



IBM

International Systems Centers

SNA ENA and Virtual Storage Guide

MAR 17 1986

SNA ENA and Virtual Storage Guide

Document Number GG24-1664-00

November 1984

International Systems Center
Raleigh

This document offers guidelines and considerations in planning for the SNA ENA (Extended Network Addressing) and VSCR (Virtual Storage Constraint Relief) offered by ACF/VTAM V3 and ACF/NCP V4.

It also considers VTAM's use of Virtual Storage in the MVS/XA and MVS/SP environments.

Although it describes many new functions offered by ACF/VTAM V3 and ACF/NCP V4, it does not provide complete descriptions of the products.

This guide is intended to assist Communications Network Planners, Systems Programmers and IBM Systems Engineers when planning for the use of ENA and/or VSCR.

CSYS LSYS AIDS

(200 pages)

First Edition (November 1984)

This document applies to ACF/VTAM V3 and ACF/NCP V4 primarily.

It is possible that this material may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

The information contained in this document has not been submitted to any formal IBM test and is distributed on an 'as is' basis **without any warranty either expressed or implied**. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the ability of customers to evaluate and integrate them into their operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Performance data contained in this document was determined in a controlled environment; and therefore, the results which may be obtained in other operational environments may vary significantly. Users of this document should verify the applicable data to their own environment.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used.

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 reader's comments is provided at the back of this document. If the form has been removed, comments may be addressed to:

IBM International Systems Center, D986/B622-3,
P.O. Box 12195,
Research Triangle Park, N.C. 27709
U.S.A.

© Copyright International Business Machines Corporation 1984

Acknowledgements

The authors of this document are:

- Rick Beatty, IBM Canada
- Vj Gopal, IBM Australia

The advisor on this project is:

- Bernd Kampmann, IBM International Systems Center, Raleigh

With gratitude, the contributions of the following persons are acknowledged:

- Wilhelm Gardt and Edi Kaufmann: who established the project.
- Marie Alberts, Glynn Furr and Fred Ostrowski: who reviewed and advised on ENA.
- Larry Corbet and Mark Ross: who reviewed and advised on storage estimation.
- Tom Hackett: who reviewed and advised on ACF/VTAM V3.
- Pat Walker: who provided an invaluable analysis on VTAM's use of virtual storage.

This document is intended to assist in planning for the use of the SNA Extended Network Addressing (ENA) and Virtual Storage Constraint Relief (VSCR) functions. These functions are provided by ACF/VTAM V3 and ACF/NCP V4. It was written primarily for networks with MVS operating systems. It also describes the new features of ACF/VTAM V3 and ACF/NCP V4.

The reader should refer to official documentation (i.e., IBM SRL publications) when evaluating, planning or implementing these functions. This will provide the required level of detail and help ensure up-to-date information.

The examples in this document do not cover all possible combinations of requirement, design, planning, implementation and operation.

The reader is assumed to have a thorough understanding of IBM SNA, familiarity with related products and experience in planning for SNA networks.

The document is structured as follows:

"INTRODUCTION"

- 1.0 "Business Requirements": Discusses the requirements of very large SNA networks and their impact on resource addressing and virtual storage.
- 2.0 "Providing SNA Solutions": Describes the technical requirements that arise with SNA networks in order to satisfy the business requirements discussed in the previous chapter.

"PRODUCT ENHANCEMENTS"

- 3.0 "MVS/XA Extended Virtual Storage Exploitation": Describes the exploitation of the MVS/XA 31-bit addressing capability by ACF/VTAM V3.
- 4.0 "Extended Network Addressing (ENA)": Describes the extended network addressing capability provided by ACF/VTAM V3 and ACF/NCP V4. This section also discusses the differences between the previous address structure and ENA. In addition, various session flows are provided to illustrate the difference between pre-ENA and ENA node communication.
- 5.0 "Other Enhancements": Describes the enhancements in ACF/VTAM V3 and ACF/NCP V4 that are not related to ENA or VSCR. However, these should be considered when planning the installation of ACF/VTAM V3 and ACF/NCP V4.

"PLANNING"

- 6.0 "Planning and Migration Considerations": Describes the considerations in:
 - determining if ENA and VSCR are required; and
 - planning for ENA and VSCR exploitation.

- 7.0 "Migration Scenarios": Describes a number of pre-ENA to ENA migration scenarios. Not all possible combinations are considered. However, there are ten scenarios which should illustrate the principles required to plan for most migrations to ENA capable networks.

"VIRTUAL STORAGE INFORMATION"

- 8.0 "Overview of Virtual Storage and Addressing" on page 101: Provides basic concepts on Virtual Storage and Addressing in the MVS/370 and MVS/XA environments.
- 9.0 "Storage Management Services" on page 119: Provides an introduction to VTAM's Storage Management Services.
- 10.0 "Virtual Storage in VTAM" on page 127: Provides technical information on VTAM Virtual Storage Allocation, and Use in the MVS environment.
- 11.0 "VTAM Virtual Storage Estimation" on page 149: Provides technical information on VTAM Virtual Storage Estimation.
- 12.0 "Monitoring Virtual Storage" on page 159: Provides an approach to monitoring Virtual Storage.

The following conventions are used in this document:

ENA	Extended Network Addressing - the process of using 8 bits to address subareas and 15 bits to address elements in an SNA network
ENA_NCP	an IBM 3725 communications controller node that contains ACF/NCP V4 and therefore is ENA capable
ENA NETWORK	an SNA network where all host nodes contain ACF/VTAM V3 and all communications controller nodes contain ACF/NCP V4
ENA_SSCP	an SSCP in a host node that contains ACF/VTAM V3 and therefore is ENA capable
Pre-ENA_NCP	an IBM 3705/3725 communications controller with ACF/NCP V2 or ACF/NCP V3
Pre-ENA_SSCP	an SSCP in a host node that contains ACF/VTAM V2R1, ACF/VTAM V2R2, ACF/TCAM V2R3 or ACF/TCAM V2R4
VSCR	Virtual Storage Constraint Relief - the use by ACF/VTAM of virtual storage beyond the 16 megabyte address in host nodes operating under MVS/XA

This document is the result of a residency project conducted at the IBM International Systems Center, Raleigh.

The following publications are referenced and/or related:

ACF/VTAM V3

VTAM Installation and Resource Definition	SC23-0111
VTAM Customization	SC23-0112
VTAM Operation	SC23-0113
VTAM Messages and Codes	SC23-0114
VTAM Programming	SC23-0115
VTAM Diagnosis Guide	SC23-0116
VTAM Diagnosis Reference	SC23-0117
VTAM Data Areas	SC23-0121
VTAM Reference Summary	SC23-0135

ACF/NCP V4 and SSP V3

NCP and SSP Messages and Codes	SC30-3169
NCP and SSP Migration	SC20-3252
NCP and SSP Installation and Resource Definition Guide	SC30-3253
NCP and SSP Resource Definition Reference	SC30-3254
NCP and SSP Diagnosis Guide	SC30-3255
NCP Diagnosis Reference	LY30-5569
SSP Diagnosis Reference	LY30-5564
NCP and SSP Reference Summary and Data Areas	LY30-5570
NCP and SSP Customization	LY30-5571

CCP

CCP Users Guide	SC30-3261
CCP Diagnosis and Installation	SC30-3262

Network Program Products

Network Program Products General Information	GC23-0108
Network Program Products Bibliography and Master Index	GC23-0134
Network Program Products Planning	SC23-0110
Network Program Products Samples	SC23-0159
NCCF Installation and Resource Definition	SC27-0660
NCCF Customization: Command Lists	SC27-0661
NCCF Operation	SC27-0664
NCCF Reference Summary	SC27-0039
NCCF Customization: Exit Routines, Command Processors	SC27-0662
NCCF Diagnosis	SY30-5562
NLDM Diagnosis Guide	SC30-3166
NLDM Installation and Operation	SC30-3165
NPDA Users Guide	SC34-2112
NPDA Recommended Action Guide	SC34-2113
NPDA User Reference	SC34-2114
NPDA Messages	SC34-2115
NPDA Installation	SC34-2117
NPDA Diagnosis	SC34-2130

This document is intended to assist in planning for the use of the SNA Extended Network Addressing (ENA) and Virtual Storage Constraint Relief (VSCR) functions provided by ACF/VTAM V3 and ACF/NCP V4. It was written primarily for networks with MVS/XA operating systems.

This document describes the major new features of ACF/VTAM V3 and ACF/NCP V4.

ENA provides the resource addressing capabilities required by very large networks. Prior to ACF/VTAM V3 and ACF/NCP V4, network addressing utilized a total of 16 bits. There was also an accompanying trade-off between subarea and element addressing. With ENA, a network may contain up to 255 subareas and 32,767 elements in each subarea. This is achieved through exploitation of the FID4 Transmission Header introduced in SNA 4.2. In addition, some of the SNA Control Vectors are different.

Within this document, the changes are discussed and described through a number of network operating scenarios. These scenarios also consider the coexistence of pre-ENA nodes within the ENA network:

- SSCP-SSCP Session Activation
- SSCP-PU Session Activation
- LU-LU Session Activation
- ER/VR Activation

The issues in migrating to exploit ENA are also considered. This is illustrated with migrational scenarios that correspond to the following objectives:

- To increase the number of channel-attached terminals.
- To increase the number of host applications.
- To migrate hosts to ACF/VTAM V3.
- To migrate NCPs to ACF/NCP V4.
- To migrate a configuration with CMC host(s).
- To migrate an application host only.
- To migrate SNI connected networks.
- To retain pre-ENA nodes as IRNs within a network.
- To attempt backup of an ENA SSCP with a pre-ENA SSCP.

VSCR provides the virtual storage required for the definition of large networks. ACF/VTAM V3 exploits the 31-bit storage addressing offered by MVS/XA (which allows a problem program in MVS to exploit up to 2 gigabytes of virtual

storage). Most of VTAM's control blocks, buffers and modules have been moved above the 16 megabyte virtual storage boundary. This removes many of the virtual storage constraints experienced today in large networks.

The document examines the extent of VSCR offered by ACF/VTAM V3. It also provides a means of assessing the impact of VSCR at any given installation. The following virtual storage topics are covered:

- Overview of virtual storage and addressing in the MVS environment.
- VTAM Storage Management Services.
- VTAM data areas and Use.
- Estimating virtual storage.
- Monitoring virtual storage.

In addition to ENA and VSCR, other new functions in ACF/VTAM V3 are described:

- ITLIM enhancements
- VTAMOBJ elimination
- Automatic SSCP-SSCP session restart
- SSCP selection exit
- VR selection exit enhancements
- NMVT enhancements
- Elimination of message flooding
- USS enhancements
- API enhancements
- VTAM message enhancements
- Dump formatting enhancements

There is some consideration of changes in ACF/NCP V4.

CONTENTS

INTRODUCTION	1
1.0 Business Requirements	3
2.0 Providing SNA Solutions	5
2.1 Large Network Host Storage	5
2.2 Large Network Addressing	5
PRODUCT ENHANCEMENTS	11
3.0 MVS/XA Extended Virtual Storage Exploitation	13
3.1 VTAM Virtual Storage Constraint Relief (VSCR)	13
3.2 VTAM Application Programs	14
3.3 VTAM Installation Exits	15
3.4 Dumping Improvements	16
4.0 Extended Network Addressing (ENA)	17
4.1 Request/Response Unit Changes	18
4.2 ENA Related Service Aid Changes	18
4.3 Session Activation	18
4.3.1 SSCP-SSCP Session Activation	19
4.3.2 SSCP-PU Session Activation	23
4.3.3 ER/VR Activation	28
4.3.4 LU-LU Session Activation	31
5.0 Other Enhancements	35
5.1 ACF/VTAM V3	35
5.1.1 ITLIM Enhancements	35
5.1.2 Elimination of VTAMOBJ	35
5.1.3 Automatic SSCP-SSCP Session Restart	36
5.1.4 SSCP Selection in Session Management Exit Routine	37
5.1.5 PLU/SLU Names in Virtual Route Selection Exit Routine	38
5.1.6 Network Management Vector Transport (NMVT)	38
5.1.7 Elimination of Message Flooding	38
5.1.8 LUNAME in USS Messages	39
5.1.9 TSO Logon Fix	39
5.1.10 API Enhancements	39
5.1.11 NCP DISPLAY Enhancement	40
5.1.12 NCP DUMP COMPLETE Message Improvement	40
5.1.13 Dump Formatting	41
5.1.14 Symptom String Subset	42
5.1.15 Descriptor Codes	42
5.1.16 NLDM Version 1 Release 3.1 Support	42
5.2 ACF/NCP V4	43
5.2.1 NEO WRAP	43
5.2.2 NCP SYSGEN/NDF	43
5.2.3 Service Aid Changes	44
5.2.4 Show Cause	44
5.2.5 Secondary Network Support	44
5.2.6 Session Information Retrieval (SIR)	44
5.2.7 Dynamic Threshold Alteration	45

5.2.8	Dynamic LPDA	45
5.2.9	Modulo 128 for boundary network nodes (BNN)	45
PLANNING		47
6.0	Planning and Migration Considerations	49
6.1	Identify Requirements	50
6.1.1	Business Requirements	50
6.1.2	Technical Requirements	50
6.1.2.1	Extended Network Addressing	50
6.1.2.2	Extended Virtual Storage	54
6.1.2.3	Other Requirements	58
6.2	High Level Network Design	59
6.2.1	General Design Guidelines	59
6.2.2	Migration Recommendations	60
6.2.3	Network Element Manipulation	61
6.2.4	ACF/VTAM V3 - Operating System Compatibility	62
6.2.5	ACF/VTAM V3 - SSCP and NCP Compatibility	63
6.2.6	Network Program Product Compatibility	64
6.2.6.1	NCP Related Products	65
6.2.6.2	Network Performance Analyzer (NPA)	65
6.2.7	Determine Storage/Cycle Requirements	65
6.3	Low Level Network Design	65
6.3.1	Pre-requisite Maintenance	65
6.3.2	VTAM/NCP Parameter Changes	66
6.3.3	VTAM API Changes	66
6.3.4	Operations Planning Considerations	66
6.3.5	Support Considerations	66
6.3.5.1	Network Planner	66
6.3.5.2	Implementation Coordinator	67
6.3.5.3	Systems Programmer	67
6.3.5.4	Application Programmer	67
6.3.5.5	Network Operator	67
6.3.5.6	Terminal User	68
7.0	Migration Scenarios	69
7.1	Scenario One: Increase Number of Channel-Attached Terminals	70
7.2	Scenario Two: Increase Number of Host Applications	73
7.3	Scenario Three: Migrating Multiple Hosts to ACF/VTAM V3	76
7.4	Scenario Four: Migrating NCPs to ACF/NCP V4	80
7.5	Scenario Five: CMC Host Migration to ACF/VTAM V3 and ACF/NCP V4	82
7.6	Scenario Six: Migrate Application Host Initially	86
7.7	Scenario Seven: SNI and ENA Interconnection	88
7.8	Scenario Eight: ENA/SNI/ENA	91
7.9	Scenario Nine: Pre-ENA Intermediate Node	94
7.10	Scenario Ten: ENA Network/Pre-ENA Backup SSCP	96
VIRTUAL STORAGE INFORMATION		99
8.0	Overview of Virtual Storage and Addressing	101
8.1	31-bit Addressing	101
8.2	Bi-modal Addressing Modes	102
8.3	Virtual Storage	102
8.4	MVS/370 Virtual Storage Areas	103
8.4.1	Common Area	103

8.4.2	Private Area	106
8.4.3	System Area	108
8.5	MVS/XA Virtual Storage Areas - Changes from MVS/370	108
8.5.1	Quantity of Virtual Storage	110
8.5.2	Paging	110
8.5.3	Extensions to Virtual Storage Areas	110
8.5.4	Segment Boundaries	111
8.5.5	REGION, IEALIMIT and IEFUSI	111
8.5.6	Relocation of Virtual Storage Areas	112
8.6	Allocation of Extended Virtual Storage	112
8.7	Fixing Virtual Storage	112
8.8	Storage Protection	113
8.9	Subpools	116
9.0	Storage Management Services	119
9.1	Building and deleting buffer pools	120
9.2	Obtaining and releasing buffers	122
9.3	Expanding and contracting buffer pools	122
9.4	Getting and freeing variable length storage	124
10.0	Virtual Storage in VTAM	127
10.1	Overview of VTAM Data Areas	128
10.2	VTAM Data Area Dictionary	130
11.0	VTAM Virtual Storage Estimation	149
11.1	Introduction	149
11.2	VTAM Virtual Storage Estimation Tables	150
12.0	Monitoring Virtual Storage	159
12.1	Why Monitor Virtual Storage?	159
12.2	An Approach	159
APPENDICES		165
A.0	VTAM Parameter Changes	167
A.1	VTAM Start Options	167
A.2	VTAM Definition Statements	168
B.0	VTAM Application Program Interface (API) Changes	171
B.1	MVS/XA Extended Virtual Storage Exploitation Related	171
B.2	Other Changes	173
B.3	No Change	173
C.0	New VTAM Sense Codes And Messages	175
D.0	Request/Response Unit Changes	177
E.0	SNA Control Vector Changes	179
F.0	SSP Enhancements	181
F.1	NCP Dump Formatter Enhancements	181
F.2	Configuration Report Program (CRP)	182
G.0	Example of Symptom String Subset Data	183

LIST OF ILLUSTRATIONS

Figure 1. SNA Subarea/Element Addressing 6
Figure 2. Pre-ENA Addressing Range 7
Figure 3. ENA Addressing Range 9
Figure 4. ENA Addressing Structure 17
Figure 5. Pre-ENA_SSCP;Pre-ENA_SSCP Session Activation 19
Figure 6. Pre-ENA_SSCP;ENA_SSCP Session Activation 20
Figure 7. ENA_SSCP; Pre-ENA_SSCP Session Activation 21
Figure 8. ENA_SSCP; ENA_SSCP Session Activation 22
Figure 9. Pre-ENA_SSCP; Pre-ENA_NCP ACTPU Flow 23
Figure 10. ENA_SSCP; Pre-ENA_NCP ACTPU Flow 24
Figure 11. ENA_SSCP; Pre-ENA_NCP; ENA_NCP ACTPU Flow 25
Figure 12. Pre-ENA_SSCP; ENA_NCP; Pre-ENA_NCP ACTPU Flow 26
Figure 13. ENA_SSCP; ENA_NCP ACTPU Flow 27
Figure 14. ER/VR Activation - ENA Network 28
Figure 15. ER/VR Activation - Pre-ENA/ENA Network 29
Figure 16. ER/VR Activation - Pre-ENA INN 30
Figure 17. LU-LU: ENA_SSCP; ENA_NCP; ENA_NCP; ENA_SSCP 31
Figure 18. LU-LU: ENA_SSCP; ENA_NCP; Pre-ENA_NCP; Pre-ENA_SSCP 32
Figure 19. LU-LU: Pre-ENA_NCP; ENA_SSCP; ENA_NCP; ENA_SSCP 33
Figure 20. LU-LU: Pre-ENA_NCP; ENA_SSCP; ENA_NCP; ENA_SSCP 34
Figure 21. Major Planning Activities 49
Figure 22. Pre-ENA Addressing Structure (16-bit Network Addresses) 51
Figure 23. SNA Subarea/Element Planning Chart 53
Figure 24. VTAM Virtual Storage Planning Chart 56
Figure 25. Sequencing Activation List (One) 61
Figure 26. Sequencing Activation List (Two) 62
Figure 27. ACF/VTAM V3 - Operating System Compatibility Chart 62
Figure 28. ACF/VTAM V3 SSCP and NCP Compatibility Chart 63
Figure 29. Network Program Product Compatibility Chart 64
Figure 30. Expanding number of Local Terminals 71
Figure 31. Expanding Applications Sessions 74
Figure 32. Migrating Multiple Hosts (Before) 77
Figure 33. Migrating Multiple Hosts (After) 78
Figure 34. Migrating NCP to ENA (Before) 80
Figure 35. Migrating NCP to ENA (After) 81
Figure 36. CMP and NCP Migration (Before) 83
Figure 37. CMP and NCP Migration (After) 84
Figure 38. Application Host Migration (Before) 86
Figure 39. Application Host Migration (After) 87
Figure 40. Use of SNI and ENA (Before) 89
Figure 41. Use of SNI and ENA (After) 90
Figure 42. Co-existence of ENA + SNI (Before) 92
Figure 43. Co-existence of ENA + SNI (After) 93
Figure 44. IRN Configuration (Before) 94
Figure 45. IRN Configuration (After) 95
Figure 46. Complete ENA Configuration 96
Figure 47. MVS Addressing Formats 101
Figure 48. MVS/370 Virtual Storage Layout 104
Figure 49. MVS/XA Virtual Storage Layout 109
Figure 50. MVS Virtual Storage Subpools 0-238 114
Figure 51. MVS Virtual Storage Subpools 239-255 115

Figure 52. Virtual Storage Area - Subpool Relationships	118
Figure 53. SMS Data Area Relationships after BLDPOOL	121
Figure 54. SMS Data Area Relationships after Pool Expansion	123
Figure 55. SMS Variable Length Storage Pools	125
Figure 56. ACF/VTAM Data Areas	129
Figure 57. VTAM Data Area - MVS Subpool Cross Reference	131
Figure 58. ERT and ERTE Relationships	137
Figure 59. RMF Virtual Storage Activity Report.....Page 1.	160
Figure 60. RMF Virtual Storage Activity Report.....Page 2.	161
Figure 61. RMF Virtual Storage Activity Report.....Page 3.	162
Figure 62. RMF Virtual Storage Activity Report.....Page 4.	163
Figure 63. CSA Usage Activity Report (produced by user program)	164
Figure 64. VTAM API Exits - Addressing Modes	171
Figure 65. Control Vector X'06'	179
Figure 66. Control Vector X'0B'	180

In this part, the sections are:

- 1.0 "Business Requirements" on page 3: Discusses the requirements of very large SNA networks and their impact on resource addressing and virtual storage.
- 2.0 "Providing SNA Solutions" on page 5: Describes the technical requirements that arise for SNA networks in order to satisfy the business requirements discussed in the previous chapter.

1.0 BUSINESS REQUIREMENTS

The role of Information Systems in many business enterprises has undergone considerable change. It was typically an expensive 'back room' operation that failed to demonstrate its value to the enterprise. With advances in the understanding of computer technology, the role has assumed a new significance. It is now seen as an essential and vital part of the business capable of making significant contributions to the objectives of profitability and increasing market share.

The exchange of information is an integral part of business processes. Computer technology has revolutionized these processes.

The computer terminal has increased the availability of service outlets where transactions may be initiated. The networks (that connect these terminals to host systems) make it viable and efficient to transport information over vast distances. The centralized host systems make it possible and economical to process large volumes of transactions. This technology allows the enterprise to expand its markets in a cost-effective way without losing control over the processing and storage of vital business information. The barriers imposed by geographical dispersion and remoteness have been diminished, opening the way for new markets beyond state and national boundaries.

This leads to the creation of large terminal networks. In many enterprises, this growth is accelerated with the thrust to automate existing corporate information and customer services. New requirements, such as credit card authorization and electronic funds transfer transactions from point-of-sale terminals at retail outlets, will lead to significant growth in the network and its terminals.

Large networks exist for other reasons as well. Technology allows the consolidation of independent telecommunications networks (i.e., voice, data, image, file) within an enterprise. The consolidation of two or more networks, as a result of a corporate merger or takeover, will cause significant expansion.

Network consolidation (i.e., the creation of a single network by rationalizing two or more networks) can be achieved with IBM SNA Multi System Networking (MSNF). The cost in replicating functions and services at each business location can be avoided. The redundancy required in telecommunications circuits and equipment can be minimized. The cost and complexity of managing and operating multiple networks will be avoided.

The large single network and centralized transaction processing facilities combine to provide the business enterprise with economies of scale and levels of control and security that are difficult to achieve otherwise. The unprecedented rate of growth in large networks has introduced these two significant technical requirements:

1. The increase in the number of terminals and applications has increased the demand for storage by the host telecommunications access method.
2. The current structure of SNA network addresses has to be expanded to accommodate this growth.

2.0 PROVIDING SNA SOLUTIONS

ACF/VTAM V3 and ACF/NCP V4 provide solutions to a large network's requirements in the areas of host storage and network addressing.

2.1 LARGE NETWORK HOST STORAGE

Prior to ACF/VTAM V3, VTAM's use of virtual storage was based on the S/370 Architecture. That architecture specified the use of 24 bits for storage addressing, which provided a theoretical virtual storage addressing capability of up to 16 megabytes. This meant that the VTAM address space could only contain a maximum of 16 megabytes of modules, buffers and control blocks.

In reality, the amount of virtual storage available to the VTAM was far less. The actual amount available to a particular installation would depend on how much storage was consumed by MVS system components and its subsystems (see "Overview of Virtual Storage and Addressing" on page 101). This could range between 6-10 megabytes.

Given that VTAM's requirements for control block and buffer storage increases proportionally with the number of network resources, several large networks soon exhausted the virtual storage available. This constrained the ability to add network resources. It also affected the ability to recover from the concurrent failure of several resources. There have been attempts at circumventing these constraints (e.g., distributing network resource ownership over several hosts). However, the measures taken were often costly, limited in effectiveness and accompanied by operational and other problems.

