMRFLAGC
 formatted command 2-61
 message 2-63
 unformatted command 2-59

NMRFLAGD
 formatted command 2-61
 message 2-64
 unformatted command 2-59

NMRFLAGJ
 formatted command 2-61
 message 2-63

NMRFLAGR
 formatted command 2-61
 message 2-63

NMRFLAGS
 formatted command 2-61
 message 2-64
 unformatted command 2-59

NMRFLAGT
 formatted command 2-61
 message 2-63
 unformatted command 2-59

NMRFLAGU
 formatted command 2-61
 message 2-63
 unformatted command 2-59

NMRFLAGW
 formatted command 2-61
 message 2-63
 unformatted command 2-59

NMRFM
 formatted command 2-62
 message 2-65
 unformatted command 2-60

NMRFMNOD
 formatted command 2-62
 message 2-65
 unformatted command 2-60

NMRFMQUL
 formatted command 2-62
 message 2-65
 unformatted command 2-60

NMRFOP A-20
 formatted command 2-62

NMRFOPA A-20

NMRFOPC A-20

NMRFOPD A-20

NMRFOPH A-20

NMRFOPR A-20

NMRFORGN
 formatted command 2-62

NMRFR
 formatted command 2-63

NMRLEVEL
 formatted command 2-61
 message 2-64
 unformatted command 2-59

NMRLINET
 message 2-65

NMRML
 formatted command 2-61
 message 2-64
 unformatted command 2-60

NMRMSG
 formatted command 2-62
 message 2-65
 unformatted command 2-60

NMRMSG contents 2-63

NMRROUT 2-58
 formatted command 2-62
 message 2-64
 unformatted command 2-60

NMRPRIO
 formatted command 2-61
 message 2-64
 unformatted command 2-59

NMRRMT
 message 2-65

NMRROUT
 message 2-65

NMRTO
 formatted command 2-62
 message 2-64
 unformatted command 2-60

NMRTONOD
 formatted command 2-62
 message 2-64
 unformatted command 2-60

NMRTOQUL

- formatted command 2-62
- message 2-64
- unformatted command 2-60

NMRTYPE

- formatted command 2-61
- message 2-64
- unformatted command 2-60

NMRYPED

- formatted command 2-61
- message 2-64
- unformatted command 2-60

NMRYPEF

- formatted command 2-61
- message 2-64
- unformatted command 2-60

NMRYPET

- formatted command 2-61
- message 2-64
- unformatted command 2-60

NMRYPEX

- formatted command 2-61
- message 2-64
- unformatted command 2-60

NMRYPE4

- formatted command 2-61
- message 2-64
- unformatted command 2-60

NMRUCM

- message 2-64

NMRUCMA

- message 2-65

NMRUSER

- message 2-65

node

- NJE primary
 - determining 5-3
- NJE secondary node
 - definition 5-3

nodes

- addressing 1-6

non SNA buffer format 3-2

notification 2-5, 2-6

- user
 - definition of 1-5

Null Buffers A-1

O

OPTB

- structure definition 2-41

OPTIMIZE 2-45

originating

- node
 - definition of 1-3
- user
 - definition of 1-5

P

PAGEDEF 2-43

Path Manager A-1, A-10

Path Manager Records A-1

PIMSG 2-43

presentation services

- introduction 1-10

products

- supported for these protocols 1-2

protocols

- allowed by network facilities 1-7
- supported by the network facility 1-7

PRTY 2-43

R

receive

- definition of 1-3

record control byte 3-9

record identifier (RID)

- with SNA records 3-15

records

- connection control 2-67

RESERVED 2-19, 2-20, 2-52, 2-53

reset signon record

- format 2-69

response signon record

- format 2-68

S

satellite 6-4

SHEETS 2-45

signon 6-6

- concurrency feature 5-2

- external interface 5-2

- implementation example 5-2

- protocol description 5-3

- with NPM 5-3

- without NPM 5-3

signon record, initial

- format 2-67

SNA A-1

- buffer fields 3-14

- transmission buffer 3-13

SNA data

- data compression 3-13

SOH ENQ 6-4, 6-8

Spanned A-1

spanned data

- record format 3-18

SRCB

- I type 5-3

- J type 5-3

store-and-forward function

- definition of 1-3

Streams A-1

- string control byte 3-8
- string control byte (SCB)
 - with SNA buffer 3-14
- string control byte(SCB)
 - introduction 1-11
- STX 6-6
- subrecord control byte (SRCB)
 - description 3-11
- SYN NAK 6-8
- SYSAREA 2-44

T

- target
 - node
 - definition of 1-3
- TERMSESS 5-8
- THRESHLD 2-44
- TRACE 2-45
- transmission buffer
 - compressing 1-11
 - definition of 1-10
 - format of
 - CTC and BSC 1-11
 - preparing 1-11
 - preparing for transmission 1-10
 - size of
 - for BSC or CTC 3-2
 - for SNA 3-13
 - SNA 1-13, 3-13
- transmit
 - definition of 1-3
- types of nodes
 - definitions 1-3

U

- UCW 6-8
- unspanned data record format 3-17

V

- VERIFY 2-45

W

- Wait-A-Bit A-1
- work
 - NJE job 1-2
 - NMR 1-2
- WRITER 2-43




Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES



BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department D58, Building 921-2
PO BOX 950
POUGHKEEPSIE NY 12602-9935



Fold and Tape

Please do not staple

Fold and Tape



Program Number

5740-XYS
5740-XC6
5685-001
5740-XYN
5665-291
5685-002
5749-DMT
5664-188
5666-273

File Number
S370-38

SC23-0070-0