In the S/370 Extended Architecture and its MVS implementation, virtual storage addressing was extended to 31 bits. This provided VTAM with the potential to use up to 2 gigabytes (i.e., 128 x 16 megabytes) of virtual storage for its modules, control blocks and buffers. This will remove the current virtual storage constraints and allow the implementation of very large SNA networks.

The growth in subsystems and user application programs that use the VTAM Application Program Interface (API) has also led to an increase in virtual storage demands.

ACF/VTAM V3 exploits the extended virtual storage in MVS/XA and provides this capability to programs that use the VTAM API.

2.2 LARGE NETWORK ADDRESSING

Prior to ACF/VTAM V3 and ACF/NCP V4, the SNA addressing structure had a fixed length of sixteen bits. The two bytes were divided into two portions:

Regardless of where the division is drawn between the subarea and element, the theoretical maximum number of addresses available in each SNA network was 65,536. However, due to other factors which dictate the required number and placement of communications controllers and hosts, this range still remains somewhat theoretical for most networks!

SUBAREA PORTION		ELEMENT PORTION	
Bits Used	Available Subareas	Bits Used	Available Elements
8	255	8	255
7	127	9	512
6	63	10	1024
5	31	11	2048
4	15	12	4096
3	7	13	8192
2	3	14	16384

Figure 2. Pre-ENA Addressing Range

The interdependency between the subarea range required and the resulting number of network elements per subarea has constrained network design in many enterprises. In many organizations, this interdependency has adversely affected:

1. End User Productivity

- To avoid the response delays incurred by accessing a system remotely, many installations channel-attach terminals into the application host.
- This reduces the number of elements available to that host subarea. Further growth (over time) in the number of channel-attached terminals and applications reduces element numbers to critical levels.
- This compels the installation to satisfy latent terminal demand with remote lines, install additional host processors and/or align the sub-area/element address split in favor of the 'element'.

2. Remote Network Growth

- Business growth (i.e., new branches, new business territories) requires the installation of additional lines, control units and terminals.
- Each terminal requires an element address. The constraints imposed by the subarea/element addressing structure often forces more communications controllers, lines, control units, and terminals than would otherwise be necessary.

3. Consolidation of Existing SNA Networks

- Two large enterprises have merged.
- Both organizations have large SNA networks which are inherently different (i.e., different subarea/element splits, naming conventions).
- The networks were to be interconnected and gradually consolidated as part of a business strategy to reduce the corporate telecommunications operating expense.
- The subarea range required imposes a subarea/element split that provides an undesirably low number of 'elements per subarea'.
- Although SNI (see below) would alleviate some of these problems, it would also introduce additional levels of complexity.

SNA Network Interconnect (SNI)

Prior to ACF/VTAM V2R2 and ACF/NCP V3, interconnecting multiple SNA networks was possible if the subarea/element address split was identical across the networks. However, interconnection standards (such as network resource names) were required to preserve the integrity of each network (e.g., by avoiding duplicate names and subareas).

ACF/VTAM V2R2 and ACF/NCP V3 introduced the SNA Network Interconnect (SNI) facility. SNI provides an SNA gateway between multiple SNA networks with different attributes (i.e., subarea/element splits, naming standards). SNI was designed to:

- address the inability to connect multiple SNA networks that have different subarea/element address splits.
- provide a solution for an installation that has exhausted the subarea addressing range as a result of a large number of interconnected hosts and communications controllers. SNI split the network into multiple logical networks and provided a means of mapping between each other.
- the coexistence of multiple networks - each with different network naming standards. The SNI ALIAS Mapping function allows duplicate resource names to exist within each network by translating the duplicate to a unique name during cross-network communication.

SNA Extended Network Addressing (ENA)

The SNA addressing structure has been expanded with the SNA Extended Network Addressing (ENA) feature. ENA is available in ACF/VTAM V3 and ACF/NCP V4 to address:

- the requirements of installations that would like to evolve the SNA network as a single logical entity; and
- generally, the requirement to define a larger number of resources in each subarea (i.e., host processor, communications controller).

ENA eliminates the direct dependency between the selection of the subarea range and the number of network elements available.

- Eight bits are used for the subarea address.
- Fifteen bits are used for the element address.
- ENA will allow up to 255 subareas (0 is reserved) to coexist within the same network.
- Each subarea has the potential of addressing up to 32,767 network elements (e.g., terminals, applications, lines, cluster controllers, etc.).

Theoretically, the maximum number of addressable resources in an ENA-capable network is 8,355,840. However, requirements in processing power, main storage, DASD, geographic placement of communication controllers and facilities constraints (such as floor space) will all combine to determine the actual number of network resources possible in any given production environment.

SUBAREA PORTION		ELEMENT PORTION	
Bits Used	Available Subareas	Bits Used	Available Elements
8	255	15	32,767

Figure 3. ENA Addressing Range

In this part, the sections are:

- 3.0 "MVS/XA Extended Virtual Storage Exploitation" on page 13: Describes the exploitation of the MVS/XA 31-bit addressing capability by ACF/VTAM V3.
- 4.0 "Extended Network Addressing (ENA)" on page 17: Describes the extended network addressing capability provided by ACF/VTAM V3 and ACF/NCP V4. This section also discusses the differences between the previous address structure and ENA. In addition, various session flows are provided to illustrate the difference between pre-ENA and ENA node communication.
- 5.0 "Other Enhancements" on page 35: Describes the enhancements in ACF/VTAM V3 and ACF/NCP V4 that are not related to ENA or VSCR. However, these should be considered when planning the installation of ACF/VTAM V3 and ACF/NCP V4.

3.0 MVS/XA EXTENDED VIRTUAL STORAGE EXPLOITATION

ACF/VTAM V3 will use the 31-bit virtual storage addressing capability provided by MVS/XA. This capability allows a program to address 2 gigabytes of virtual storage within a single MVS address space; increasing the amount of virtual storage available to VTAM by a multiplication factor of 128 over previous releases.

This change is available only to MVS/XA systems with ACF/VTAM V3 installed.

The chapter "Overview of Virtual Storage and Addressing" on page 101 provides an introduction to the discussion that follows.

3.1 VTAM VIRTUAL STORAGE CONSTRAINT RELIEF (VSCR)

Prior to ACF/VTAM V3, each VTAM address space was limited to 16 megabytes of virtual storage. It was difficult to:

- define and activate a large number of resources (e.g., 6,000 LUs or more) within a single VTAM address space; and
- recover from the failure of large nodes (e.g., a Host or NCP Intermediate Routing Node (IRN) with several hundred resources).

In this release, VTAM can address up to 2 gigabytes of virtual storage within its address space. These are some of the benefits:

- Definition and ownership of very large networks (e.g., more than 10,000 LUs) from a single host.

While there may be some advantage in owning a network from several hosts, 'single host ownership':

- avoids the operational and network management complexity inherent in running a network from multiple 'owners'; and
 - allows the customer to exploit the processing power and economic advantages offered by IBM's largest central processors.
- Avoidance of 'VTAM short of storage' conditions.

The virtual storage shortage occurred, most frequently, during recovery from large node failures. In addition to extending recovery time, the shortage prevented completion of resource 'cleanup'.

Most VTAM control blocks, buffers and modules will be moved from the VTAM private area and the common area into their corresponding extensions above the 16 megabyte virtual storage boundary. Where there are 24-bit addressing dependencies, the data area or module will remain below the boundary.

The major areas that will remain below the boundary are:

- IO Buffers (IOBUF) will remain below because Format 0 CCWs are still used by VTAM. The Format 1 CCW defined in System/370 Extended Architecture has not been implemented in ACF/VTAM V3.
- Modules that:
 - must branch to user exits; or
 - must branch to interfaces that do not exploit 31-bit addressing; or
 - may receive control from modules that do not exploit 31-bit addressing.
- Parameter lists for restricted or hollow interfaces to MVS Systems Services (i.e., services that partially support or do not support 31-bit addressing).
- SSP interfaces.

VTAM will allocate the storage that is usually obtained from an application program's private address space above the boundary. However, any storage that may be referenced by the application program directly will remain below the boundary.

VTAM's use of CSA can be controlled through the CSALIMIT parameter. This has been changed to allow specification of a limit up to 2 gigabytes.

This enhancement alleviates the VTAM virtual storage constraint that has inhibited growth and expansion in large SNA networks.

3.2 VTAM APPLICATION PROGRAMS

In ACF/VTAM V3, programs can be coded to use the VTAM RECORD Application Program Interface (API) under the 24-bit or 31-bit addressing modes in MVS/XA.

Therefore, changes to application programs are only necessary if additional virtual storage is required.

Most of the VTAM API control blocks can reside on either side of the 16 megabyte virtual storage boundary. However, only programs executing in 31-bit addressing mode can access data above the boundary.

The ACB and related storage (i.e., application name and password fields) must remain below the boundary due to the need to maintain compatibility with VSAM.

For similar reasons, the GENCB, MODCB, SHOWCB and TESTCB macro instructions will not manipulate data areas located above the 16 megabyte virtual storage boundary (although these 'macros' may be executed in the 31-bit addressing mode).

VTAM API user exit routines can execute in either the 24-bit or 31-bit addressing mode. The addressing mode in which an exit receives control will be determined as follows:

- Exits specified through the EXLST - the addressing mode used when the ACB was opened.
- RPL exits - the addressing mode used when the original RPL-based request was issued.
- SYNAD/LERAD exits - the addressing mode used when the CHECK request was issued (if OPTCD=SYN is used, this will be the mode of the original RPL-based request for which the CHECK was issued).

VTAM macro instructions can be invoked and executed in either addressing mode. However, the following points are significant:

- The ACB must be in 24-bit storage.
- The parameter lists for OPEN and CLOSE must be in 24-bit storage.
- The manipulative macro instructions can only operate on data areas in 24-bit storage.
- The data areas referenced by any macro instruction must be in storage that is compatible with the program's addressing mode (i.e., a program in 24-bit mode will ignore the high order address byte with unpredictable results).
- VTAM will clear the high order byte of addresses in control blocks used by 24-bit programs.

There are no changes in ACF/VTAM V3 for the following VTAM RECORD API functions:

- Authorized Path
- Parallel Sessions
- Multi-memory Applications
- Programmed Operator Interface (POI)

There are changes to the Communications Network Management Interface (CNMI). As these are due to support for ENA, NMVT and the IBM 3710 (i.e., they do not result from MVS/XA exploitation), they are discussed later.

3.3 VTAM INSTALLATION EXITS

The VTAM Installation Exits are:

- Session Management
- Virtual Route Selection
- VR Pacing Window Size
- Session Accounting

- **Session Authorization**

In ACF/VTAM V3, these can be coded to execute in either the 24-bit or 31-bit addressing mode under MVS/XA.

These routines will be given control in the addressing mode specified through the AMODE attribute of the load module. Existing routines need not be changed; these will run in 24-bit mode.

3.4 DUMPING IMPROVEMENTS

Due to the potential for large virtual storage dumps in 31-bit addressing mode (i.e., 2 gigabytes), ACF/VTAM V3 will use the improved dumping facilities of MVS/XA (specifically MVS/SP V2R1).

VTAM will dump those storage areas which have been allocated in the VTAM key (VTAM uses MVS Supervisor Storage Key 6 for storage protection). The dump will include the following:

- Storage in Key 6
- Subpools 0, 227-231, 239, 241, 245, 252-255
- Storage in Key 0 allocated by VTAM
- ALLPSA

This will limit the dump to private and common storage used by VTAM.

These dumping options are specified via the SDUMP parameter list ('hard coded' in a VTAM module). If it is necessary to dump other areas of storage, these options may be modified through the MVS CHNGDUMP command.

These improvements will reduce the size of dumps created (via SDUMP):

- for a VTAM module abend under a IRB, SRB or TCB in VTAM's address space; or
- under a VTAM FRR/ESTAE routine used by a VTAM application program

while ensuring adequate problem determination information is collected. Dumps created by the Z NET command will not be affected.

4.0 EXTENDED NETWORK ADDRESSING (ENA)

Throughout this document, the term **ENA network** refers to a network of hosts and communications controllers that have been migrated to ACF/VTAM V3 and ACF/NCP V4. The term **ENA node** refers to a host with ACF/VTAM V3 or a communications controller with ACF/NCP V4.

Background

The preparations to extend SNA network addressing started with SNA 4.2 (i.e., ACF/VTAM V1R3 and ACF/NCP V1R3). Additional space was reserved in the FID4 Transmission Header (TH). The FID4 TH introduced in ACF/VTAM V1R3 and ACF/NCP V1R3 contained two non-contiguous network address fields consisting of four bytes for subarea address and two bytes for element addresses.

The Extended Network Addressing function of ACF/VTAM V3 and ACF/NCP V4 utilizes:

- eight bits for subarea addressing; and
- fifteen bits for element addressing.

The following diagram illustrates the change in the SNA addressing structure:

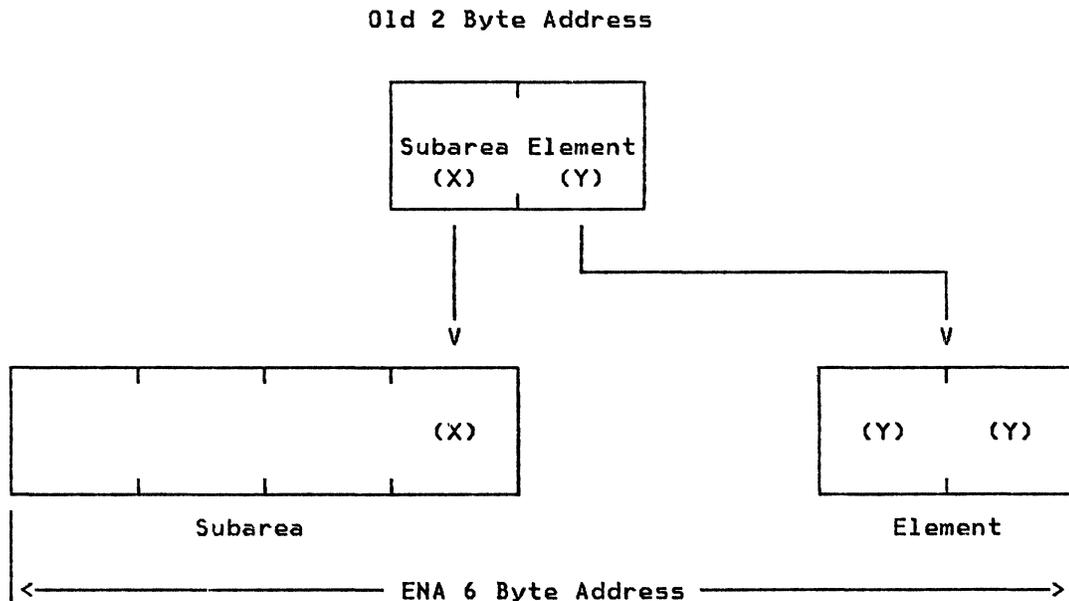


Figure 4. ENA Addressing Structure

4.1 REQUEST/RESPONSE UNIT CHANGES

There have been no new RUs introduced in ACF/VTAM V3 and ACF/NCP V4 for ENA support. To accommodate extending the network address structure from two to six bytes, many of the Request/Response Unit (RUs) have changed.

The change in the RU format - bytes three and four of many of the RUs, in the past, referred to the network address (subarea/element). In ACF/VTAM V3 and ACF/NCP V4, bytes three and four now contain only the element address. The sub-area value of the network address has been removed from the RU.

Please refer to "Request/Response Unit Changes" on page 177 if you are interested in more detail.

4.2 ENA RELATED SERVICE AID CHANGES

VTAM Internal Trace

The VTAM Internal Trace record has been rewritten to allow for the ENA expansion to the network address. In addition, the default options have changed.

The Dump Formatter has been rewritten with new trace formats and a number of dumping improvements have been made.

ACFTAP

The ACFTAP program has been changed to reflect the network address expansion.

4.3 SESSION ACTIVATION

ENA and pre-ENA nodes can coexist within the same network. However, for a pre-ENA/ENA mixed environment to be possible, the ENA node must be able to determine the level of support existing in the adjacent nodes. This is achieved during the ACTPU/ACTCDRM operations and their subsequent responses.

An 'Activate Physical' (ACTPU) is sent by the SSCP to activate a session with the PU and to obtain certain information about the PU. Control Vector X'0B' provides the ACTPU operation, the 'ENA support' indicator.

An 'Activate Cross-Domain Resource Manager' (ACTCDRM) is sent by an SSCP to activate a session with another SSCP. Just as in the case of the ACTPU, ACTCDRM is used to pass information related to the CDRM-CDRM session. Control Vector X'06' provides the ACTCDRM operation, the 'ENA Support' indicator.

"SNA Control Vector Changes" on page 179 describes the layout of Control Vectors X'06' and X'0B'.

After the ENA and pre-ENA support is determined, it is the ENA node's responsibility to transform all RUs into the appropriate format for the adjacent node. The pre-ENA node expects the pre-ENA RU format during the ACTPU and ACTCDRM operations.

4.3.1 SSCP-SSCP Session Activation

The following scenarios present the effect on SSCP activation in pure ENA and mixed (pre-ENA/ENA) environments.

Pre-ENA SSCP; Pre-ENA SSCP

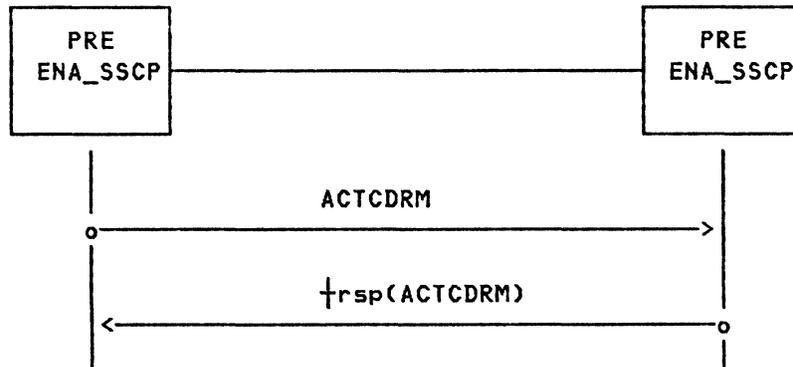


Figure 5. Pre-ENA_SSCP;Pre-ENA_SSCP Session Activation

- Both SSCPs have determined that communication is possible. The exchanges will be performed using the pre-ENA addressing format as each SSCP is a pre-ENA SSCP.

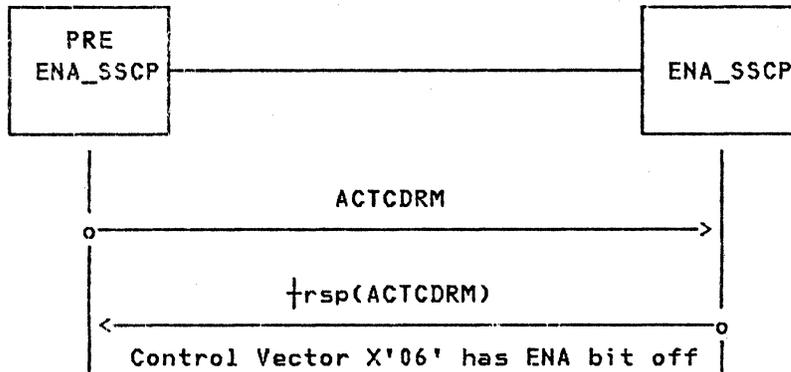


Figure 6. Pre-ENA_SSCP;ENA_SSCP Session Activation

- The SSCPs will communicate using the pre-ENA format (Cross-Domain Session Services) RUs.
- For SSCP-SSCP communication to be possible, the subarea address of the ENA_SSCP must be within the range specified by MAXSUBA.
- For subsequent LU-LU sessions to be possible, the ENA_SSCP LU must be within the addressability range used by the pre-ENA host.

ENA SSCP; Pre-ENA SSCP

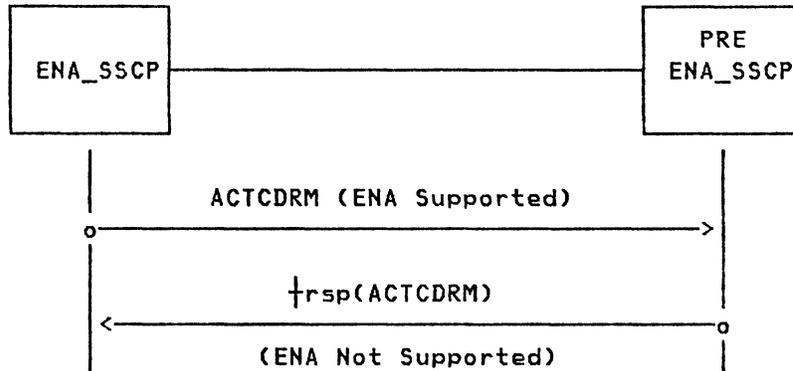


Figure 7. ENA_SSCP; Pre-ENA_SSCP Session Activation

- The CDRM-CDRM session will follow the pre-ENA format.
- Just as in the previous scenario, the ENA resources must be within the address capability of the pre-ENA resources for communication to be possible.

ENA_SSCP; ENA_SSCP

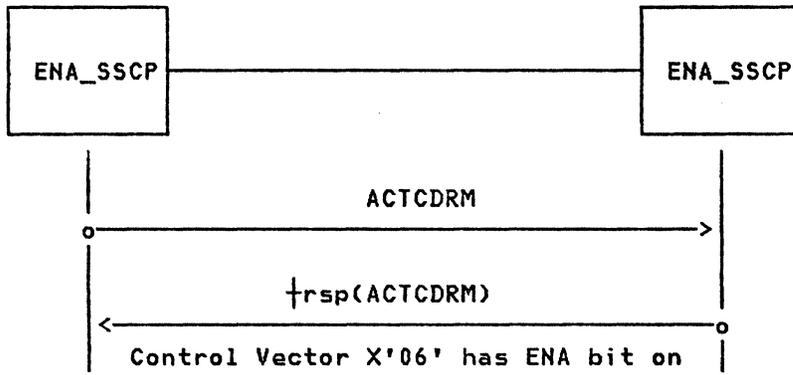


Figure 8. ENA_SSCP; ENA_SSCP Session Activation

- Both hosts support ENA addressing formats.

4.3.2 SSCP-PU Session Activation

During the ACTPU operation, Control Vector X'0B' is used to indicate if either node can support ENA. The subsequent exchange between the SSCP and the NCP determine the level of support possible on the SSCP_PU session.

Pre-ENA SSCP; Pre-ENA NCP

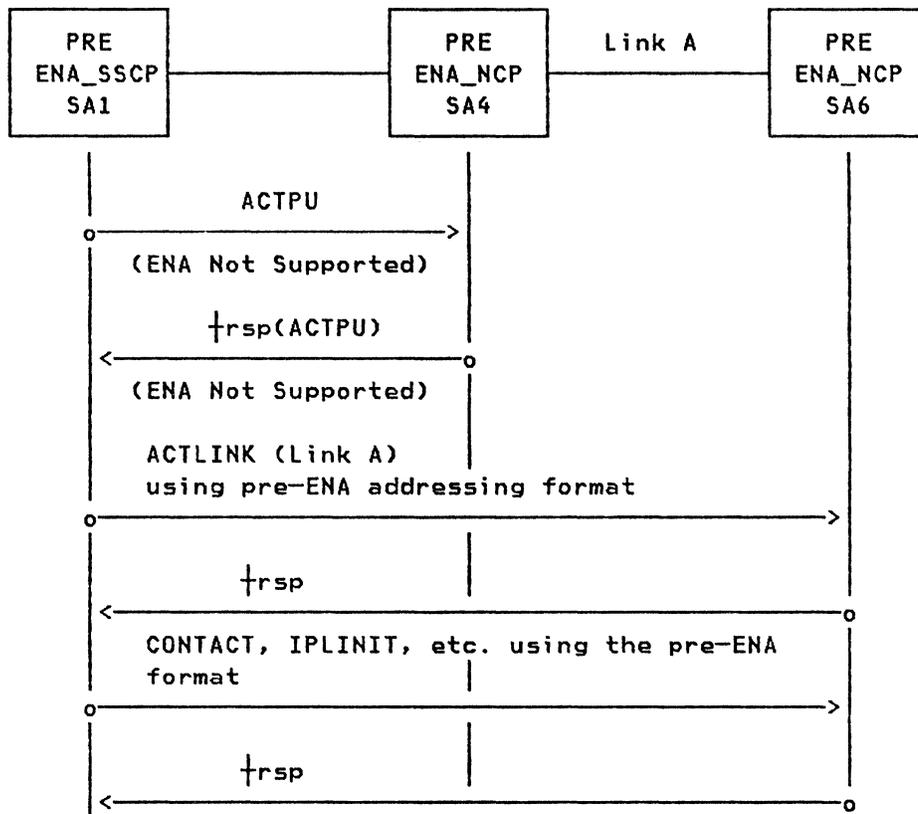


Figure 9. Pre-ENA_SSCP; Pre-ENA_NCP ACTPU Flow

- During the SSCP->PU ACTPU operation, the SSCP passes the Control Vector X'0B' but ENA support is not present in the pre-ENA release of ACF/VTAM.
- The NCP responds positively to the ACTPU and communicates with the SSCP using the pre-ENA format.

ENA_SSCP; Pre-ENA_NCP

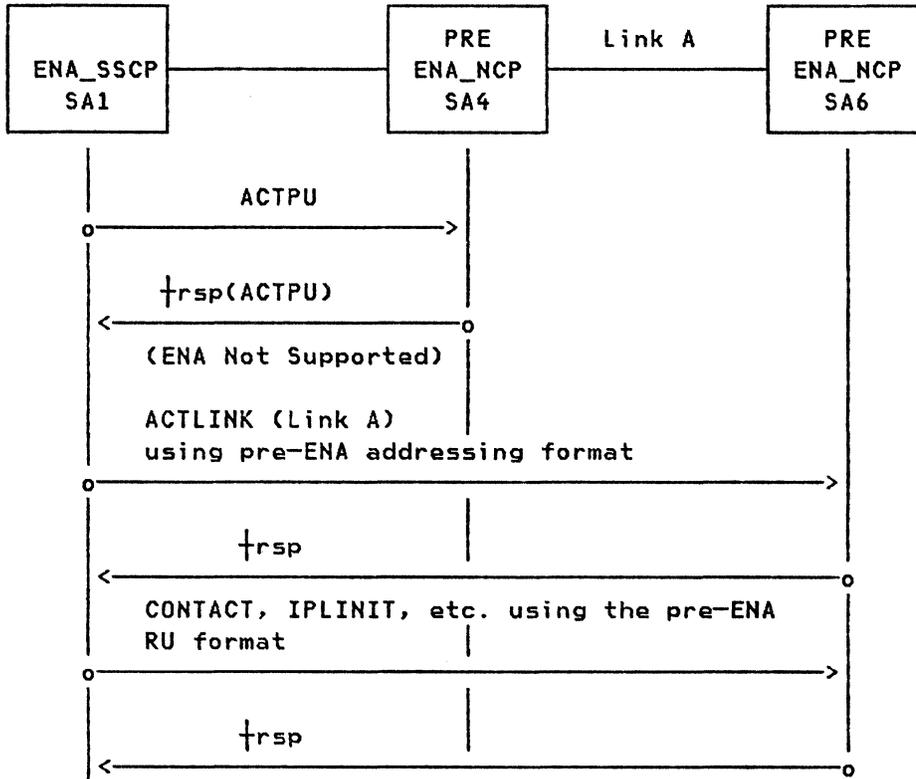


Figure 10. ENA_SSCP; Pre-ENA_NCP ACTPU Flow

- During the SSCP->PU ACTPU operation, the SSCP indicates that 'ENA Support' is present for the host.
- SA4 is at a pre-ENA level; the response to the ACTPU carries Control Vector X'0B' with the 'ENA Support' bit turned off.
- Since the NCP is incapable of supporting ENA, the SSCP must keep the MAXSUBA operand in the ACF/VTAM start deck to avoid an exception response to ACTPU (X'0809 0024').
- Bytes three and four of the ACTLINK and CONTACT contain the pre-ENA element addresses of Link A and the adjacent link station for SA6.
- The SSCP must use element addresses that comply with the subarea/element address range used by NCPs - SA4 and SA6.

ENA_SSCP; Pre-ENA_NCP; ENA_NCP

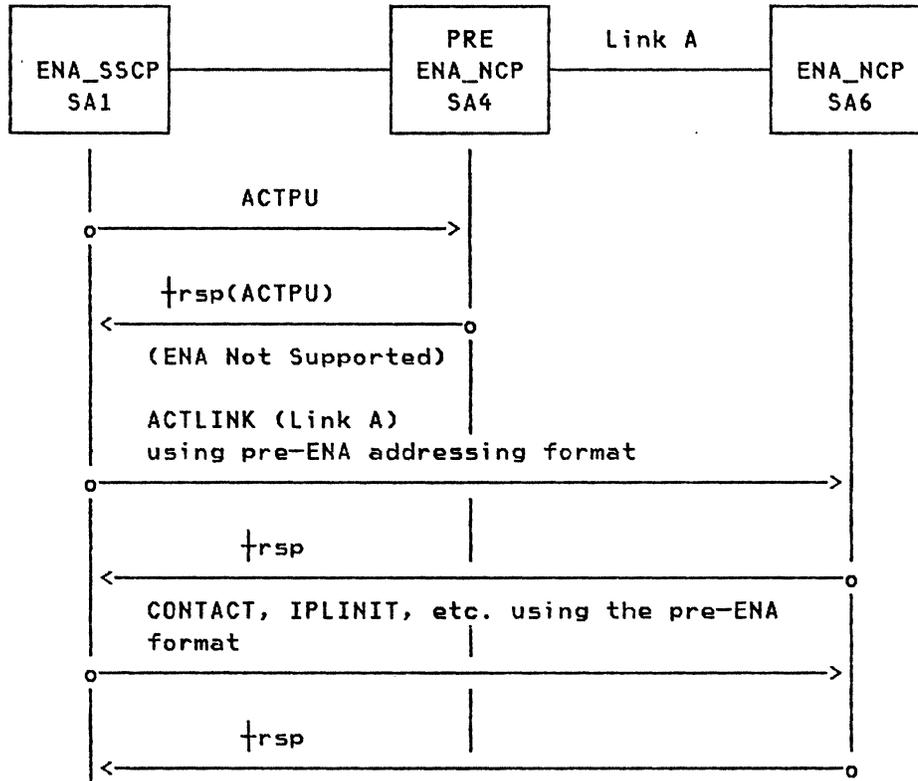


Figure 11. ENA_SSCP; Pre-ENA_NCP; ENA_NCP ACTPU Flow

- During the SSCP->PU ACTPU operation, the SSCP indicates that 'ENA Support' is present for the host.
- SA4 is a pre-ENA_NCP; the response to the ACTPU carries Control Vector X'0B' with the 'ENA Support' bit turned off.
- Since the NCP SA4 is incapable of supporting ENA, the SSCP and NCP SA6 must keep the MAXSUBA operand.
- SSCP and SA6 elements may be assigned higher than the pre-ENA range provided SA4 does not have to communicate with them.

Pre-ENA_SSCP; ENA_NCP; Pre-ENA_NCP

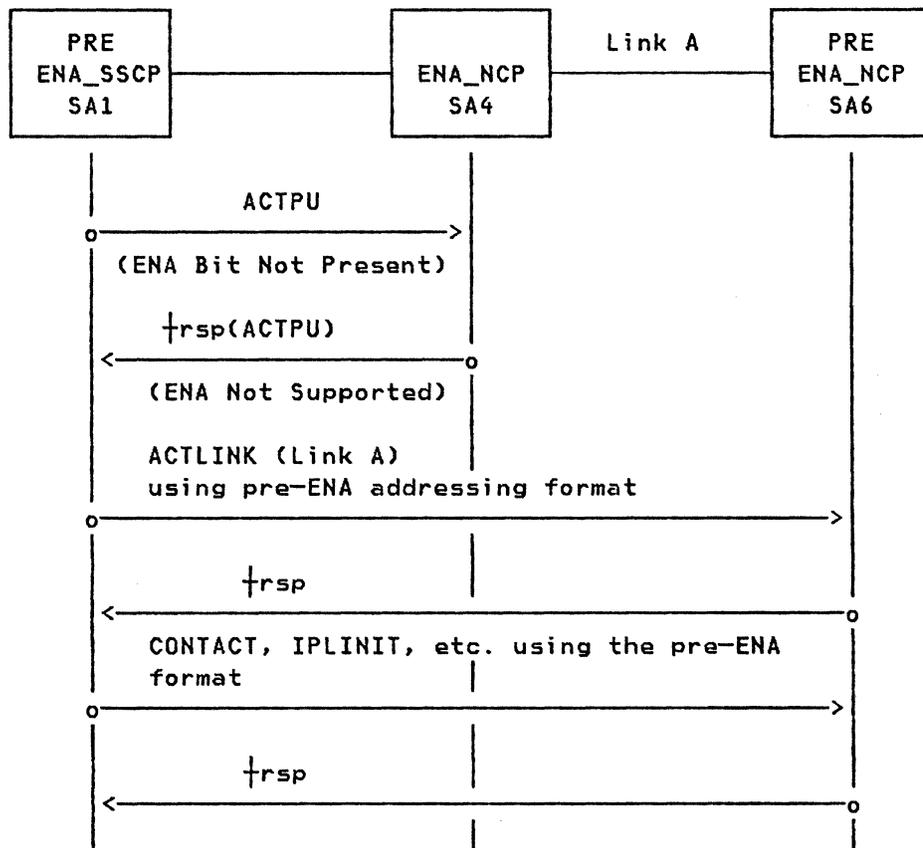


Figure 12. Pre-ENA_SSCP; ENA_NCP; Pre-ENA_NCP ACTPU Flow

- The SSCP is incapable of passing the ENA indicator on the ACTPU operation.
- SA4, responds passing Control Vector X'0B' with the 'ENA Supported' bit set off.
- Resources not conforming to the addressing range will be unknown to VTAM. Consequently, no session will be attempted with those resources. The activation of SA4 would continue normally without them.

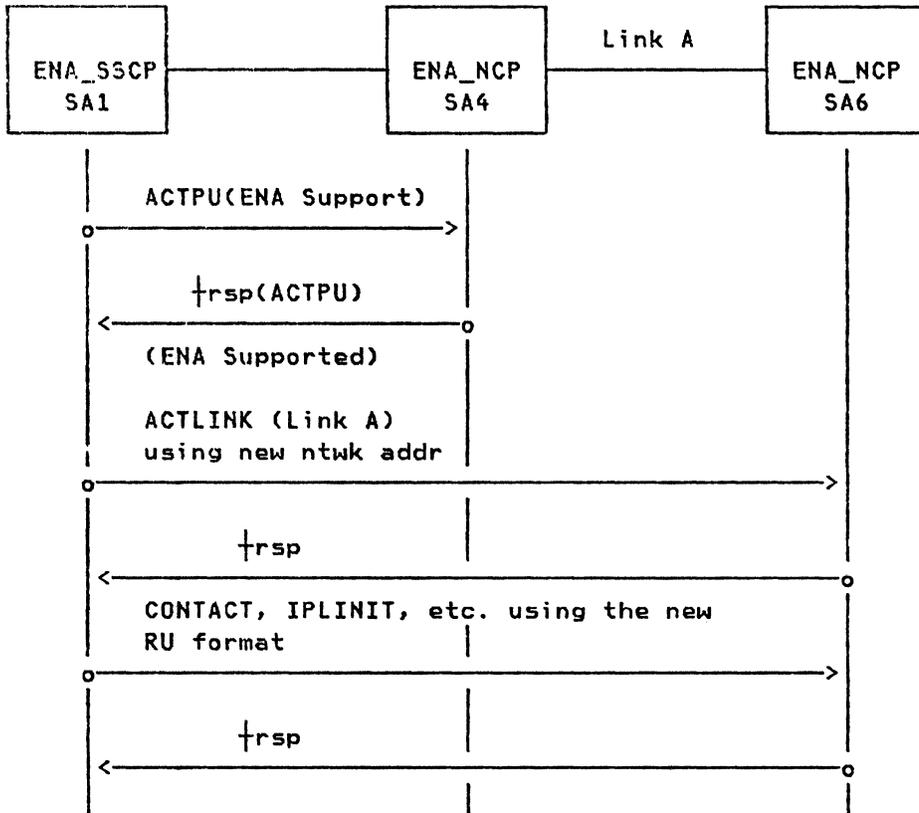


Figure 13. ENA_SSCP; ENA_NCP ACTPU Flow

- The ENA_NCP confirms that it supports ENA by responding with a positive response to the ACTPU/CV-X'0B'.
- The SSCP determines that SA4 does support ENA, and sends 'ENA format RUs'.
- Bytes three and four of the ACTLINK and CONTACT contain the ENA element addresses of Link A and the adjacent link station for SA6.

4.3.3 ER/VR Activation

In a 'pure ENA Network', ER/VR activation is straightforward. In a mixed ENA and pre-ENA environment, ER/VR activation is possible provided the subarea numbers of the end points are within the subarea range of MAXSUBA.

ENA_SSCP; ENA_NCP; ENA_NCP; ENA_SSCP

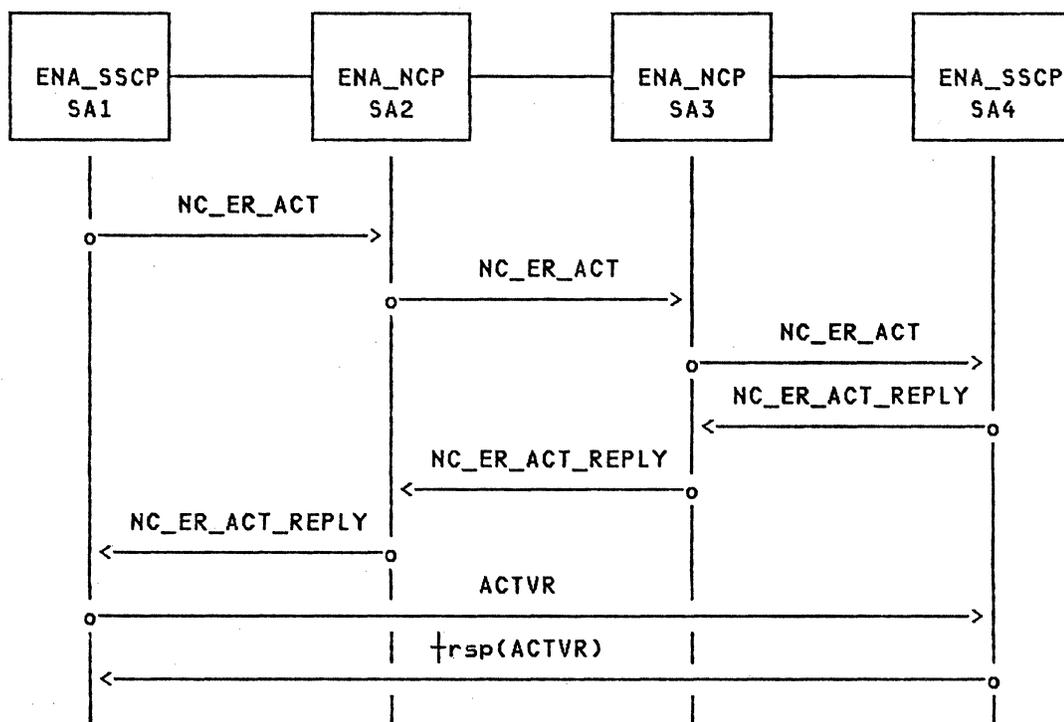


Figure 14. ER/VR Activation - ENA Network

- The ER and VR Activation was successful (shown by the +rsp(ACTVR) in the figure above).

ENA_SSCP; ENA_NCP; Pre-ENA_NCP; Pre-ENA_SSCP

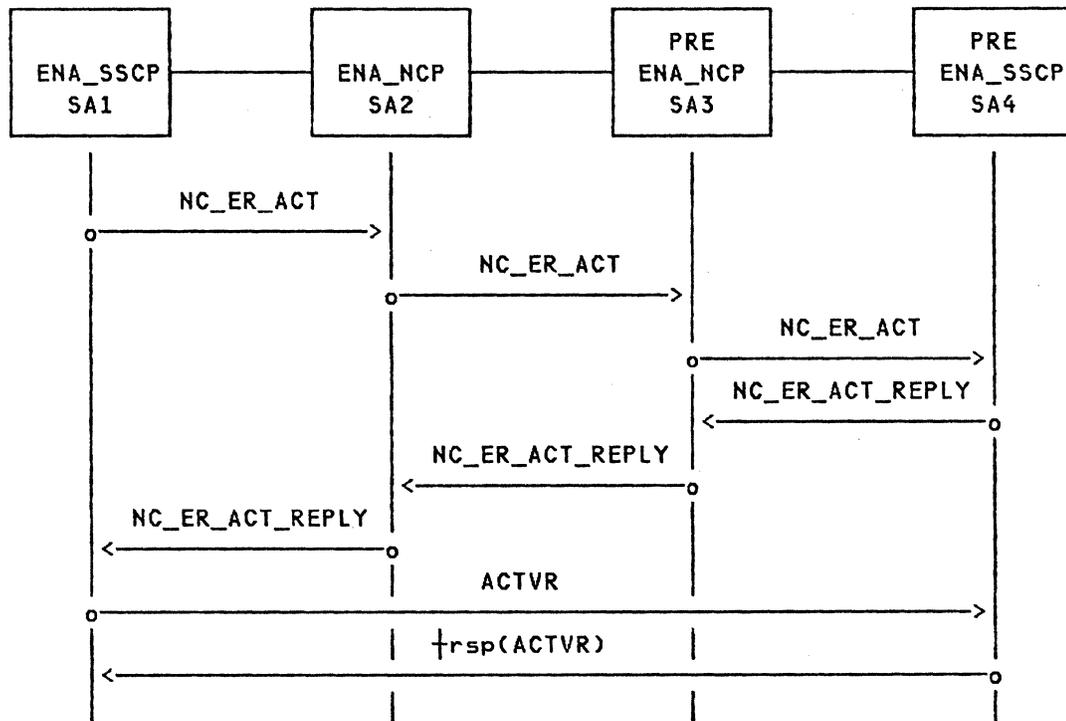


Figure 15. ER/VR Activation - Pre-ENA/ENA Network

- The ER/VR activation would be successful provided the subarea address of the origin (one endpoint) was within the MAXSUBA value used by the other endpoint.
- Assume NCP SA3 were owned by a third host (neither SA1 nor SA4). Then if SA4 is an ENA node with subarea number above MAXSUBA, it could contact (but not establish a VR to) NCP SA3 and establish a VR to SA1.

ENA_SSCP; ENA_NCP; Pre-ENA_NCP; ENA_NCP

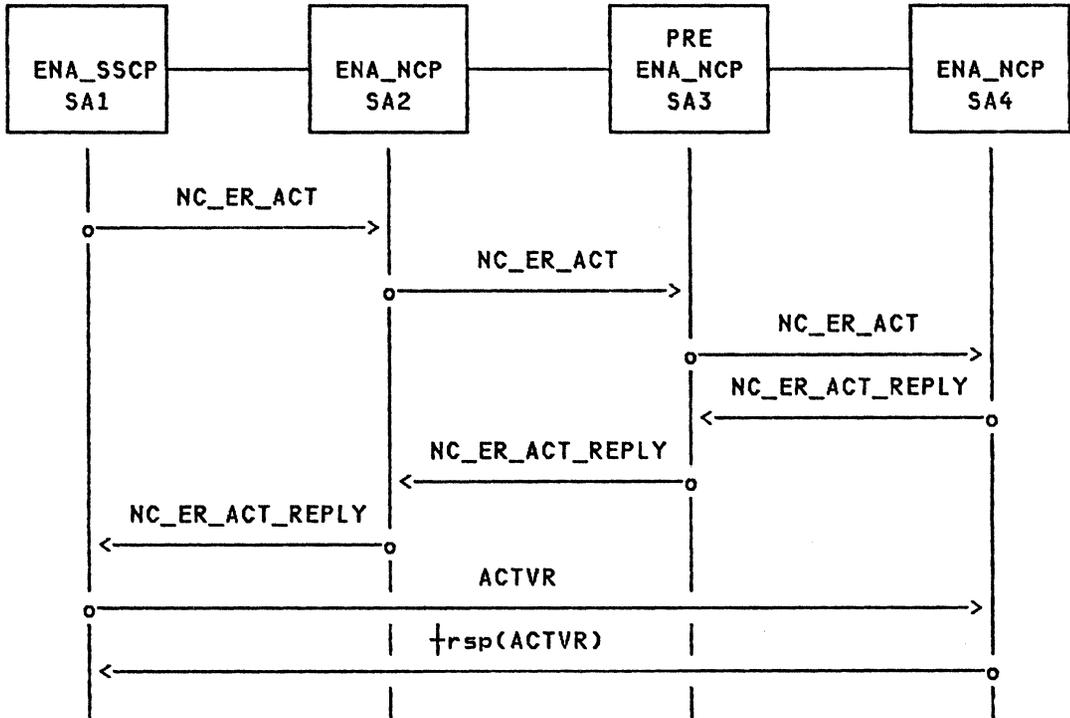


Figure 16. ER/VR Activation - Pre-ENA INN

- The ER/VR activation would be successful even if the subarea addresses of all ENA nodes are higher than MAXSUBA.
- If that is the case the pre-ENA NCP can act as Intermediate Network Node only. (A PTF is required for this case.)

4.3.4 LU-LU Session Activation

If a SSCP detects a mismatch between ENA and pre-ENA nodes during the INIT (and CDINIT), it will return a negative response with a sense code X'088E'.

If an unconditional request (during RNAA) for a pre-ENA compatible address cannot be satisfied, a negative response with sense code X'0812' will be returned.

ENA_SSCP; ENA_NCP; ENA_NCP; ENA_SSCP

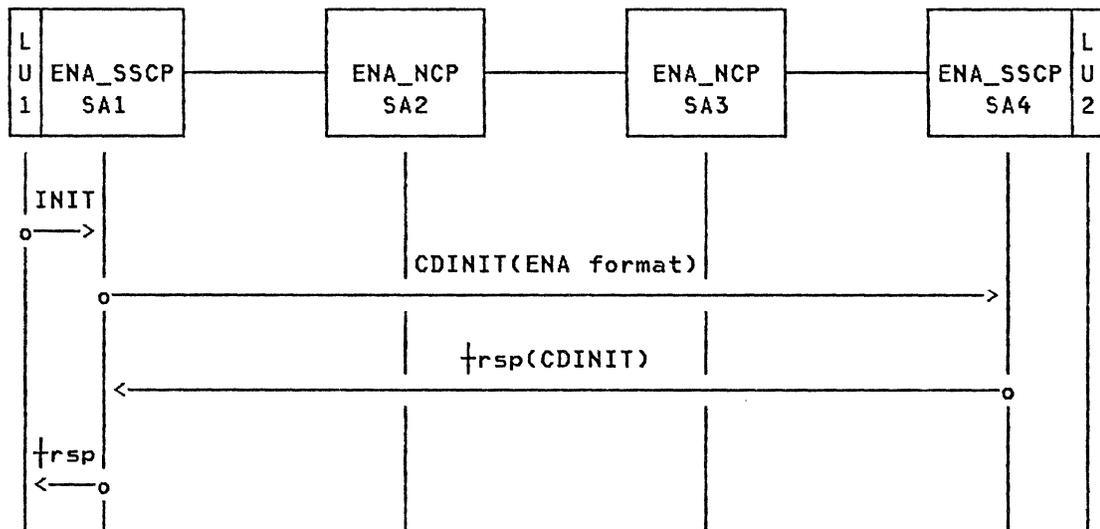


Figure 17. LU-LU: ENA_SSCP; ENA_NCP; ENA_NCP; ENA_SSCP

- Both SSCP's previously agreed to support ENA (through the ACTCDRM operation).
- In this environment the LU-LU session set-up is identical to pre-ENA LU-LU session set-up.

ENA_SSCP; ENA_NCP; Pre-ENA_NCP; Pre-ENA_SSCP

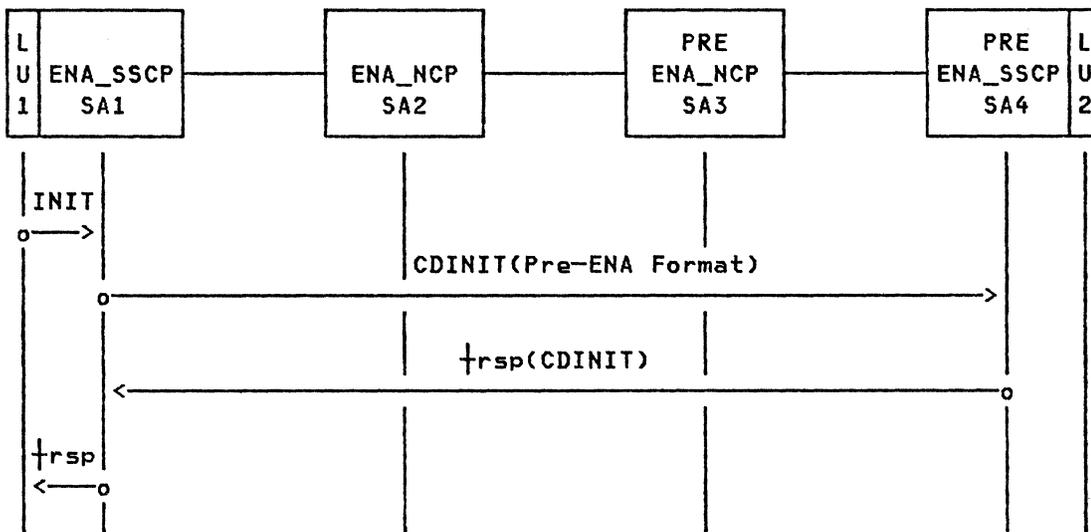


Figure 18. LU-LU: ENA_SSCP; ENA_NCP; Pre-ENA_NCP; Pre-ENA_SSCP

- The SSCP-SSCP had previously determined that the pre-ENA format was applicable.
- The CDINIT operation was successful.

PRE-ENA_NCP; ENA_SSCP; ENA_NCP; ENA_SSCP

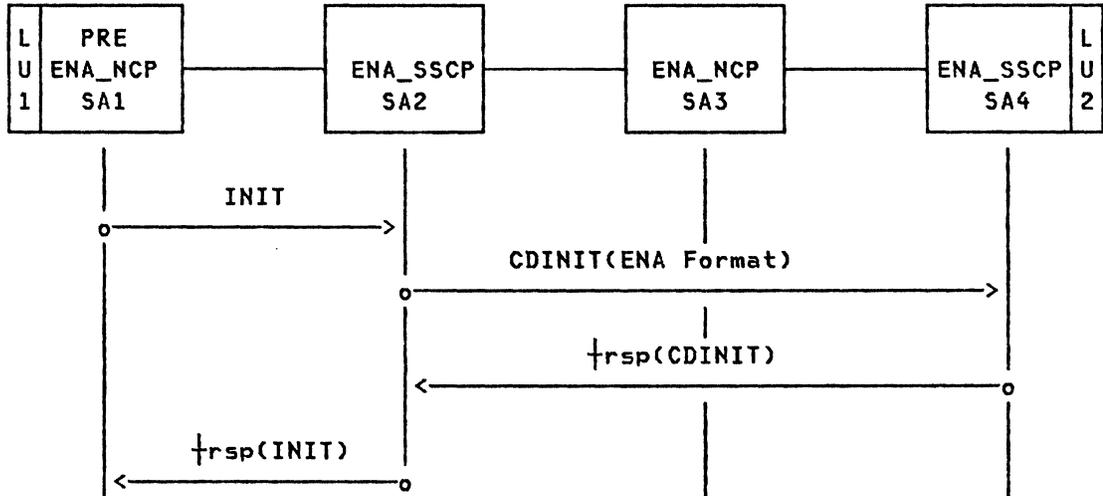


Figure 19. LU-LU: Pre-ENA_NCP; ENA_SSCP; ENA_NCP; ENA_SSCP

-
- All nodes in the network conform to the subarea split used by SA1.
 - The CDINIT operation was successful.

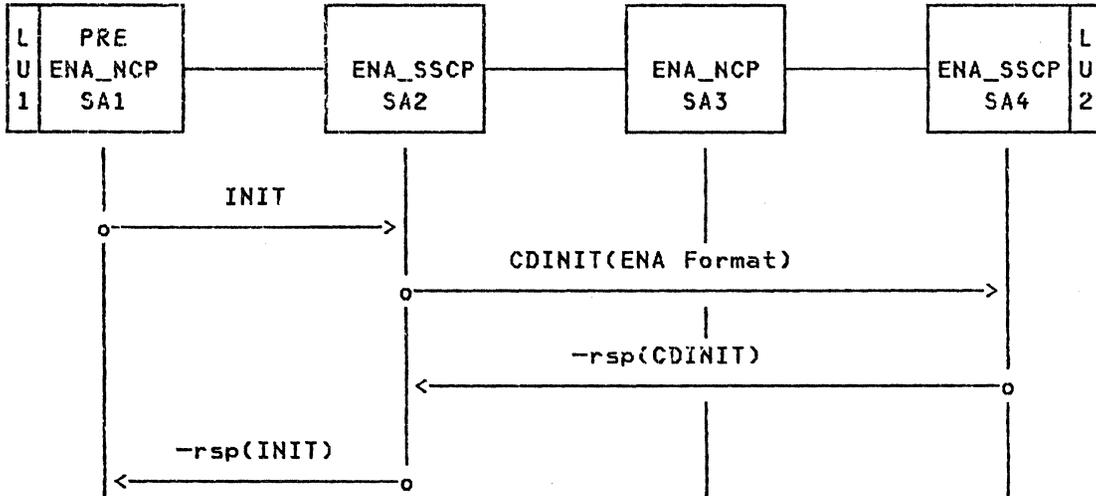


Figure 20. LU-LU: Pre-ENA_NCP; ENA_SSCP; ENA_NCP; ENA_SSCP

- The network element address of LU2 is above the range addressable by pre-ENA NCP.
- The LU-LU session setup failed (sense 088E).

5.1 ACF/VTAM V3

5.1.1 ITLIM Enhancements

ITLIM is a decimal integer (0-65535) that specifies:

- the maximum number of Session Services (SS) requests (i.e., USS requests like LOGON/LOGOFF and formatted requests like INITIATE/TERMINATE); and
- the maximum number of Logical Units Services (LUS) requests that originate from the VTAM API (e.g., SIMLOGON, OPNDST, CLSDST)

that can be processed simultaneously by VTAM.

VTAM will queue all requests so that the concurrent number of requests processed does not exceed the ITLIM value. This value specified will be treated internally as two independent pacing values (i.e., ITLIM of SS requests and ITLIM of LUS requests may be processed concurrently).

In ACF/VTAM V3, this limit will apply to same-domain, cross-domain and cross-network requests. This will reduce VTAM's requirement for storage during session establishment/termination and some error recovery conditions. ITLIM will not pace the establishment/deactivation of SSCP-PU or SSCP-LU sessions.

In order to prevent deadlocks, the queues will be checked every three seconds. If no requests have been processed since the last interval, five requests will be allowed to run.

5.1.2 Elimination of VTAMOBJ

The VTAMOBJ function will be discontinued in ACF/VTAM V3. However, the VTAMOBJ DD statement may be left in the VTAM startup procedure. It will be ignored.

Although VTAMLST major node definitions will be recompiled at each and every activation, this enhancement will eliminate DASD space problems (e.g., X37 abends) associated with VTAMOBJ.

5.1.3 Automatic SSCP-SSCP Session Restart

From ACF/VTAM V1R3 through to ACF/VTAM V2R2, SSCP-SSCP failures were treated in this manner:

- Active LU-LU sessions (i.e., INIT,CINIT,BIND,SESSST completed) continue without interruption, provided the paths used are not affected.
- Queued LU-LU session requests (i.e., INIT completed but CINIT to drive the PLU's logon exit has not been sent) managed by the failed CDRMs are discarded. The initiator is notified of 'session setup failures'.
- Pending active LU-LU session requests (i.e., INIT,CINIT completed but BIND and SESSST not completed) may succeed if the route is not affected. However, knowledge of these sessions is discarded (DISPLAY of LU will not reflect this session status).
- The SSCP-SSCP session can be re-activated by network operator command (i.e., V NET,ACT,ID=cdmname). The session will be re-bound using ACTCDRM(ERP) on:
 - an alternate route defined within the COS; or
 - the same route, if it has recovered.

In ACF/VTAM V3, SSCP-SSCP session error recovery will proceed in much the same way. However, the network operator 'V NET,ACT,ID=cdmname' function will now be performed automatically by VTAM. NCCF CLISTS and Command Processors invoked on SSCP-SSCP session failure should be reviewed.

For this function to work, one SSCP must be at ACF/VTAM V3 and the other may be at an earlier release. The ACTCDRM(ERP) needs to flow successfully from one SSCP only.

The user can disable this function through the CDRM definition statement in VTAMLST.

In single-network environments (i.e., non-SNI), this function should always work, provided there is a virtual route (restored or alternate) available. The ACTCDRM(ERP) can be queued.

In multiple-network environments (i.e., SNI connected networks), failures in configurations where SSCPs are in adjacent networks will recover. As there will be cases where the first recovery attempt is unsuccessful, and a subsequent recovery attempt from the other SSCP is required, recovery will not always be immediate. In the case where one of the SSCPs has lost its session with the Gateway NCP, the session may remain PACDR until a valid Gateway NCP becomes available.

Failures in back-to-back configurations will not recover. In those cases, the session will be INACT or INACTX, and an operator V ACT command should be issued (when the necessary resources become available). This will not disrupt existing LU-LU sessions.

5.1.4 SSCP Selection in Session Management Exit Routine

Prior to ACF/VTAM V3, the Session Management Exit Routine (ISTEXCAA) was invoked for these LU-LU session management functions:

- Accounting
- Authorization
- Gateway Path Selection

In this release, the exit may be also used for SSCP Selection.

SSCP Selection is useful in establishing cross-domain and cross-network LU-LU sessions where the SSCP owning the DLU (Destination LU) is not known.

- For a cross-domain session request, a default list of SSCPs can be provided. The session request will be routed to each SSCP in order until the owning SSCP is found or the list is exhausted.
- For a cross-network session request, the GW_SSCP will use a list of adjacent SSCPs. The session request will be routed as described above.

The SSCP Selection function operates as follows:

- VTAM will invoke the exit during LU-LU session establishment. (If no adjacent SSCPs exist, this function will not be invoked.)
- The exit receives as input:
 - a default list of SSCPs; or
 - a list of adjacent SSCPs built during the previous LU-LU session setup for the same resource.
- The exit can be coded to reorder or shorten the list of SSCPs. When the exit returns, this list will be presented to VTAM.
- VTAM will route the session setup request to the next SSCP in the list.

The SSCP Selection Function and Gateway Path Selection Function are related. When processing a cross-network LU-LU session request, the GW_SSCP will:

- invoke SSCP Selection; and
- then invoke Gateway Path Selection to determine the GW_NCP to use in the LU-LU session path.

This will help improve performance in some networks, by avoiding lengthy trial and error processing during session establishment.

5.1.5 PLU/SLU Names in Virtual Route Selection Exit Routine

The names of the PLU and SLU will be provided as input to the Virtual Route Selection Exit (ISTEXCVR) when it is driven during session establishment.

This exit allows the user to modify the order of virtual routes as specified in the COS entry.

Installations with naming conventions that identify the subarea within the LU name will find it easier to exploit this exit.

5.1.6 Network Management Vector Transport (NMVT)

NMVT is the vehicle used to transfer network management data between the device collecting the data and the application program using the VTAM Communications Network Management Interface (CNMI).

NMVT replaces the REQMS, RECFMS, and RECMS. All the new functions (i.e., Session Information Retrieval, Dynamic LPDA) provided by ACF/VTAM V3 and ACF/NCP V4 are based on the NMVT.

5.1.7 Elimination of Message Flooding

Some network error conditions flood the operator console with replicated messages. In ACF/VTAM V3, only the first occurrence of certain messages (see "Non-Replicated VTAM Messages" on page 176) will be displayed. Any subsequent messages that are identical in text and appear within 30 seconds of the first occurrence will be suppressed from display and logging and will not be passed to VTAM POI programs.

EXAMPLE: When an Explicit Route fails, each PU_T4 subarea in the network detecting the condition will:

- report with a ROUTE.INOP(NS) to its owning SSCP; and
- notify adjacent PU_T4s of the condition.

The process will continue until all eligible PU_T4s have been notified.

Prior to ACF/VTAM V3R1, each ROUTE.INOP(NS) received by the SSCP would produce a multi-line message (IST526I). In this release, only the first occurrence of the message will be displayed.

5.1.8 LUNAME in USS Messages

ACF/VTAM V3 allows the terminal's LU name to be included in USS messages. When &&LUNAME is coded VTAM will replace it with the terminal's LU name.

5.1.9 TSO Logon Fix

Prior to ACF/VTAM V3, TSO could not process LOGON requests requiring an ACB beyond one that had been deactivated by operator command. TSO will now bypass the deactivated ACB and use the next available.

5.1.10 API Enhancements

In addition to API changes related to MVS/XA (see "VTAM Application Programs" on page 14), ACF/VTAM V3 provides the following enhancements:

- Sense Code on CINIT Response

Programs that issue CLSDST OPTCD=RELEASE can specify a SENSE code to be provided with a negative response to CINIT.

- SON Code on UNBIND

Programs may specify a SON code to be provided with the UNBIND for:

- TERMSESS
- TERMSESS OPTCD=UNBIND
- CLSDST OPTCD=RELEASE

The architected SON codes are:

01	Normal end of session.
02	BIND forthcoming.
03	Talk
04	Restart mismatch.
05	LU not authorized.
06	Invalid session parameters.
07	Virtual route inoperative.
08	Route extension inoperative.
09	Hierarchical reset.
0A	SSCP gone.
0B	Virtual route deactivated.
0C	LU failure - unrecoverable.
0E	LU failure - recoverable.
0F	Cleanup
FE	Invalid session protocol (application should supply SENSE).

- Synchronous LOGON/SCIP Exit Scheduling

SETLOGON HOLD will cause all BINDs and CINITs to be queued; therefore the LOGON and SCIP exits will not be driven immediately. SETLOGON START will resume exit scheduling.

This may prevent the over-allocation of virtual storage during application session establishment.

- Queued Response Notification

SEND OPTCD=RSPQUED|NRSPQUED

When OPTCD=RSPQUED is coded on the SEND macro instruction, VTAM will look for queued responses. When the SEND is posted complete:

- the flag RPLRSPNM will be set if there are any responses on the normal flow inbound response queue; or
- the flag RPLRSPQR will be set if there are any responses on the normal flow inbound data queue.

The application program can test these flags to see if there are any queued responses.

When OPTCD=NRSPQUED is coded on the SEND macro instruction, VTAM will not look for queued responses.

5.1.11 NCP DISPLAY Enhancement

The result of an NCP DISPLAY (i.e., D NET,ID=ncpname) will indicate:

- the type of dump (i.e., CSP, DYNA or MOSS) that is active; and
- whether the NCP is in slowdown condition.

This should improve operator awareness.

5.1.12 NCP DUMP COMPLETE Message Improvement

The NCP DUMP COMPLETE message will indicate the type of dump that has been completed (i.e., CSP, DYNA or MOSS).

This will reduce confusion if two or more dumps are requested concurrently for the same NCP.

5.1.13 Dump Formatting

In ACF/VTAM V2R2, the VTAMMAP option of the MVS Service Aid PRDMP caused the formatting of the following data areas:

ATCVT	PST	MPST	ACDEB	LUCB	FMCB
FMCB Ext	BPDY	BPCB	PXB	CONFT	QAB
RDTE	NCB	CRA	ITTRC		

In ACF/VTAM V3, the following areas will be formatted in addition to those shown above.

LQAB	WRE	EID	ERT	VRT	SIB
------	-----	-----	-----	-----	-----

In addition, the VTAMMAP option has been enhanced to improve selectivity in dump formatting. One or more of these options may be specified:

- ALL** Provide the VTAM, STORAGE, RDTFULL, SES and ROUTE options as described below.
- ONLY** Format the ATCVT and provide the SES and ROUTE options.
- RDTFULL(name)** Format the Resource Definition Table entries (RDTE) in their entirety along with their Node Control Blocks (NCB). If (name) is specified, then format RDTE for that node only. Else, format all RDTEs.
- RDYHIER(name)** Format the specified RDTE and all RDTEs that are below the one specified in the RDTE hierarchy. If (name) is omitted, all RDTEs are formatted (like RDTFULL).
- RDTSUM(name)** Format RDTEs in a summarized form. This provides a one-line entry per RDTE.
- ROUTES** Format the explicit and virtual route tables (i.e., ERT, VRT and associated VR blocks).
- SES(name)** Format RDTEs for all major/minor nodes along with LUCB, CDEB, FMCB, FMCBE and SIB. The name can refer to session end points (i.e., APPL, LU or TERM).
- STORAGE** Format buffer pool control blocks (i.e., BPCB, PXB, SPANC, PGTE).
- VTAM** Format the RDT, RDTEs, MPST, PST, NCBs, BPCB, PXB, LQAB, WRE, and EIDs. In addition, provide the module names and addresses in the ATCVT.

5.1.14 Symptom String Subset

ACF/VTAM V3 will provide symptom string data via the MVS/XA System Diagnostic Work Area Variable Recording Area (SDWAVRA). This will be provided when a VTAM ESTAE or FRR exit routine is invoked on a program ABEND and can be used in RETAIN type searches.

See "Example of Symptom String Subset Data" on page 183 for an example.

5.1.15 Descriptor Codes

ACF/VTAM V3 will use MVS MCS Descriptor Code 3 for 'critical eventual action' WTO messages. These are messages suffixed with message type code E (i.e., ISTxxxE).

The 'critical eventual action' message results from an error condition that:

- requires operator action; but
- does not stop VTAM from executing until the action is complete.

Previously, these messages used MVS MCS Descriptor Code 11 (which 'freezes' messages on the operator console). In some network error situations, the operator console was saturated with 'frozen' messages.

In this release, these messages will not be frozen but may be recalled through the MVS 'D R,L' command.

5.1.16 NLDM Version 1 Release 3.1 Support

ACF/VTAM V3 will provide these enhancements for the SESSION AWARENESS and TRACE functions in NLDM Version 1 Release 3.1.

Negative BIND Response Data: When a BIND failure notification is received by the SSCP, NLDM will be notified with the BIND failure function code and sense code (from the BINDF/CDESSSF_RU). This will help in investigating session set-up failures.

Session End Reason Code: When a SESSEND notification is received by the SSCP, the session end data sent to NLDM will include the SESSEND reason code (from the SESSEND RU). This will help in investigating abnormal session terminations.

Non-truncated FMD.NS_PIU Trace Data: Function Management Data PIUs that carry Network Services related data for SSCP-SSCP, SSCP-PU and SSCP-LU sessions will be captured without truncation for NLDM. This will help in investigating session setup failures.

NLDM Trace Command Support for Undefined Resources: ACF/VTAM V3 will accept START_TRACE/STOP_TRACE requests from NLDM that refer to resources not currently defined and/or activated:

- VTAM will queue the request until the named resource is activated.
- When the resource is ACTIVE, the trace will be started automatically.
- When the resource is deactivated, VTAM will re-build the trace request which will be re-queued.
- When the resource is reactivated, the trace will be restated. This process will continue until NLDM terminates or a STOP_TRACE request is issued.

This assures continuity in tracing while the resource is in the ACTIVE state, despite intervening deactivation(s) of the resource or its major node.

5.2 ACF/NCP V4

5.2.1 NEO WRAP

In the previous release of 3725/NCP, Wrap tests, SIT, and LTRACE functions were not supported for the Network Extensions Option (NEO) user. NCP was position dependent on certain line control block fields. When using NEO 'code' these fields are positioned differently.

The following changes have been made in ACF/NCP V4 to provide the NEO user with additional function:

- a compatibility operand on the GROUP macro for NEO line groups that are compatible to NCP's.
- an indication in the CCBTYPE which line Type the user code will emulate - when the compatibility operand is specified (e.g., NSI could possibly use the BSC line protocol in the NCP processing).

This facility provides the NEO user the option to emulate NCP supported resources and perform the same tests available as the emulated device.

5.2.2 NCP SYSGEN/NDF

The entire NCP generation process has been changed with SSP V3. The Systems Programmer and IBM Systems Engineer should refer to NCP and SSP Installation and Resource Definition Guide, SC30-3253 for further information.

5.2.3 Service Aid Changes

There have been a number of improvements made to the Dump Formatter and the Configuration Report Program. Refer to "SSP Enhancements" on page 181 for a summary of these changes.

ACFTAP: ACF/TAP has been changed to reflect the addressing changes introduced by ENA as well as a series of network management enhancements.

5.2.4 Show Cause

Whenever the threshold for a station counter is exceeded or a Deactivation request is received, NCP builds one of these unsolicited RECMSs and forwards it to NPDA:

- BSC/SS Station Statistics RECMS - Recording Mode X'81' (e.g., error count threshold exceeded)
- SNA Statistics RECMS - Recording Mode X'86' (e.g., total retries threshold exceeded)

NPDA builds trend information for the network operator from these RUs.

SHOW CAUSE will use a byte in these records to indicate which threshold has been exceeded or that a Deactivation request was received.

5.2.5 Secondary Network Support

This support allows:

- the definition of a Secondary Network (e.g., the IBM 3710 Network Controller and its SDLC multi-dropped PU_T2 devices),
- LPDA tests to be run by station instead of link (to prevent the return of false problem determination information from certain devices - e.g., pseudo PU_T2 devices downstream of the IBM 3710 Network Controller); and
- ACF/TAP to format traces for SDLC, BSC and SS devices attached to the IBM 3710 Network Controller.

5.2.6 Session Information Retrieval (SIR)

When a session passes through a Gateway NCP (GW_NCP):

- the last two sequence numbers into the gateway transform; and

- the last two sequence numbers out of the gateway transform

are recorded by NCP. Prior to ACF/NCP V4, these numbers were not available to the operator for cross-network problem determination. Session Information Retrieval (SIR) gives the operator the ability to obtain these numbers for each cross-network session.

Network Logical Data Manager (NLDM) Version 1 Release 3.1 is required to activate/deactivate SIR for a specific session and to display SIR data. Obviously, an NCP with SNI function (i.e., ACF/NCP V3 +) is required.

5.2.7 Dynamic Threshold Alteration

Commands from NCCF Version 2 Release 2 are accepted by ACF/NCP V4 to query and modify threshold values. This allows increasing/decreasing the rate of statistical records to NPDA without requiring a system generation (SYSGEN) of NCP.

5.2.8 Dynamic LPDA

Commands from NCCF Version 2 Release 2 are accepted by ACF/NCP V4 to display or alter the LPDA parameters specifying the execution of LPDA. This allows changing the LPDA parameters when installing IBM 386x modems without requiring a SYSGEN of the NCP.

5.2.9 Modulo 128 for boundary network nodes (BNN)

Modulo 128 support for intermediate network node (INN) links was previously announced and is available with ACF/NCP Version 3 for the 3725. INN link support is also provided with ACF/NCP V4. Modulo 128 is now supported for boundary network node (BNN) links. This capability allows sending or transmitting up to 127 blocks of data in one direction before requiring a response, resulting in more efficient use of high-speed links or satellite transmissions.

In this part, the sections are:

- 6.0 "Planning and Migration Considerations" on page 49: Describes the considerations in:
 - determining if ENA and VSCR are required; and
 - planning for ENA and VSCR exploitation.
- 7.0 "Migration Scenarios" on page 69: Describes a number of pre-ENA to ENA migration scenarios. Not all possible combinations are considered. However, there are ten scenarios which should illustrate the principles required to plan for most migrations to ENA capable networks.

6.0 PLANNING AND MIGRATION CONSIDERATIONS

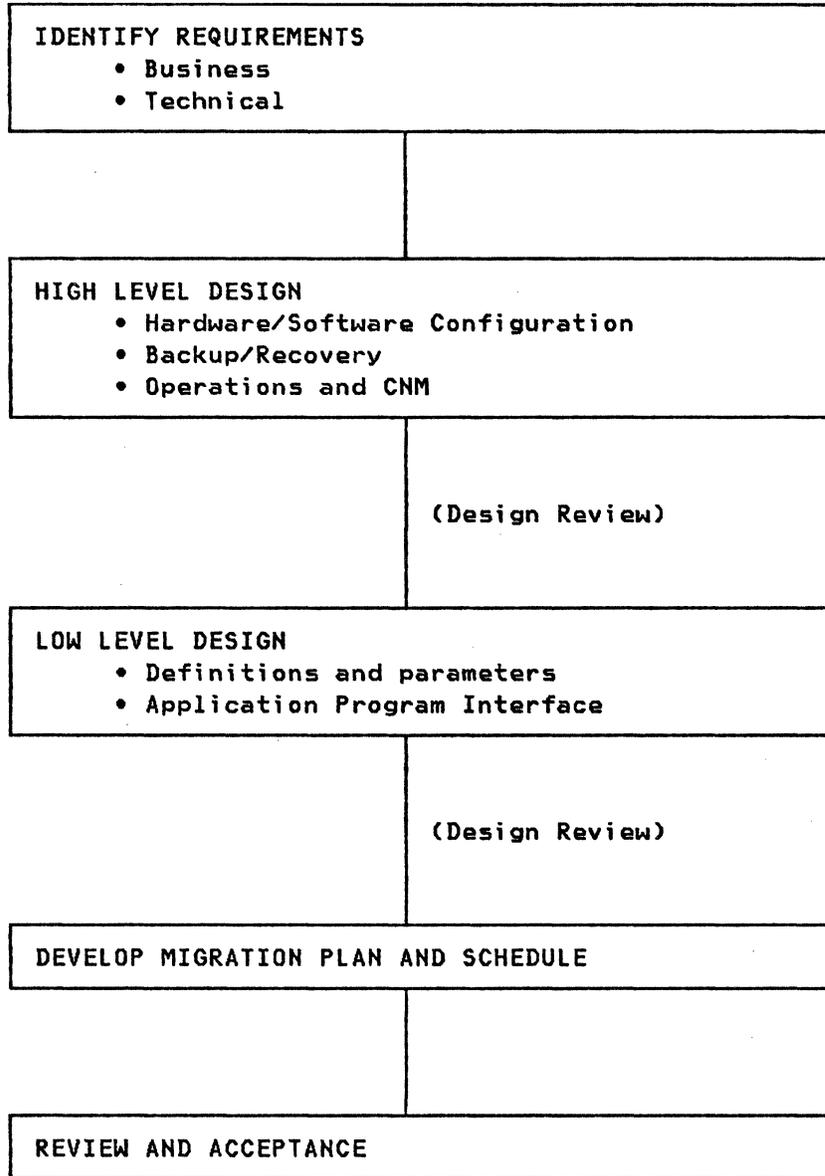


Figure 21. Major Planning Activities

6.1 IDENTIFY REQUIREMENTS

6.1.1 Business Requirements

What are the business requirements that the network must serve?

The network has a role in helping the enterprise meet its business objectives. Identifying this role will help prepare the business justification for the project and determine the design requirements.

From the business planning function, obtain information on the business expansion and business activities planned over the next 2-3 years. This information should indicate:

- the types and quantities of terminals,
- new service requirements; and
- changes to existing services.

A forecast of terminal and application requirements is often difficult to obtain. It is hard enough trying to determine terminal requirements from month to month.

Therefore, consider these alternative approaches:

1. Identify network growth to accurately determine if and when ENA and VSCR are required.
2. Order and install ACF/VTAM V3 and ACF/NCP V4. Treat the ENA and VSCR migration as another upgrade to the latest software levels. This will eliminate the costly (and often nebulous) exercise of trying to determine specific terminal requirements for the next 2-3 years.

6.1.2 Technical Requirements

6.1.2.1 Extended Network Addressing

The identified business requirements should help predict if and when the 16-bit network addressing scheme (used in releases prior to ACF/VTAM V3 and ACF/NCP V4) will be an inhibitor to growth and expansion. Figure 22 on page 51 shows the number of elements per subarea available for each MAXSUBA value under that scheme.

MAXSUBA (Number of Subareas)	Elements per Subarea
7	8192
15	4096
31	2048
63	1024
127	512
255	256

Figure 22. Pre-ENA Addressing Structure (16-bit Network Addresses)

Determine Remote Element Requirements: The following is an approach to determining network element and subarea addressing requirements.

1. Terminals/Workstations

- Do plans exist to migrate from the current terminal types to new function workstations?
- Is there a business strategy indicating the ratio of terminals to employees? Is there a business mandate to change the current ratio within a specified period?
- Is the enterprise investigating new business opportunities that could dramatically change the current network configuration?
- Are there business cases identifying productivity improvements through the use of terminals?
- The terminal requirements should be documented by location.

2. Cluster Controllers

- What is the expected branch office or 'service outlet' growth over the next 2-3 years? Plans for the relocation or creation of branches can often be determined from the organization's real estate function.
- Establish network design guidelines to determine the number of terminals to install on a control unit. In particular, look at line and control unit capacity.

3. Communications lines (i.e., switched, leased, virtual)

- Establish network design guidelines to determine the number of control units to pack on one line.
- Determine the serving area for each line. A serving area is a region defined by the network planner to be a group of locations which, when placed together on a line, offers the optimal price solution.

4. Other Requirements

- Generally speaking, you have to take into account that every addressable unit will take up one element address (e.g., PU/LU dynamic reconfiguration definitions, switched and virtual resources, half-sessions in Gateway NCPs, backup resources).

Determine Communications Controller Subarea Requirements

- Consider the degree of line diversification. If possible, spread lines going to the same destinations across multiple central site communications controllers for backup and recovery.
- Determine if growth in communications controllers will be centralized, distributed, or both.
- Are all services available from all communications controllers? Are multiple lines from the same branch required?
- The IBM Systems Engineer should run the 37x5 configurator and other IBM aids to determine the mix of resources that will meet the response time objectives defined in the service level agreements.

Determine Local Element Requirements

1. Applications

- Determine the requirements of operations, development and other internal DP functions (e.g., NCCF sessions, TSO access, TAF sessions, etc.).
- Determine the number of applications required for production, development and testing.
- Is the application development organization planning to exploit new functions (i.e., distributive processing, parallel sessions)?

DATE:	Current		+6 mth		+12 mth		+18 mth		+24 mth	
	Inc	TTD	Inc	TTD	Inc	TTD	Inc	TTD	Inc	TTD
A. No. of LUs and TERMINALS	-									
B. No. of PUs and CLUSTERS	-									
C. No. of LINES	-									
D. Total of Remote Elements	-									
E. Total Number of 37x5s	-									

DATE:	Current		+6 mth		+12 mth		+18 mth		+24 mth	
	Inc	TTD	Inc	TTD	Inc	TTD	Inc	TTD	Inc	TTD
F. No. of Applications	-									
G. Support Terminals	-									
H. Backup Elements	-									
I Total Of Local Elements	-									
J. Total Number of Hosts	-									

DATE:	Current		+6 mth		+12 mth		+18 mth		+24 mth	
	Inc	TTD	Inc	TTD	Inc	TTD	Inc	TTD	Inc	TTD
K. Elements in Largest Subarea										
L. Total Subarea Requirement										

Inc.....Increment for this period.
TTD.....Total to date.

Figure 23. SNA Subarea/Element Planning Chart

2. Terminals

- Determine the number of channel-attached terminals required.
- What type of terminal controller is used? Non-SNA control units require two elements to address each terminal - one for the UCB and one for the network name.
- Could there be plans to channel attach more terminals?
- Are new workstations using multiple LU sessions (i.e., 3270/PC, 3290) required?

Determine Host Subarea Requirements

- How many host processors will be required over the next 2-3 years?
- Are existing non-SNA host processors going to participate in the SNA network over the next 2-3 years?

Having estimated the network addressing requirements, the feasibility of the proposed network should be examined. In addition to considering other requirements necessary to achieving service levels (e.g., availability, backup/recovery, capacity, performance, problem management, change control, security), the virtual storage required by each VTAM in the network must be determined.

6.1.2.2 Extended Virtual Storage

Figure 24 on page 56 provides a structure for estimating VTAM's virtual storage requirements. Like the subarea/element plan, this should be projected over the next two years. Some installations may require more frequent milestones than the bi-annual ones used here.

These estimates should be made for every host CPU in the network.

The calculations in Figure 24 should be performed for each milestone.

The quantities of virtual storage (e.g., 0.82K of CSA for each LU in recovery) were derived from measurements taken in MVS/370 systems. Due to the asynchronous nature of VTAM scheduling, virtual storage requirements will fluctuate considerably over the duration of each VTAM process (e.g., session establishment, error recovery for large node failures). The values in Figure 24 on page 56 are only averages and should be treated as approximations useful for most properly configured and balanced VTAM host systems. A **balanced system** is one where the quantities of CPU power, main storage, channels and DASD are:

- available in the correct proportions (i.e., there are no 'bottlenecks' in the systems); and

* adequate to process its total peak workload.

In the description below, the terms CMP and AHP are used. These are from the IBM Communications Management Configuration (CMC) concept which advocates the ownership of all remote resources in an organization's network from a host CPU that is dedicated to that function and Communications Network Management.

CMP Communications Management Processor

The host processor that is a central point for operation and control of the network and 'owns' (i.e., activates/deactivates the terminals, controls error recovery, etc.) all lines, remote controllers and terminals. There are no application subsystems in this host other than network management subsystems. The CMP is also known as CMC Host.

AHP Application Host Processor

Other host processors in an environment that has implemented the IBM CMC concept. The AHP is also known as Data Host.

The calculations required to estimate VTAM's virtual storage requirements are described below:

- A-G Calculate the size of each VTAM buffer pool with the product of 'baseno' and 'bufsize'.
- H The number of terminal and application LUs (including CDRSCs) defined to this VTAM.
- I The number of TSO users logged on to this host.
- J This is the approximate amount of storage in the private area for VTAM modules.
- K The number of PUs, owned by this VTAM, that will be active during normal operation.
- L The number of LUs, owned by this VTAM, that will be active during normal operation. Use number in H above.
- M The number of LUs, owned by this VTAM, that will be in session with applications during normal operation.
- N The amount of private area virtual storage available (to all address spaces) in this MVS system. See "Overview of Virtual Storage and Addressing" on page 101.
- O The amount of CSA, in this MVS system, available for use by VTAM.
- AA The amount of CSA required by this VTAM under normal circumstances (i.e., all nodes activated, all session establishment complete and no error recovery in progress).

VS REQUIREMENT IN K BYTES	Current	+6 mth	+12 mth	+18 mth	+24 mth
A. CRPLBUF (baseno*bufsize)					
B. IOBUF (baseno*bufsize)					
C. LFBUF (baseno*bufsize)					
D. LPBUF (baseno*bufsize)					
E. SFBUF (baseno*bufsize)					
F. SPBUF (baseno*bufsize)					
G. WPBUF (baseno*bufsize)					
H. (0.25K * #LU active)					
I. (2.50K * #TSO Users)					
AA. CSA STEADY STATE (A+B+C+D+E+F+G+H+I)					
BB. CSA RECOV. (1K * #LU)					
CC. CSA PEAK (AA+BB)					
J. Base Modules in Private	1300K	1300K	1300K	1300K	1300K
K. (0.60K * #PU active)					
L. (0.40K * #LU active)					
M. (0.20K * #LU in session)					
DD. PRIVATE STEADY (J+K+L+M)					
EE. PRIVATE ERP (3.5K * #LU)					
N. Total Private Available					
FF. PRIVATE PEAK (DD+EE)					
O. CSA Available to VTAM					
GG. MAX NO. LUS RECOVERABLE (N-DD/3.5) OR (O-AA)					

Figure 24. VTAM Virtual Storage Planning Chart

BB For CMPs, the amount of CSA required by this VTAM for the recovery of LUs that it owns AND that were in session with applications (after a failure of the application or its host). The number of LUs can only be 'guesstimated'. As a starting point, use the number of sessions envisaged for the largest application host.

For AHPs, the amount of CSA required by this VTAM for the recovery of LUs with which it was in session (after a failure of the application). Use the number of sessions envisaged for the largest application subsystem on this host.

CC The maximum of CSA required by this VTAM (e.g., during network activation/deactivation without pacing controls, session establishment/termination with a major application host or error recovery for a large node failure).

DD The amount of storage in the REGION component of the private area (see "Overview of Virtual Storage and Addressing" on page 101) required by VTAM under normal circumstances (i.e., all nodes activated, all session establishment complete and no error recovery in progress).

EE For CMPs, the amount of storage in the REGION component of the private area required by this VTAM for the recovery of LUs that it owns and that were in session with applications (after a failure of the application or its host). The number of LUs can only be 'guesstimated'. As a starting point, use the number of sessions envisaged for the largest application host.

For AHPs, the amount of storage in the REGION component of the private area required by this VTAM for the recovery of LUs that it is in session with (after a failure of the application). Use the number of sessions envisaged for the largest application subsystem on this host.

FF This is the total storage that should be made available in the REGION component of the private area for recovery, as described above.

GG The maximum number of terminal LUs in session with applications that can fail concurrently and be recovered. The constraint will be either the CSA or private storage area; whichever can handle fewer LUs during recovery.

Calculate the CSA requirement:

$$\text{No. of LUs recoverable with available CSA} = (O-AA)$$

Calculate the private requirement:

$$\text{No. of LUs recoverable with available private} = ((N-DD)/3.5)$$

The maximum number of LUs recoverable will be the lower of the two values.

It should then be possible to determine if and when more virtual storage will be required by VTAM. These are the indicators:

- Insufficient CSA for steady state requirements (AA).
- Insufficient CSA for peak requirements (CC).

- Insufficient private storage for steady state requirements (DD).
- Insufficient private storage for recovery (EE).
- The number of LUs that could potentially fail concurrently (see BB or EE) exceeds 'maximum number of LUs recoverable' (GG).

6.1.2.3 Other Requirements

Apart from network addressing and virtual storage addressing constraints, ACF/VTAM V3 and ACF/NCP V4 should be installed for one or more of the following reasons (see "PRODUCT ENHANCEMENTS" on page 11 for details):

- ITLIM enhancements
- Elimination of VTAMOBJ
- Automatic SSCP-SSCP Session Restart
- SSCP Selection Exit
- Elimination of message flooding
- API enhancements
- Operator command and message interface enhancements
- NCP SYSGEN/NDF enhancements
- NCP service aid changes
- Support for the latest Communications Network Management products
- Support for the IBM 3710 Network Controller product
- To remove the design and operational complexity in the SNA Network Interconnection (SNI) function by using ENA

Whatever the reason, it is important to determine when ACF/VTAM V3 and ACF/NCP V4 should be installed.

6.2 HIGH LEVEL NETWORK DESIGN

Integrating ENA into the existing SNA network can easily turn into a complex problem. However, careful and thorough planning can minimize risks and complexity. The following sections offer an approach to planning for the installation of ACF/VTAM V3 and ACF/NCP V4.

6.2.1 General Design Guidelines

The following guidelines apply to the transitional steps of migrating from an exclusively pre-ENA network to a mixed environment:

1. ACF/NCP V4 is not available for 3705. This does not mean that a 3705 cannot coexist in the same network with ENA nodes - it can!
2. ACF/VTAM V3 is not available for OS/VS1.
3. ACF/VTAM V3 and ACF/NCP V4 will provide ENA support for VSE. SNI support is not provided for VSE with ACF/VTAM V2R2 and ACF/NCP V3 and will not be provided in ACF/VTAM V3 and ACF/NCP V4.
4. Check the general availability dates for the products that must be installed. The availability dates for the products will vary.
5. All pre-ENA nodes throughout the network must use a common MAXSUBA. This is no change to the current environment. If different MAXSUBA values are necessary (i.e., separate networks must interconnect), the SNI function must be used.
6. Pre-ENA nodes always send and receive the pre-ENA format of RUs. It is the responsibility of the ENA supported node to perform the appropriate transformation.
7. An ENA_SSCP or ENA_NCP can be intermediate routing nodes for pre-ENA RUs.
8. A pre-ENA_SSCP can act as intermediate routing node, if the endpoints of the route are below MAXSUBA.
9. A pre-ENA_NCP can act as intermediate routing node. If the endpoints of the route are above MAXSUBA, a PTF is required.
10. The subarea number of an ENA_SSCP or ENA_NCP that is adjacent to a pre-ENA SSCP cannot be greater than the MAXSUBA value.
11. The subarea number of an ENA_SSCP or ENA_NCP that is adjacent to a pre-ENA_NCP can be greater than the MAXSUBA value, provided the pre-ENA NCP is only used as an Intermediate Routing Node (IRN). (A PTF is required in this case.)
12. LU-LU communication can only exist between pre-ENA and ENA nodes if the element addresses are within the pre-ENA address range.

6.2.2 Migration Recommendations

1. There are no subarea restrictions, as long as all subarea numbers remain below MAXSUBA.
2. Install ACF/VTAM V3 on all hosts, prior to migrating to ACF/NCP V4. This will:
 - a. make VSCR and the other ACF/VTAM V3 enhancements available to the host;
 - b. provide the base to perform required Network Program Product upgrades;
 - c. provide operations with access to the new network management facilities;
 - d. reduce the planning and operations complexity associated with the coexistence of different addressing structures.
3. Install SSP Version 3 on all hosts to allow for loading and dumping of pre-ENA NCPs as well as ENA NCPs.
4. If there are pre-ACF/VTAM V2 or pre-ACF/NCP V2 nodes within the network, these must be brought up to at least Version 2 to co-exist within the ENA network. (ACF/NCP V1R2.1 nodes that cannot be upgraded are supported through Gateway NCP only.)
5. Keep the SSCP subarea address below the MAXSUBA value until all communications controllers and hosts have been converted to ACF/VTAM V3 and ACF/NCP V4 (i.e., migrate to a homogeneous ENA network prior to exploiting the ENA addressing extensions). If this is not possible, PTFs are required to generate pre-ENA NCPs.
6. SNI can be used to map together incompatible pre-ENA and ENA networks. SNI will convert a subarea number (greater than MAXSUBA) to an address compatible with the pre-ENA network. This should not be viewed as the initial attempt to solve the addressing problem as it does introduce an additional level of complexity - especially for customers that do not have the SNI function implemented. If BSC devices are used cross-network, PTFs are required for VTAM and NCP.
7. Continue to specify the MAXSUBA parameter until all communications controllers and hosts have been converted to ENA.
8. If a boundary NCP suffers from element id shortage, its migration should be given priority. Once the BNN conversion is complete, the Intermediate Nodes should be migrated to ENA.
9. Ensure network management products are reviewed for ACF/VTAM V3 and ACF/NCP V4 compatibility.
10. Host element address assignment can be controlled by the order in which the minor nodes are activated. By activating the key applications early in the network initialization process, one can ensure a lower element address

assignment. This may help to ensure that the ENA based applications are within the address range that coexisting pre-ENA based terminals use.

11. ACF/VTAM will reuse element addresses when allocating for parallel session usage.
12. To help determine the elements assigned to each resource within an NCP, you can use the Configuration Reporting Program (CRP).

6.2.3 Network Element Manipulation

Suggestions for sequencing host element address assignments to achieve an optimal use of addresses:

1. If local terminals use only the local processor's applications, the terminals should be activated last.
2. If applications require contact with pre-ENA applications/terminals, activate these applications early in the network initialization process. This will provide a better chance for address range compatibility.

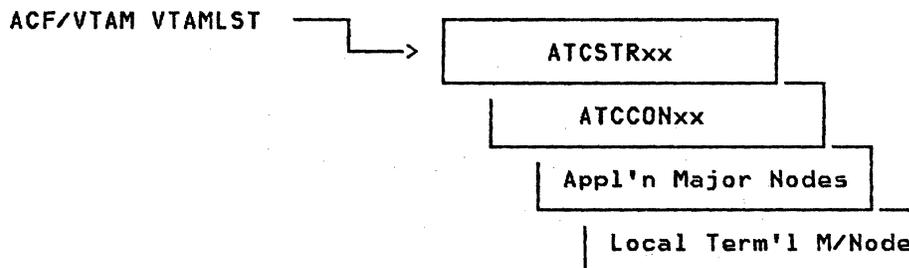


Figure 25. Sequencing Activation List (One)

3. If local terminals require contact with pre-ENA applications, activate these terminals early in the network initialization process. This will help provide a better chance for address range compatibility.
4. If parallel sessions will communicate only with applications based on ENA_SSCPs, then activate these resources last. This will force a higher element address to be used.

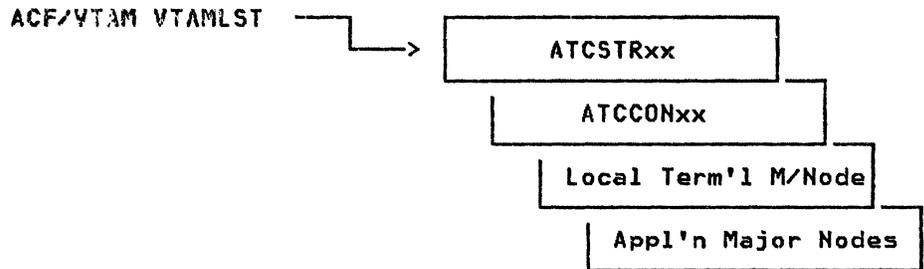


Figure 26. Sequencing Activation List (Two)

6.2.4 ACF/VTAM V3 - Operating System Compatibility

ACF/VTAM V3 is available on the following operating systems. However, the VSCR function is not provided for all operating systems.

OPERATING SYSTEM	ENA	VSCR	COMMENTS
MVS/370 Release 3.8	Yes	No	31-bit addressing requires MVS/XA.
MVS/SP V1 (MVS/370)	Yes	No	31-bit addressing requires MVS/XA.
MVS/SP V2 (MVS/XA)	Yes	Yes	
DOS/VSE AF Release 3.1 and 4	Yes	No	31-bit addressing requires MVS/XA.

Figure 27. ACF/VTAM V3 - Operating System Compatibility Chart

6.2.5 ACF/VTAM V3 - SSCP and NCP Compatibility

The following chart shows the levels of VTAM and TCAM with which ACF/VTAM V3 can have SSCP-SSCP sessions and the levels of NCP with which ACF/VTAM V3 can have SSCP-PU_T4(NCP) sessions.

SSCP LEVEL	SSCP-SSCP SESSION SUPPORT	COMMENTS
ACF/VTAM V1R2	No	Not supported
ACF/VTAM V1R3	No	Not supported.
ACF/VTAM V2R1 •	Yes	ENA/VSCR not available.
ACF/VTAM V2R2 •	Yes	ENA/VSCR not available.
ACF/VTAM V3 for VM/SP • (equal to ACF/VTAM V2R1)	Yes	ENA/VSCR not available.
ACF/TCAM V2R3 •	Yes	ENA/VSCR not available.
ACF/TCAM V2R4 •	Yes	ENA/VSCR not available.

NCP LEVEL	SSCP-NCP SESSION SUPPORT	COMMENTS
ACF/NCP V1R2.1	No	LU-LU through GW_NCP V3.
ACF/NCP V1R3	No	Not supported.
ACF/NCP V2 •	Yes	ENA not available.
ACF/NCP V3 •	Yes	ENA not available.

- The term 'pre-ENA node', used throughout this document, will apply to nodes containing these levels.

Figure 28. ACF/VTAM V3 SSCP and NCP Compatibility Chart

6.2.6 Network Program Product Compatibility

The following chart identifies the interdependencies between the IBM Network Program Products and ACF/VTAM V3 and ACF/NCP V4. The reader should consult Network Program Products Planning, SC23-0110 for a description of changes to these products.

Ensure that the level of the product shown in the chart is installed when migrating to ACF/VTAM V3 and ACF/NCP V4.

PRODUCT	LEVELS COMPATIBLE WITH ACF/VTAM V3 AND ACF/NCP V4
Network Communications Control Facility (NCCF)	V1R2, V2R1, V2R2
Network Problem Determination Application (NPDA)	V2R1, V3R1, V3R2
Network Logical Data Manager (NLDM)	R2, R3.1
Network Performance Monitor (NPM) (replaces NPA)	R1
Network Routing Facility (NRF)	R2
Network Terminal Option (NTO)	R3
Network Packet Switching Interface (NPSI)	information not yet available
Non-SNA Interconnect (NSI)	information not yet available
Systems Support Programs (SSP)	V3

Figure 29. Network Program Product Compatibility Chart

6.2.6.1 NCP Related Products

Customers with NRF, NTO, NPSI, or NSI installed must order and install the releases listed in Figure 29 on page 64. This is only if these products are to coexist within the same communications controller as ACF/NCP V4.

6.2.6.2 Network Performance Analyzer (NPA)

The NPA Host program is not supported by ACF/VTAM V3 and ACF/NCP V4. It has been superseded by the Network Performance Monitor (NPM) Program Product.

Therefore, NPA users should convert to NPM in order to obtain communications controller performance information in the ENA environment.

6.2.7 Determine Storage/Cycle Requirements

The IBM Systems Engineer has access to tools which will help determine the host and communications controller storage and cycle requirements for the customer's configuration.

6.3 LOW LEVEL NETWORK DESIGN

6.3.1 Pre-requisite Maintenance

1. VTAM

- PTFs are required to read an NCP definition with more elements than valid under pre-ENA restrictions.
- PTFs are required to activate a Link Station that contacts an NCP with subarea above MAXSUBA.

2. Systems Support Programs (SSP)

- A PTF is required for SSP to allow the specification of subarea addresses greater than MAXSUBA. This is to change five macros that check destination and adjacent subarea addresses against the MAXSUBA value.

3. SNI

- A PTF is required for support BSC devices cross-network.

6.3.2 VTAM/NCP Parameter Changes

There have been a number of changes to parameters in both ACF/VTAM V3 and ACF/NCP V4. For a summary of VTAM changes, refer to "VTAM Parameter Changes" on page 167.

For a complete description of NCP parameter changes, refer to NCP and SSP Installation and Resource Definition Guide, SC30-3253 and NCP and SSP Resource Definition Reference, SC30-3254.

6.3.3 VTAM API Changes

The VTAM Application Program Interface has been enhanced to allow an application to exploit 31-bit virtual storage addressing under MVS/XA.

There are other enhancements that are available in all operating systems supported by ACF/VTAM V3.

For further details, refer to "VTAM Application Program Interface (API) Changes" on page 171.

6.3.4 Operations Planning Considerations

Installations that have NCCF CLISTs, exits and command processors which rely on VTAM messages and sense codes should consult "New VTAM Sense Codes And Messages" on page 175 and "Non-Replicated VTAM Messages" on page 176. These provide information on the new/changed VTAM messages and sense codes.

6.3.5 Support Considerations

6.3.5.1 Network Planner

The Network Planner is responsible for establishing the network standards, the routing and connectivity topology for the network, and specifying the staging of changes to the production configuration. This person is key to migrating to ENA.

The Network Planner must understand (conceptually) the new functions available in ACF/VTAM V3 and ACF/NCP V4 and ensure that the business obtains the benefits of these enhancements.

The Network Planner must understand the pre-ENA -> ENA migration connectivity restrictions before any attempt is made to migrate the existing production service to ENA.

The Network Planner must also understand the installation's VTAM virtual storage requirements and the potential provided by the VSCR function. If this function is planned, the paging subsystem should be reviewed (see "Paging" on page 110).

Migration scenarios should be developed with an objective of preserving the current production service level commitments.

6.3.5.2 Implementation Coordinator

The implementation coordinator (in the context of this document) is the project leader responsible for the implementation of new levels of ACF/VTAM and ACF/NCP. He/she should understand the interdependencies identified by the Network Planner and ensure they are not ignored.

6.3.5.3 Systems Programmer

Systems programmers are responsible for installing and maintaining SCP based software. This includes selecting correct values for new operands and debugging any problems that arise. In addition, the system programmer should review any NCCF CLISTS, exits and command processors that are dependent on VTAM message formats and sense codes. If VSCR is required and increased paging activity is anticipated, he/she may have to reconfigure the paging subsystem.

The systems programmer must understand the changes made to the ACF/VTAM and NCP control blocks to support ENA. An understanding of the addressing technique, is also required. He/she should understand the VTAM and NCP parameter and definition changes.

6.3.5.4 Application Programmer

The network address expansion should be transparent to the application programmer.

However, there are other changes that the application programmer should consider. See "VTAM Application Program Interface (API) Changes" on page 171.

6.3.5.5 Network Operator

The network operator must now be aware that the network addresses are element addresses only. Any problem determination techniques or procedures that are dependent on the Subarea/Network Element split must be changed.

The network operator must understand (conceptually) the new functions available in ACF/VTAM V3 and ACF/NCP V4. He/she should be aware of new/changed VTAM messages and sense codes (and any changes made to NCCF CLISTs, exits and command processors, as a consequence).

6.3.5.6 Terminal User

The network address expansion should be transparent to the end user.

The following terms are used in this chapter:

ENA_NCP	an ENA supported NCP
ENA Network	a homogenous ACF/VTAM V3 and ACF/NCP V4 network
ENA_SSCP	an ENA supported SSCP
MSA	the MAXSUBA (MAXimum SUBArea) specification
Pre-ENA_NCP	ACF/NCP V2, ACF/NCP V3
Pre-ENA_SSCP	ACF/VTAM V2R1, ACF/VTAM V2R2, ACF/TCAM V2R4
SA	subarea address of the node
SNI	SNA Network Interconnect facility

7.1 SCENARIO ONE: INCREASE NUMBER OF CHANNEL-ATTACHED TERMINALS

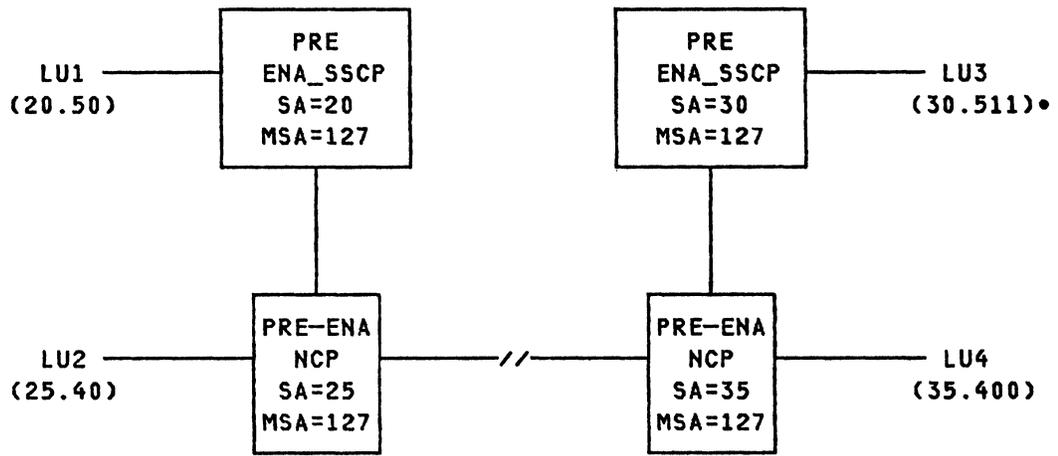
OBJECTIVE

- An installation has a requirement to install 200 additional terminals but has reached the network element addressing limits on that host. There is a preference for channel-attached terminals to reduce response time delays.
- ACF/VTAM V3 will be installed to increase the number of channel-attached terminals.

CONSIDERATIONS AND DESIGN CRITERIA

- The target processor participates in a multi-domain network and is accessed by terminals attached to any of the network nodes.
- LU1, LU2, LU3 and LU4 communicate with applications resident on all the hosts.
- The current MAXSUBA value is set at 127 - leaving the total element addressing count for each host and communications controller at 511.

PRE-ENA CONFIGURATION



ENA CONFIGURATION

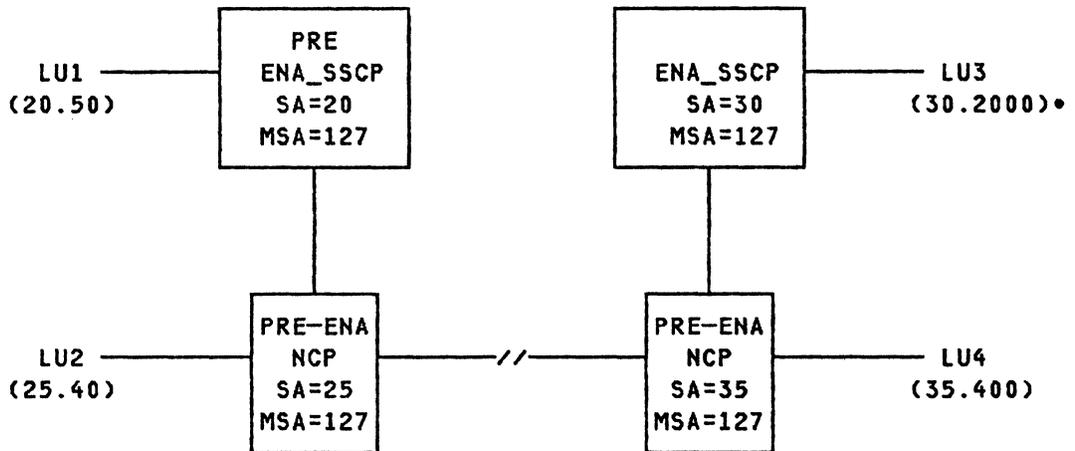


Figure 30. Expanding number of Local Terminals

OBSERVATIONS

- MAXSUBA is still specified on all nodes in the network. The ENA supported

host uses the MAXSUBA specification to determine how to address the pre-ENA nodes.

- During the SSCP-PU ACTPU session between SA30 and SA35 the result indicated that the NCP was incapable of handling ENA support. Having the MAXSUBA specified, SA30 acknowledged that pre-ENA addressing would be in effect.
- The local elements on SA30 falling between the range 511 and 2000 are not accessible for communication with a LU attached to any of the other hosts or communications controllers.
- The activation sequence on the ENA supported host was established to allow the application definitions to get assigned prior to the other network resources. This will increase the possibility that the application address assignment will fall within the pre-ENA subarea/element address range and will communicate with the terminals attached to the pre-ENA nodes.
- LU1, LU2 and LU4 can continue to communicate with the ENA based application provided the above technique was successful in setting the application element address within the pre-ENA address range.

RECOMMENDATIONS

- This is one of the easiest migrations possible. There are no complex routing configuration implications to plan or operate to. In addition, it provides immediate tangible benefits for installing ACF/VTAM V3.
- The new SA30 terminals, which have element addresses above 511, can communicate with applications in SA30 only.

7.2 SCENARIO TWO: INCREASE NUMBER OF HOST APPLICATIONS

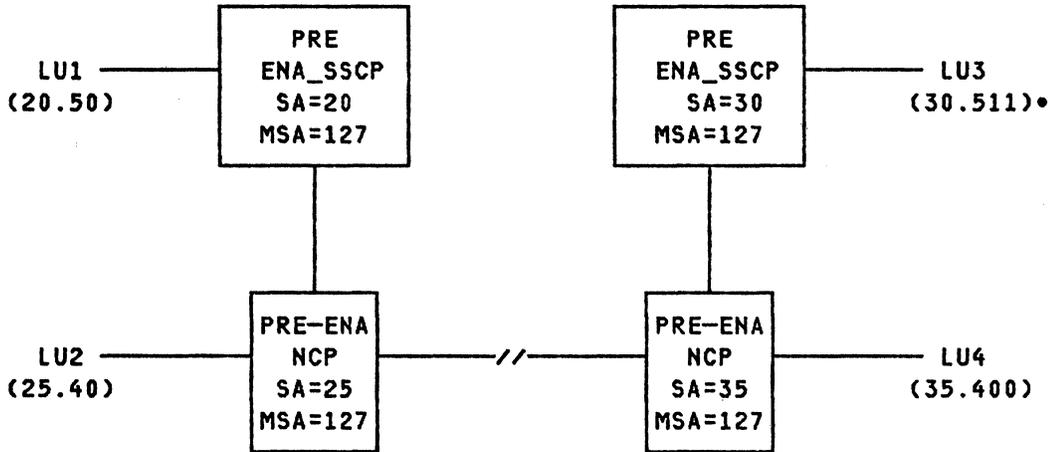
OBJECTIVE

- The installation must increase the number of concurrent TSO sessions but has reached the maximum number of local elements possible within the constraints of the subarea/element split. The number of subareas in use precludes reducing MAXSUBA to buy element relief. ACF/VTAM V3 will be installed to allow the number of channel-attached terminals and ACBs that will be required for the additional TSO sessions.

CONSIDERATIONS AND DESIGN CRITERIA

- The target processor participates in a multi-domain network and provides an application to NCP attached logical units as well as other host attached logical units.
- LU1, LU2, LU3 and LU4 communicate with applications resident on all hosts.
- The current MAXSUBA value is set at 127 - leaving the total element addressing count for each host and communications controller at 511.
- SA30 has reached 511 elements.

PRE-ENA CONFIGURATION



ENA CONFIGURATION

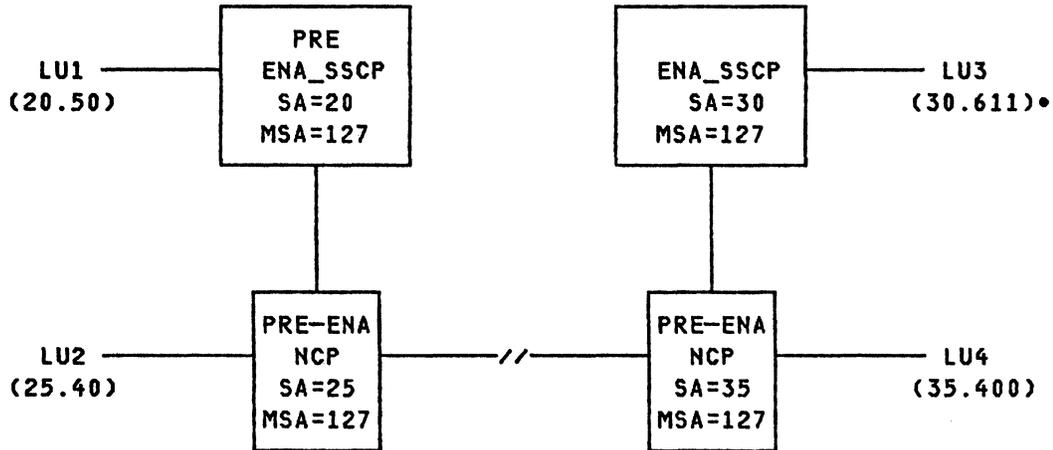


Figure 31. Expanding Applications Sessions

OBSERVATIONS

- MAXSUBA is still required on all nodes in the network.
- The application addresses which are above the MAXSUBA/ELEMENT range are not compatible with the MAXSUBA specified range existing on the pre-ENA nodes across the network. None of the terminals attached to the pre-ENA nodes will be able to access the application elements between 511-611.
- The terminals attached to the ENA supported host can continue to communicate with the applications within >MAXSUBA range (511-611).
- The activation sequence on the ENA host was setup to allow applications accessed by other nodes to be initiated first. This increased the possibility of cross-nodal communication.

RECOMMENDATIONS

- This is an easy migration step but should be viewed as a short term interim due to the increase in operations complexity.
- The next steps are to convert the remaining hosts and then the communications controllers requiring access to the ENA host applications.

7.3 SCENARIO THREE: MIGRATING MULTIPLE HOSTS TO ACF/VTAM V3

OBJECTIVE

- The requirement is to migrate an MVS/XA host to ACF/VTAM V3 to exploit the virtual storage relief feature. The ENA network feature is not a requirement (however, it is implicit in the migration to ACF/VTAM V3).

CONSIDERATIONS AND DESIGN CRITERIA

- The customer has implemented a 'local' network which provides access for a channel-attached terminal to communicate with all applications within the building.
- The hosts are a mix of VM/370 using VCNA/VS1, MVS/XA and MVS SP.
- Each host participates in the multi-domain network within the building.
- Hosts SA21 and SA40 communicate with remote devices through the attached NCP.
- The current MAXSUBA value is set at 127 - leaving the total element addressing count for each host and communications controller at 511.

PRE-ENA CONFIGURATION

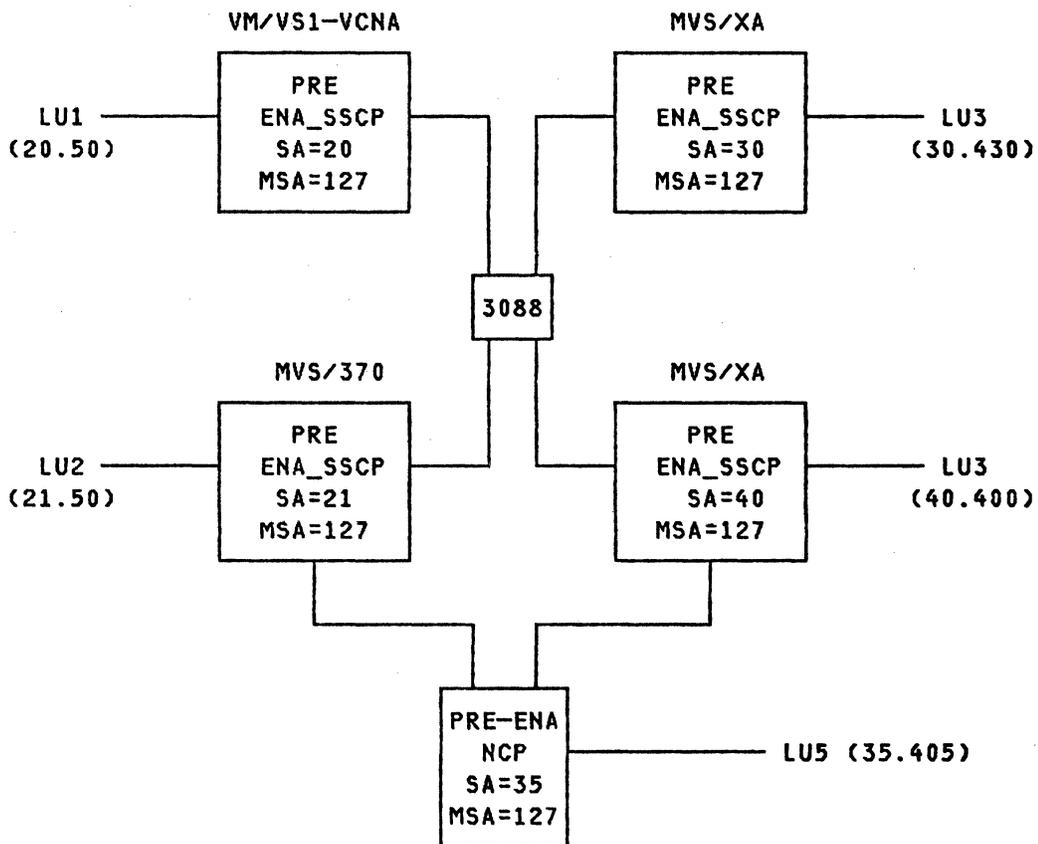


Figure 32. Migrating Multiple Hosts (Before)

ENA CONFIGURATION

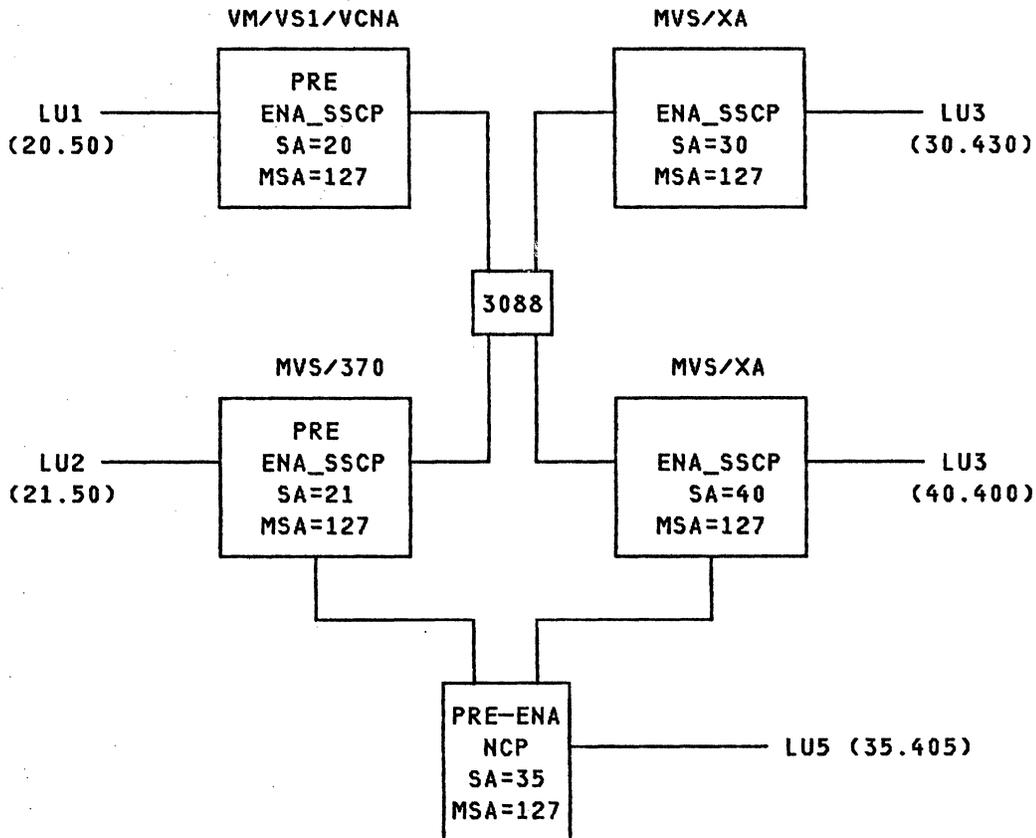


Figure 33. Migrating Multiple Hosts (After)

OBSERVATIONS

- During the ACTCDRM operations with SA20 and SA21, the responses indicated they were incapable of supporting ENA.
- During the ACTPU operation between SA40 and SA35, SA40 indicates that it supports ENA. SA35 is unable to support ENA and responds with the 'ENA not Supported' bit set in the Control Vector used by the ACTPU.

- Since the element address range and the subareas specified have remained unchanged, communication between all subareas is possible. No migration PTF is required.
- The response to the ACTCDRM between SA30 and SA40 indicated that ENA was supported.
- The MAXSUBA specified in both SA30 and SA40 is not used for communication between each of these nodes. However, it is required to specify how to communicate with the pre-ENA nodes.
- The communications controller could be a 3705 or 3725 provided an ENA compatible level of NCP is installed.

RECOMMENDATIONS

- This technique offers minimal impact to the existing environment and allows the installation to take advantage of VSCR.
- From a growth perspective, if all hosts were interested in migrating to ENA, the function provided by ENA is not available on a VS1 operating system.

7.4 SCENARIO FOUR: MIGRATING NCPS TO ACF/NCP V4

OBJECTIVES

- The installation would like to install ENA to resolve a network element constraint problem in the existing 3705 communications controllers.

CONSIDERATIONS AND DESIGN CRITERIA

- ACF/NCP V4 is supported on 3725 only - a hardware migration is required prior to the installation of ACF/NCP V4.
- The installation has an 'any-to-any' network strategy (i.e., any terminal in the network has the capability of accessing any of the host applications).

PRE-ENA CONFIGURATION

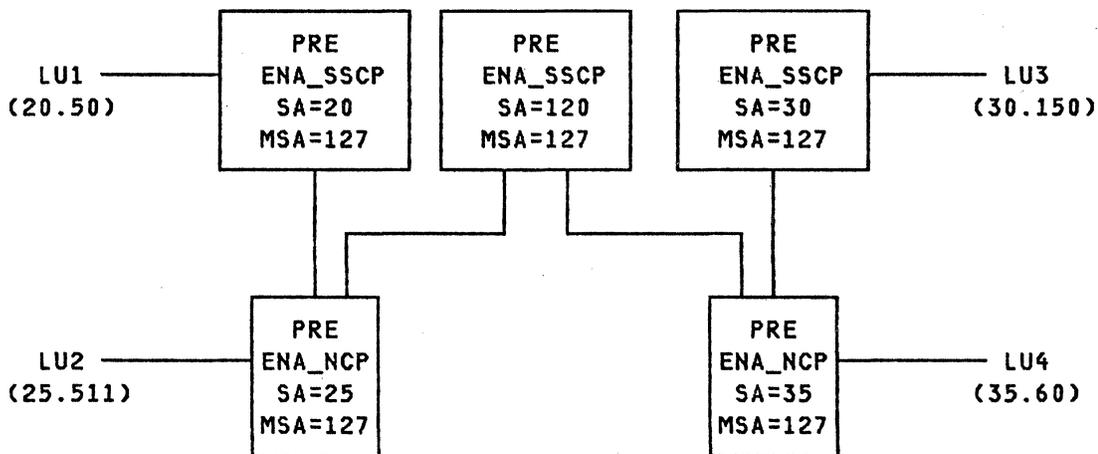


Figure 34. Migrating NCP to ENA (Before)

ENA CONFIGURATION

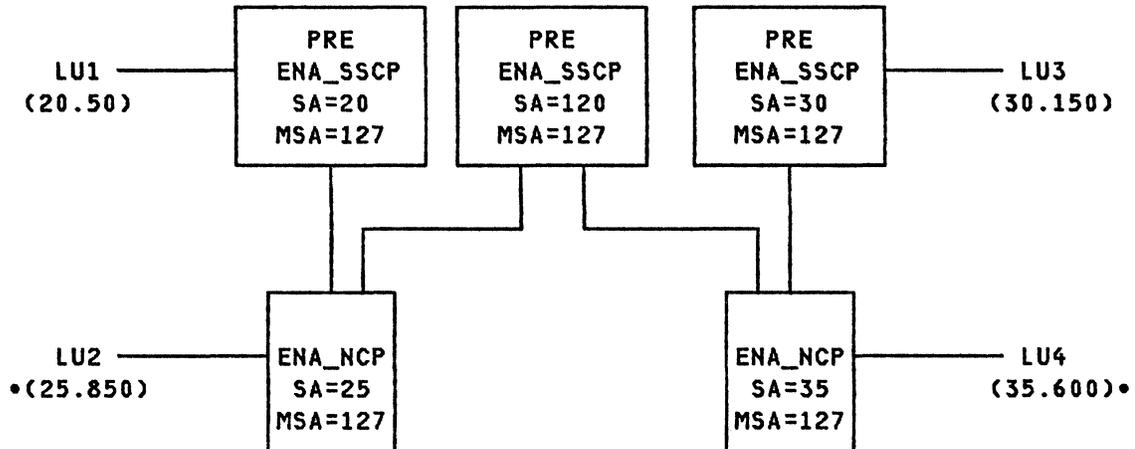


Figure 35. Migrating NCP to ENA (After)

OBSERVATIONS

- NCP generation will be successful as SSP V3 provides the expanded address capability.
- All the network resources identified on SA25 and SA35 above the 511 element address maximum are not activated by the owning SSCP. Consequently, network expansion above the pre-ENA subarea/element addressing range, is not possible.

RECOMMENDATIONS

- This approach to ENA migration is **not** recommended. Extending the network address does not provide any benefit - the pre-ENA hosts cannot communicate with new elements on the ENA_NCPs.

7.5 SCENARIO FIVE: CMC HOST MIGRATION TO ACF/VTAM V3 AND ACF/NCP V4

OBJECTIVES

- A network service organization (Domain 'B') provides an SNA networking facility for multiple customer networks (depicted by Domain 'A' and Domain 'C'). In addition, 'B' supports a major network of terminals attached to communications controllers. 'B' is constrained by the subarea/element address split currently implemented and would like to upgrade to ACF/VTAM V3 and ACF/NCP V4.
- All three networks participate as domains within a single network. For instance, the MAXSUBA and naming standards used by all domains are compatible. There is no subarea duplication.
- 'A' and 'C' do not have a mandate to install ENA support within the same timeframe.
- The objective in this scenario is to illustrate the steps and considerations when upgrading one network to support ENA while keeping the interfacing networks at pre-ENA levels.

CONSIDERATIONS AND DESIGN CRITERIA

- Domain 'B' SSCP performs the network management and ownership for a 'backbone' network consisting of a series of IRN (depicted by SA125) and BNN (depicted by SA126).
- Domains 'A' and 'C' are operating on 'pre-ENA' releases of ACF/VTAM and ACF/NCP and are not likely to migrate to ACF/VTAM V3 and ACF/NCP V4 at the same time as network 'B'.
- Terminals attached to all nodes have the capability to access applications on SA20 and SA30.

PRE-ENA CONFIGURATION

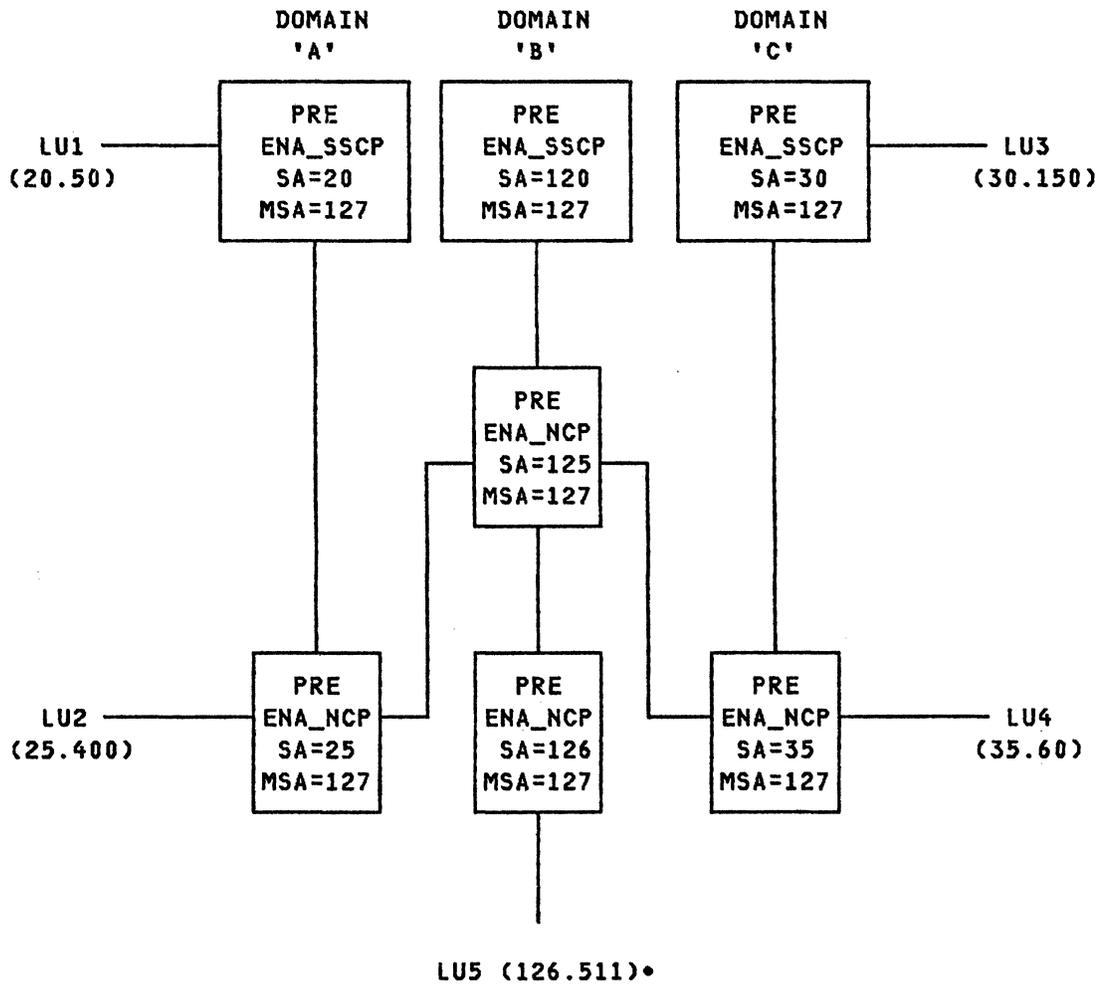


Figure 36. CMP and NCP Migration (Before)

ENA CONFIGURATION

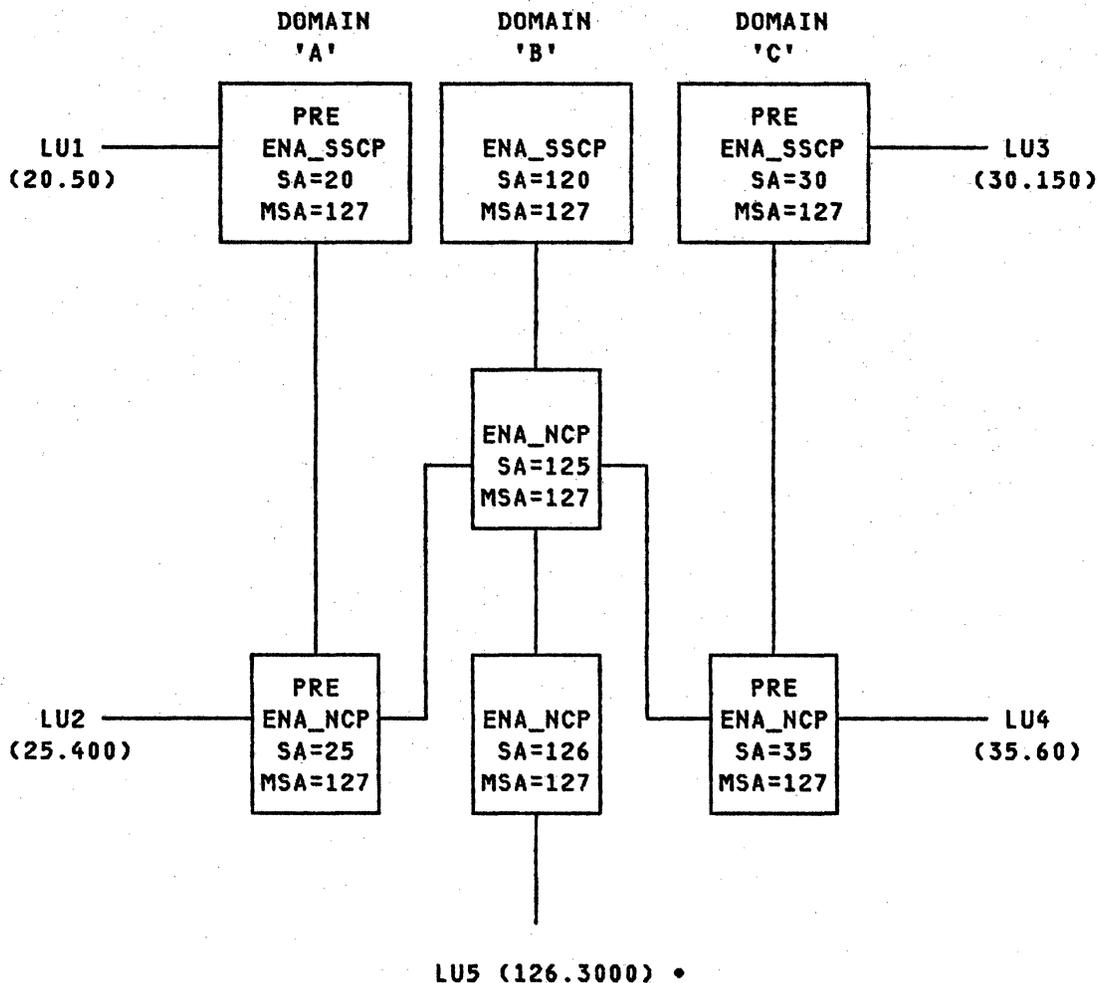


Figure 37. CMP and NCP Migration (After)

OBSERVATIONS

- Domain 'B' is on ENA level (i.e., all nodes can exercise the complete functions of ACF/VTAM V3 and ACF/NCP V4).

- Domains 'A' and 'C' are not on ENA level and communicate using the pre-ENA subarea/element addressing structure. All subarea and element addresses must fall within this range.
- The MAXSUBA specification is required in all nodes when pre-ENA addressing is present.
- The MAXSUBA value must be equivalent across all communications controllers and hosts.
- LU1, LU2, LU3 and LU4 can communicate with applications offered on SA120 provided the application has an element address assigned within the pre-ENA subarea/element address range.
- Element addresses between 511-3000 defined on SA126 are unable to communicate with the pre-ENA_SSCP (SA20 and SA30).
- The additional elements installed on SA126 can communicate with SA120 (the ENA_SSCP).
- In this scenario, SA120 provides the network management processor function (i.e., CMP). The other hosts are the application processors.

RECOMMENDATIONS

- This is **not** an advisable migration path. We achieved the objective of being able to add more terminals but were unable to access the target application processors in domains A and C.

7.6 SCENARIO SIX: MIGRATE APPLICATION HOST INITIALLY

OBJECTIVES

- An installation is interested in migrating the Application Hosts to ACF/VTAM V3 to increase the number of channel-attached terminals or for the VSCR function.
- For some reason (e.g., change control), the site has elected not to migrate the communications controllers and network management host until later.

CONSIDERATIONS AND DESIGN CRITERIA

- All terminals have access to SA20 and SA30.
- The MAXSUBA specification is required in all nodes when pre-ENA addressing is present.
- The MAXSUBA value must be equivalent across all communications controllers and hosts.

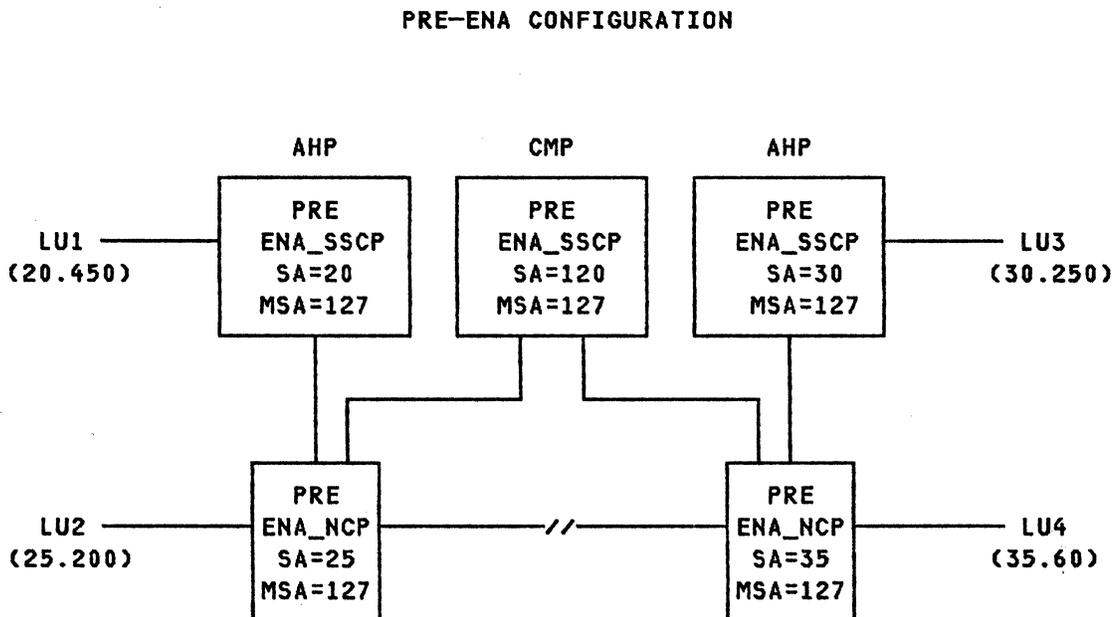


Figure 38. Application Host Migration (Before)

ENA CONFIGURATION

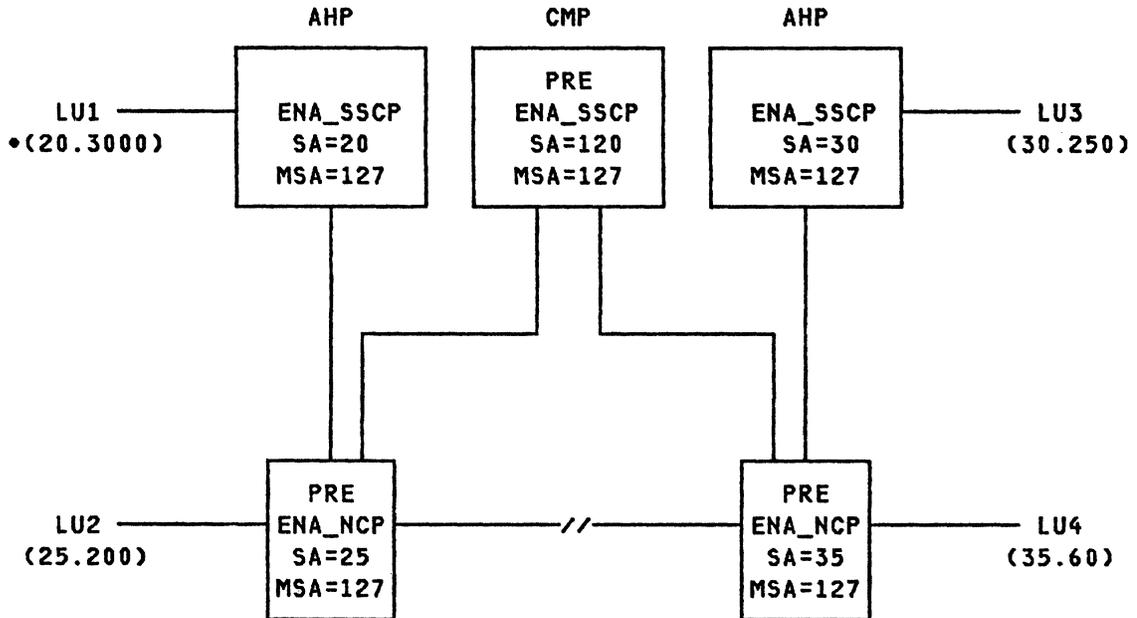


Figure 39. Application Host Migration (After)

OBSERVATIONS

- LU1 is able to communicate with all applications in SA20 and SA30.
- LU1 is unable to establish a session with the pre-ENA SA120, as a result of having an address beyond the pre-ENA addressing limit.

RECOMMENDATIONS

- To ensure that LUs across the pre-ENA nodes can continue to access the SA20 and SA30 applications, the application major node(s) should be activated prior to the terminals major node(s).

7.7 SCENARIO SEVEN: SNI AND ENA INTERCONNECTION

OBJECTIVES

- Network 'B' has reached the maximum number of network elements attainable with a MAXSUBA=127. Instead of installing an additional 3725, the installation will exploit the extended addressing provided by ACF/VTAM V3 and ACF/NCP V4.
- In the past, the installation solved this problem by implementing SNI and logically separating the two networks. Network 'A' is defined with a MAXSUBA=31, allowing 2047 element addresses on each subarea node.
- The operations staff would prefer to manage one logical network as all terminals across both networks access the same central site applications.

CONSIDERATIONS AND DESIGN CRITERIA

- The installation has implemented the SNI function provided by ACF/VTAM V2R2 and ACF/NCP V3 to interconnect two networks with different MAXSUBA parameters.
- SA15 is expected to grow from 511 network element addresses to approximately 2500.
- All terminals have the capability of accessing each of the hosts.
- SNI is used to map network 'A' resources into network 'B'. The gateway is needed to allow the coexistence of the network 'A' requirement for 2047 element addresses per subarea with the network 'B' requirement for a large number of subareas.
- Both networks use the same network naming standards.

PRE-ENA CONFIGURATION

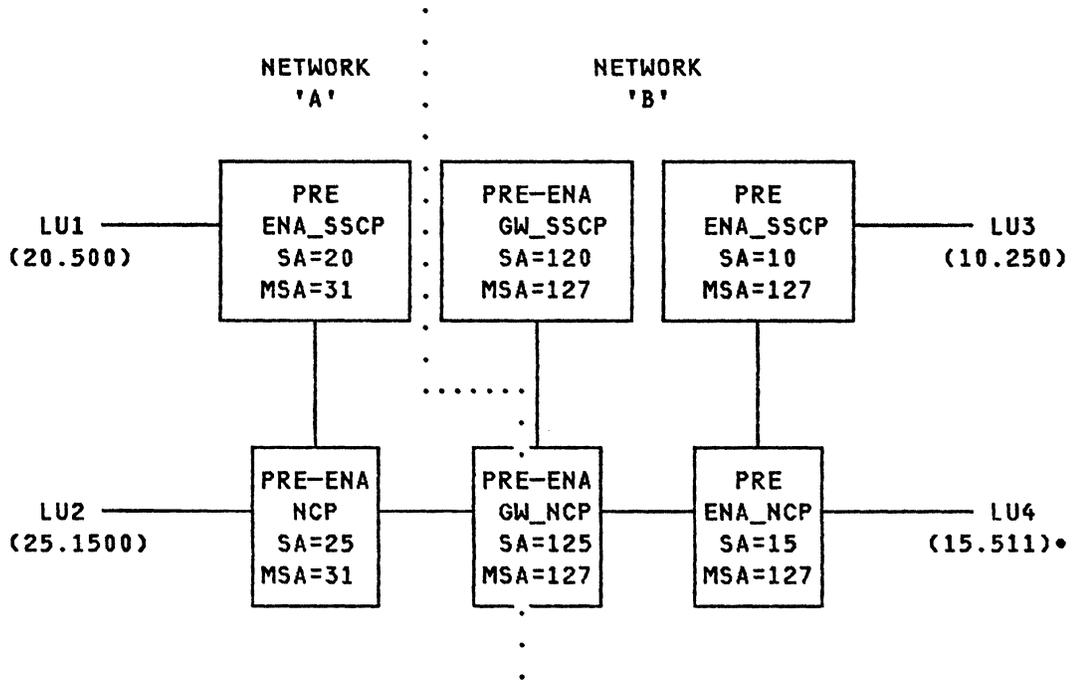


Figure 40. Use of SNI and ENA (Before)

ENA CONFIGURATION

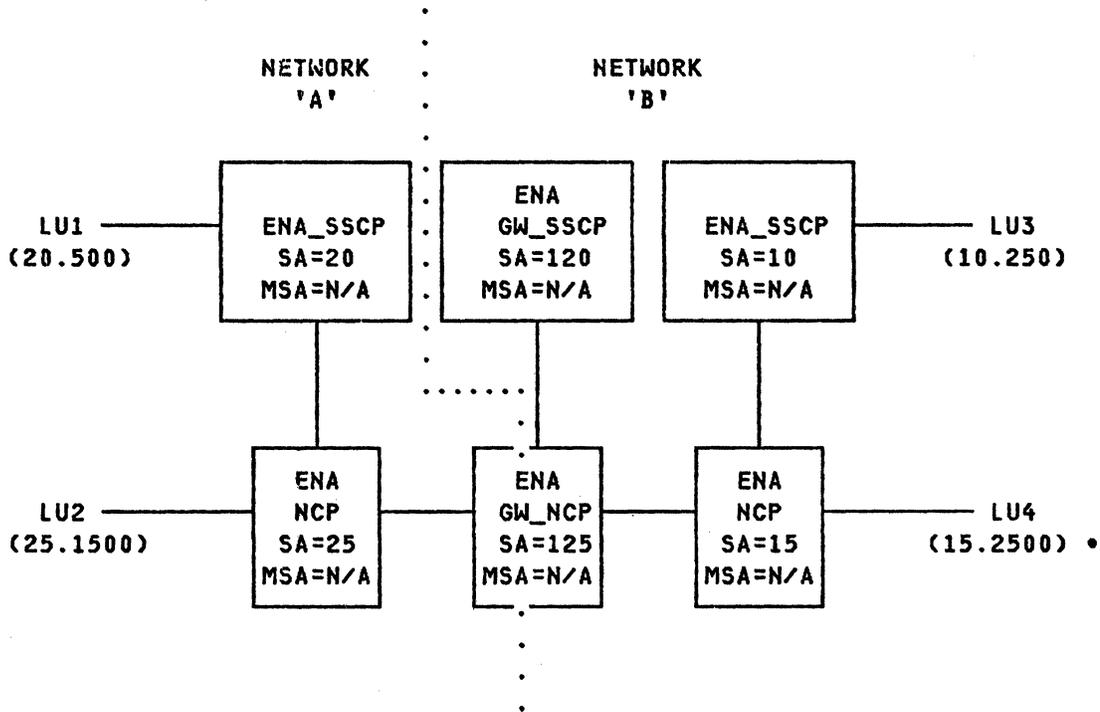


Figure 41. Use of SNI and ENA (After)

OBSERVATIONS

- SNI is not required if both networks migrate to ACF/VTAM V3 and ACF/NCP V4.
- The additional element addresses on SA15 (512-2500) will be able to access all hosts/applications.
- Phasing in these changes dictate keeping the SNI function active until both networks are on ACF/VTAM V3 and ACF/NCP V4.
- It is possible to redefine the network and withdraw the SNI facility once the same addressing split is implemented across both networks.

RECOMMENDATIONS

- One logical ENA network would be more effective for operations and support.

7.8 SCENARIO EIGHT: ENA/SNI/ENA

OBJECTIVES

- Two SNA networks that are totally diverse (i.e., each has a different sub-area split, naming standards, etc.) conventions. The installation currently uses the SNI facility of NCP and VTAM to convert one standard to another for cross-network communications.
- Each network belongs to a separate organization and would like to continue to manage each network separately.
- Network 'A' would also like to expand the number of terminals supported on SA35.

CONSIDERATIONS AND DESIGN CRITERIA

- The differences in the MAXSUBA specification can be eliminated if ACF/VTAM V3 and ACF/NCP V4 are installed.
- The network naming standard differences require the use of the SNI ALIAS mapping function.

PRE-ENA CONFIGURATION

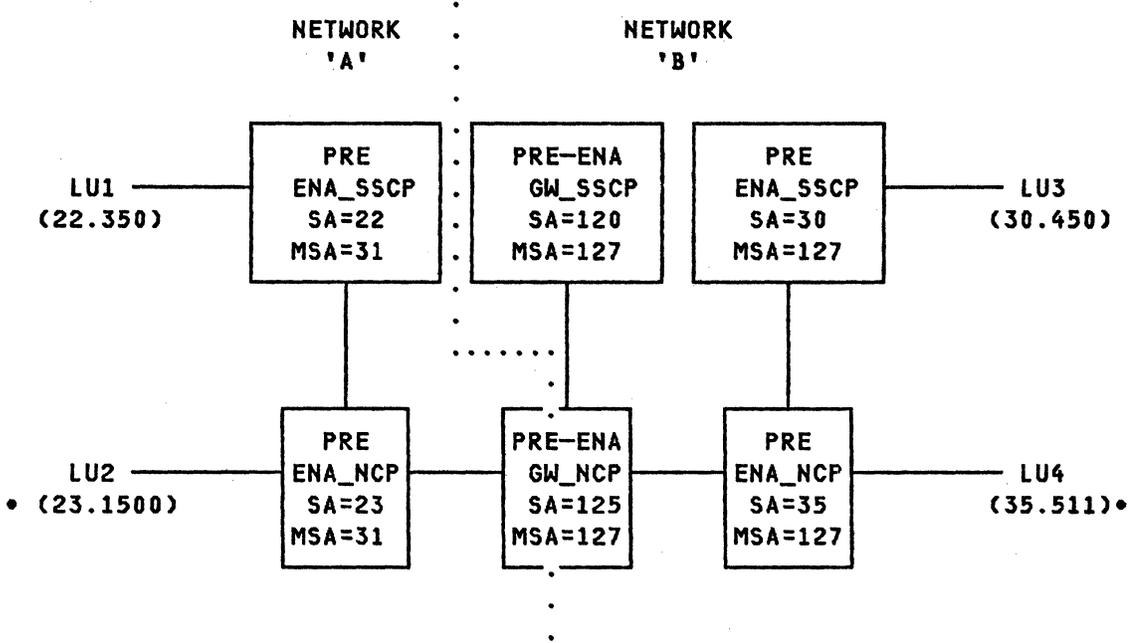


Figure 42. Co-existence of ENA + SNI (Before)

ENA CONFIGURATION

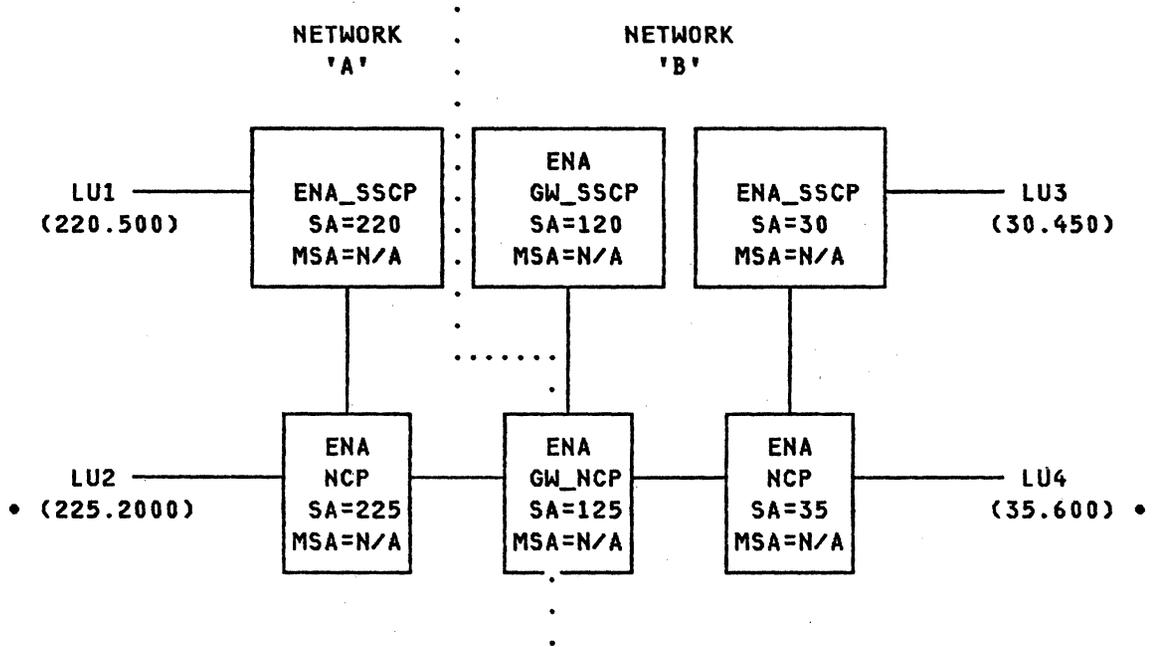


Figure 43. Co-existence of ENA + SNI (After)

OBSERVATIONS

- SNI will continue to perform the ALIAS mapping function (i.e., continue to translate the network names on one network to the standard used on the other).

RECOMMENDATIONS

- Migrating to this environment can be approached in the following sequence:
 1. Migrate Network 'B' SSCP followed by the Network 'B' NCPs. This will provide the capability to define additional resources earlier on Network 'B'.
 2. Migrate Network 'A' SSCP followed by the Network 'A' NCPs.

7.9 SCENARIO NINE: PRE-ENA INTERMEDIATE NODE

OBJECTIVES

- This scenario illustrates the effect on the network when all subarea nodes across the network support ENA with the exception of the IRN.

CONSIDERATIONS AND DESIGN CRITERIA

- The Intermediate Routing Node is running ACF/NCP V2.
- All communication controllers are 3705s.

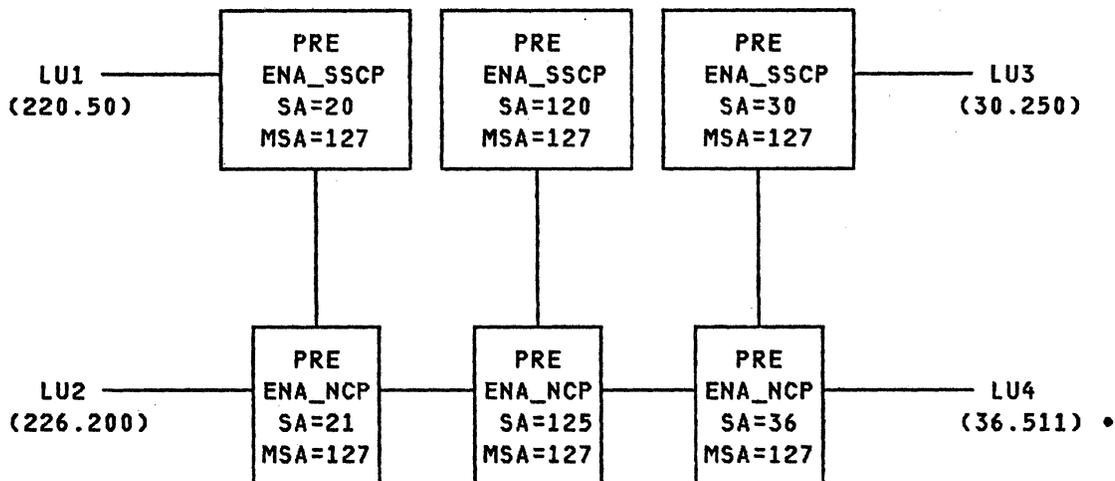


Figure 44. IRN Configuration (Before)

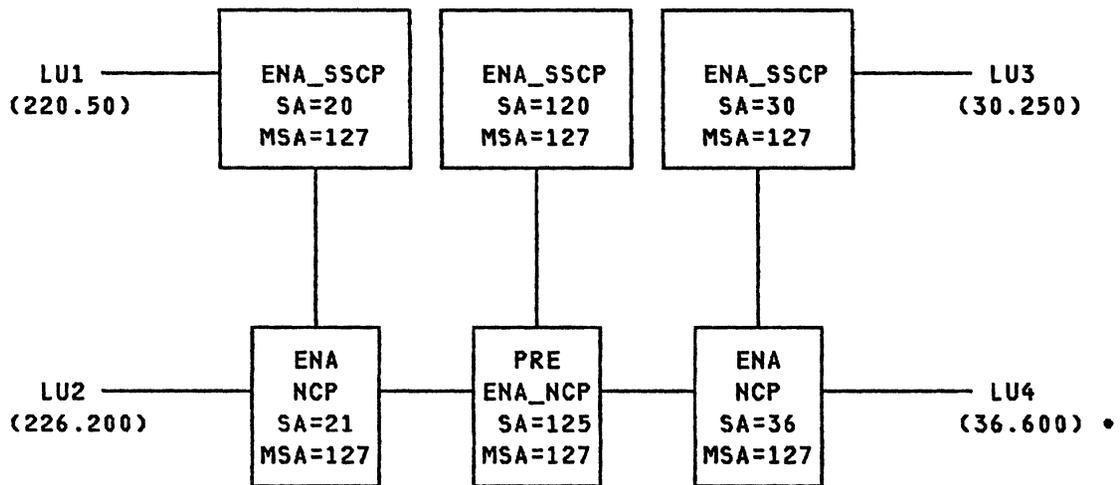


Figure 45. IRN Configuration (After)

OBSERVATIONS

- LU4 can communicate with all nodes since the pre-ENA NCP uses the extended address field (that was introduced in ACF/NCP V1R3) in the FID4 Transmission Header.
- Communication Controller capacity planning was performed to ensure that the existing IRNs could handle the additional traffic associated with the increased number of terminals on SA36.
- The BNN 3705 was replaced with 3725 R2 to support ACF/NCP V4.
- The appropriate PTFs have been applied to allow the IRN to pass the ENA format RUs through the network, before any ENA nodes are given subareas above MAXSUBA.

RECOMMENDATIONS

- This approach is attractive as the IRNs may not require an upgrade from 3705s to 3725s.

7.10 SCENARIO TEN: ENA NETWORK/PRE-ENA BACKUP SSCP

OBJECTIVES

- Having a pre-ENA_SSCP defined as the backup host to an ENA supported host introduces a number of restrictions.
- This scenario describes the communication implications of that configuration.

CONSIDERATIONS AND DESIGN CRITERIA

- Both hosts have the same operating system.
- The normal owner of NCP SA31 is SA30.
- All the applications on SA30 are accessed by all the terminals across the network.
- When SA30 is down, SA20 runs all the applications previously running on SA30.

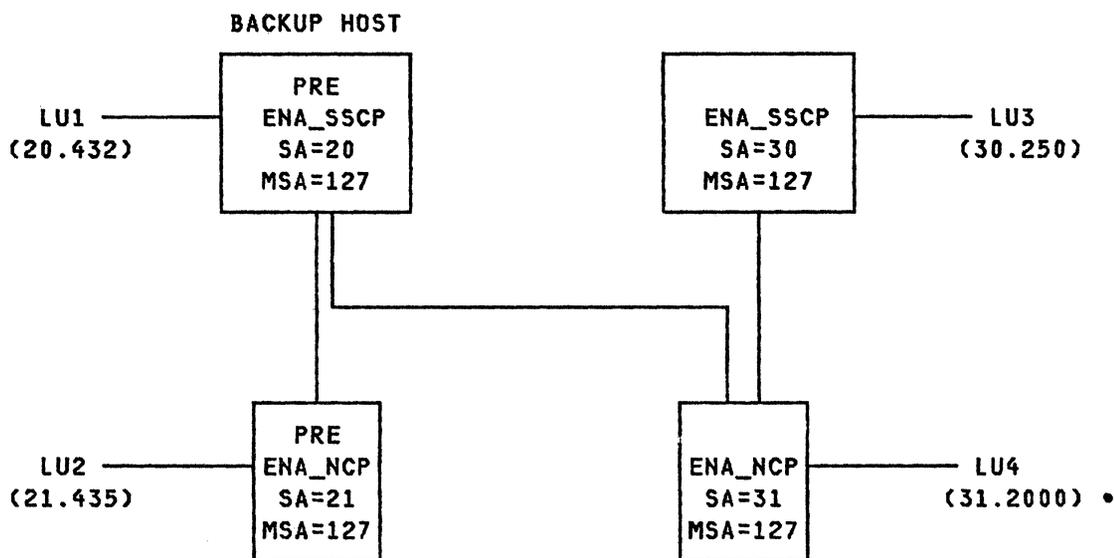


Figure 46. Complete ENA Configuration

OBSERVATIONS

- If SA31 is acquired by SA20, all of the terminals above the subarea/element split range will be unable to access the applications they would access during normal operation.
- If SA31 is activated from SA20, all elements above the subarea/element split range will not be activated.

RECOMMENDATIONS

- This is **not** recommended due to the loss of all ENA function and subsequent loss of service to the terminals assigned to the extended addressing elements.
- Although many locations have 'mirror image' strategies (i.e., the backup system and the production system must be at the same operating system level), the migration to ACF/VTAM V3 may have to occur at different times on each system due to different service requirements, change control restrictions, etc.

In this part, the sections are:

- 8.0 "Overview of Virtual Storage and Addressing" on page 101: Provides basic concepts on Virtual Storage and Addressing in the MVS/370 and MVS/XA environments.
- 9.0 "Storage Management Services" on page 119: Provides an introduction to VTAM's Storage Management Services.
- 10.0 "Virtual Storage in VTAM" on page 127: Provides technical information on VTAM Virtual Storage Allocation, and use in the MVS environment.
- 11.0 "VTAM Virtual Storage Estimation" on page 149: Provides technical information on VTAM Virtual Storage Estimation.
- 12.0 "Monitoring Virtual Storage" on page 159: Provides an approach to monitoring Virtual Storage.

8.0 OVERVIEW OF VIRTUAL STORAGE AND ADDRESSING

This chapter provides some basic concepts on Virtual Storage and Addressing in MVS environments. These concepts should be helpful in understanding the subsequent sections on VTAM's use of virtual storage.

This is not an exhaustive treatment of the subject. However, this section should provide an adequate introduction.

For more information, refer to MVS/Extended Architecture Supervisor Services and Macro Instructions, GC28-1154.

8.1 31-BIT ADDRESSING

The 24-bit address, used in releases of MVS prior to MVS/SP V2, can accommodate a maximum of 16,777,216 bytes (or 16 megabytes). Bits 8-31 of a general register are used for the address. Although the address occupies a four-byte field, the high order byte is not significant.

In MVS/SP V2, the 31-bit address is used. This allows addressing of up to 2,147,483,648 bytes (or 2 gigabytes) of virtual storage. Bits 1-31 of a general register are used for the address. Each address will require a four-byte field. All four bytes are significant in addressing.

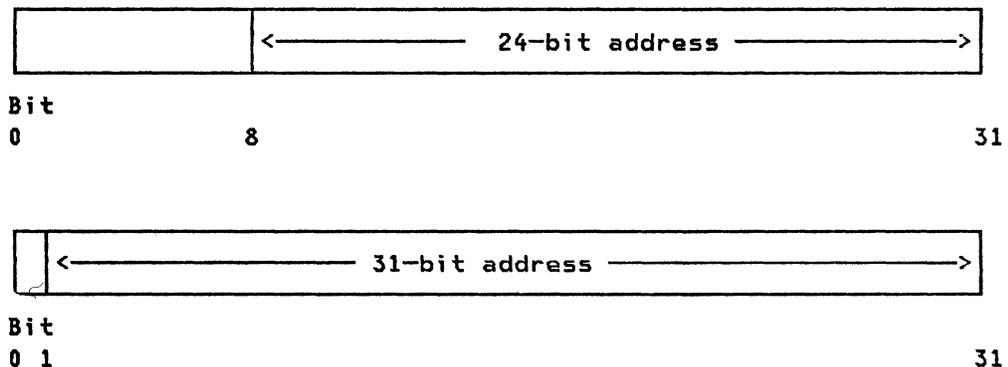


Figure 47. MVS Addressing Formats

In MVS/XA, each program is capable of addressing up to 2 gigabytes of data and instructions within a single address space.

ACF/VTAM V3 is the first release of VTAM to exploit 31-bit addressing.

8.2 BI-MODAL ADDRESSING MODES

For compatibility, programs may use the 24-bit and 31-bit addressing modes available in MVS/XA.

All addresses (i.e. 24-bit or 31-bit) will occupy four bytes in the PSW. The 24-bit address is expanded to occupy 31 bits by inserting seven binary zeros (B'0000000') into bit positions 1-7 of the address field.

The PSW contains a new indicator that defines the addressing mode in use.

Therefore, VTAM application programs can continue to run in the MVS/XA environment with the 24-bit address. These programs will be restricted to addressing up to 16 megabytes of virtual storage; while those using 31-bit addressing can address up to 2 gigabytes.

Programs can switch from one addressing mode to the other during execution.

8.3 VIRTUAL STORAGE

In any computing system, instructions and data are kept in main storage locations for some time during a program's execution. In a system without virtual storage, the program addresses these locations directly (i.e., the program's addressing range is equal to the number of addressable locations in main storage).

In MVS, a program's addressing range exceeds the number of addressable locations in main storage. MVS/370 has a theoretical addressing range of 16 megabytes, regardless of the amount of real storage installed. MVS/XA allows for 2 gigabytes.

However, the amount of virtual storage available to a program under MVS is less than the theoretical addressing range. In MVS, virtual storage may be sub-divided into the following areas:

- System Area
- Private Area
- Common Area

These cumulatively account for all virtual storage available within the theoretical addressing range. Although a program can address all of these areas, it

must be remembered that any increase in the size of any one of these areas will be at the expense of another.

Therefore, under MVS/370:

System Area + Private Area + Common Area = 16 megabytes

Under MVS/XA:

System Area + Private Area + Common Area = 2 gigabytes

8.4 MVS/370 VIRTUAL STORAGE AREAS

In MVS, virtual storage is divided into a number of functional areas. VTAM uses storage from several of these areas. The characteristics of these areas dictate how and where VTAM allocates storage.

Figure 48 on page 104 shows the layout of virtual storage areas under MVS/370. Each area will be described briefly (in order of location).

8.4.1 Common Area

The Common Area provides storage for data and programs that are used by several address spaces in an MVS system. For example, VTAM application programs (e.g., IMS, CICS, TSO) can access buffers in the Common Area created and managed by the VTAM address space.

The Common Area is allocated from the top of virtual storage (i.e., from 16 megabytes) downwards. The most significant parts of the Common Area, in terms of VTAM's usage, are:

- System Queue Area (SQA)
- Link Pack Area (LPA)
- Common Service Area (CSA)

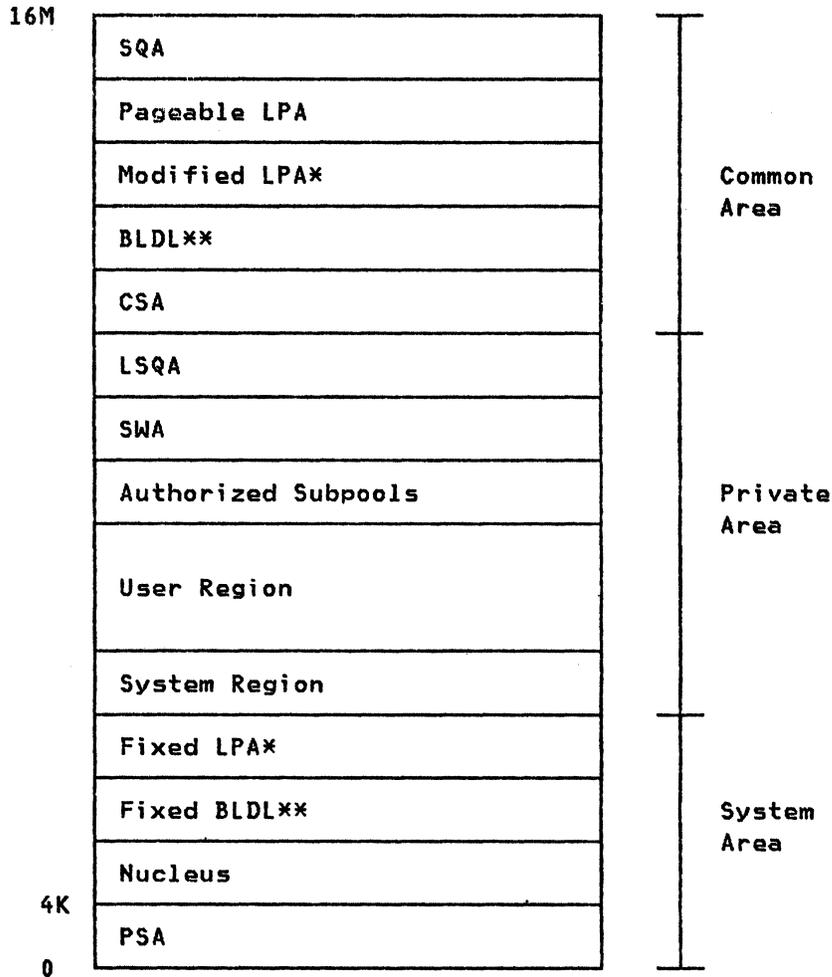
The term 'common storage', used in VTAM publications on storage, refers to storage within the Common Area.

System Queue Area (SQA): SQA contains tables and queues relating to the entire system. Its size varies with the size of the system configuration and job requirements (i.e., IMS regions, CICS regions, TSO address spaces, batch initiators, started tasks).

SQA is allocated from the top of virtual storage (i.e., from the 16 megabyte address downwards) in 64K segments. Its size may be controlled through the MVS 'SQA=' parameter in SYS1.PARMLIB(IEASYSxx) or MVS operator command. It is fixed

in real storage (i.e., not pageable) and reduces the amount of virtual storage available (refer "Virtual Storage" on page 102).

VTAM does not use the SQA directly.



* Optional areas.

** BDL and Fixed BDL are mutually exclusive.

Figure 48. MVS/370 Virtual Storage Layout

Link Pack Area (LPA): LPA contains load modules for SVC routines, access methods and routines and selected user programs. If load modules are accessed frequently and concurrently by several address spaces, they should be placed in LPA for performance reasons (e.g., avoiding program-fetch overhead).

There are three forms of LPA; all are similar in function but differ in terms of storage residency:

- Pageable LPA (PLPA)

Here, load modules may be 'paged in' on demand as these will not always be resident in main storage. The length of time that a load module will remain in main storage will depend on its frequency of use. If the page is not used for a period of time, the system will 'steal' the real storage frame occupied.

PLPA load modules are packed within page boundaries. If a module is larger than 4K, it will be placed on a new page boundary. Smaller modules will be used to fill out the unused portions of the page.

- Modified LPA (MLPA)

This is an optional extension to PLPA where modules under test may be placed. It exists only for the duration of the current IPL and has to be re-built at each IPL.

- Fixed LPA (FLPA)

FLPA is an optional part of the System Area (not the Common Area; but described here for convenience). This allows the MVS user to have LPA modules fixed in real storage to avoid the paging overhead. However, it is usually not used as highly referenced modules in PLPA tend to be resident in main storage (therefore, negating the benefit of fixing modules in real storage).

Obviously, a large FLPA would make less real storage available to the rest of the system.

PLPA is the most significant of these three areas in terms of VTAM's storage consumption. It is allocated from beneath the SQA downwards in 4K blocks. Obviously, its size is dependent on the number of modules included.

VTAM uses about 400K+ of PLPA for load modules.

BLDL: The BLDL is a list of directory entries for load modules residing on a system library. It improves systems performance by avoiding library searches that would otherwise be required to locate a load module.

The pageable form of BLDL is located below PLPA or MLPA, if present. The Fixed BLDL is located directly above the Nucleus. There can only be one of these two forms of BLDL in any system.

VTAM does not use the BLDL directly.

Common Service Area (CSA): CSA contains data areas which are, usually, accessed by several address spaces (e.g., buffers and control blocks required by IMS, CICS, TSO, VTAM, etc.).

CSA is allocated in 4K blocks directly below PLPA.

VTAM places buffers defined through ATCSTRxx (e.g., IOBUF, LFBUF, SFBUF) in CSA. The size of this area can be controlled through the MVS parameter 'CSA=' in SYS1.PARMLIB(IEASYSxx) or MVS operator command.

VTAM's use of this area can be controlled through the 'CSALIMIT=' parameter(s) in ATCSTRxx.

If VTAM cannot acquire sufficient storage in CSA, VTAM could lose messages, fail to initiate/terminate sessions, or enter an interlock condition or some other undesirable state.

In a network host that owns 10,000 nodes, VTAM would use between 1000K-2000K of CSA approximately. The amount used would vary with VTAM's activity over time.

8.4.2 Private Area

Each address space in an MVS system has its own private area. Data in the Private Area is primarily for the use of the task owning the address space. For example, VTAM keeps several control blocks and work areas that are not referenced directly by VTAM application programs in the Private Area.

The size of the Private Area will be:

$$\text{Private Area} = 16\text{M} - \text{Common Area} - \text{System Area}$$

Depending upon the installation, the size of the Private Area may vary between 5-10 megabytes. This is a major factor that determines how large a network may be defined to VTAM, for that installation.

The Private Area is allocated in two directions:

- from the top of the System Area up; and
- from the bottom of the Common Area down.

VTAM uses all parts of the Private Area:

- Local System Queue Area (LSQA)
- Scheduler Work Area (SWA)
- Authorized Subpools
- User Region
- System Region

The term 'private storage', used in VTAM publications, refers to storage within the Private Area.

Local System Queue Area (LSQA): This contains tables and queues unique to VTAM's address space (e.g., private segment and private page tables, Region Control Task control blocks) required by MVS.

LSQA is allocated from the top of the Private Area downwards; it may be intermixed with SWA and Authorized Subpool pages.

VTAM requires about 30K for its LSQA.

Scheduler Work Area (SWA): This contains control blocks that must exist from task initiation to task termination for the VTAM address space. It provides information for MVS's work scheduling mechanisms.

SWA is allocated from the top of the Private Area downwards; it may be intermixed with LSQA and Authorized Subpool pages.

VTAM requires about 60K for the SWA.

Authorized Subpools: These subpools are located high in the Private Area.

Prior to ACF/VTAM V2R2, SP229/230 was used for VTAM control blocks and work areas required in dynamic conditions (e.g., network activation/deactivation, recovery of a large node).

These subpools are allocated from the top of the Private Area downwards and may be intermixed with LSQA and SWA pages.

Without ITLIM controls, VTAM would attempt to obtain as much of SP229/230 as it required during dynamic conditions. The quantity required would vary with the number of nodes involved in activation/deactivation or recovery. Failure to obtain storage in SP229/230 resulted in activation/deactivation failure or incomplete error recovery. The maximum storage available to these subpools is:

SP229/230 = Private Area - LSQA - SWA - User Region - System Region

Since ACF/VTAM V2R2, this storage is obtained from subpools within the User Region (i.e., SP250). This removes the contention between the Authorized Subpools and the User Region during network activation/deactivation and error recovery.

User Region: This contains the VTAM USSTABs, MODETABs, user exit routines and some standard VTAM modules. It also contains control blocks that define the network configuration. VTAM uses about 1500K of the User Region for base code. Additional storage from the User Region will be required for defining the network configuration (e.g., an additional 1500K to define a network of approximately 10,000 nodes).

The User Region is allocated from the top of the System Region upwards. The User Region size can be specified through the 'REGION=' parameter on the MVS JCL JOB or EXEC statement. If this JCL parameter was not specified, the default size specified in the installation's JES parameters would apply. In MVS/370, the specified value (or the JES default) may be checked and reduced by the installation's IEALIMIT exit (an optional MVS exit). This value will determine the amount of virtual storage that may be acquired within the User Region.

From ACF/VTAM V2R2 on, VTAM obtains storage for dynamic activities from SP250 in the User Region (instead of SP229/230 which is above the User Region). While this reduces the Authorized Subpool requirement, this has the effect of increasing the User Region requirement.

Therefore, it is important to note that:

- (for releases prior to ACF/VTAM V2R2) as the User Region gets larger, the amount of storage available to SP229/230 will reduce; but MVS will use SP250 instead. The size of the User Region should not be larger than necessary.
- (for ACF/VTAM V2R2 and subsequent releases) the User Region has to be increased to accommodate the additional storage required for dynamic activities; review the size of the User Region. SP250's expandability is constrained by the 'REGION=' JCL parameter, IEALIMIT or IEFUSI specification.

System Region: This is used by MVS system functions performing work for VTAM under the Region Control Task. About 24K of storage is used here.

8.4.3 System Area

The System Area is allocated from the bottom of virtual storage during the MVS IPL process. This area is fixed in real storage and allocated prior to the Common Area and Private Areas. The Nucleus is the most significant portion of the System Area.

VTAM uses about 3K of the Nucleus for I/O modules.

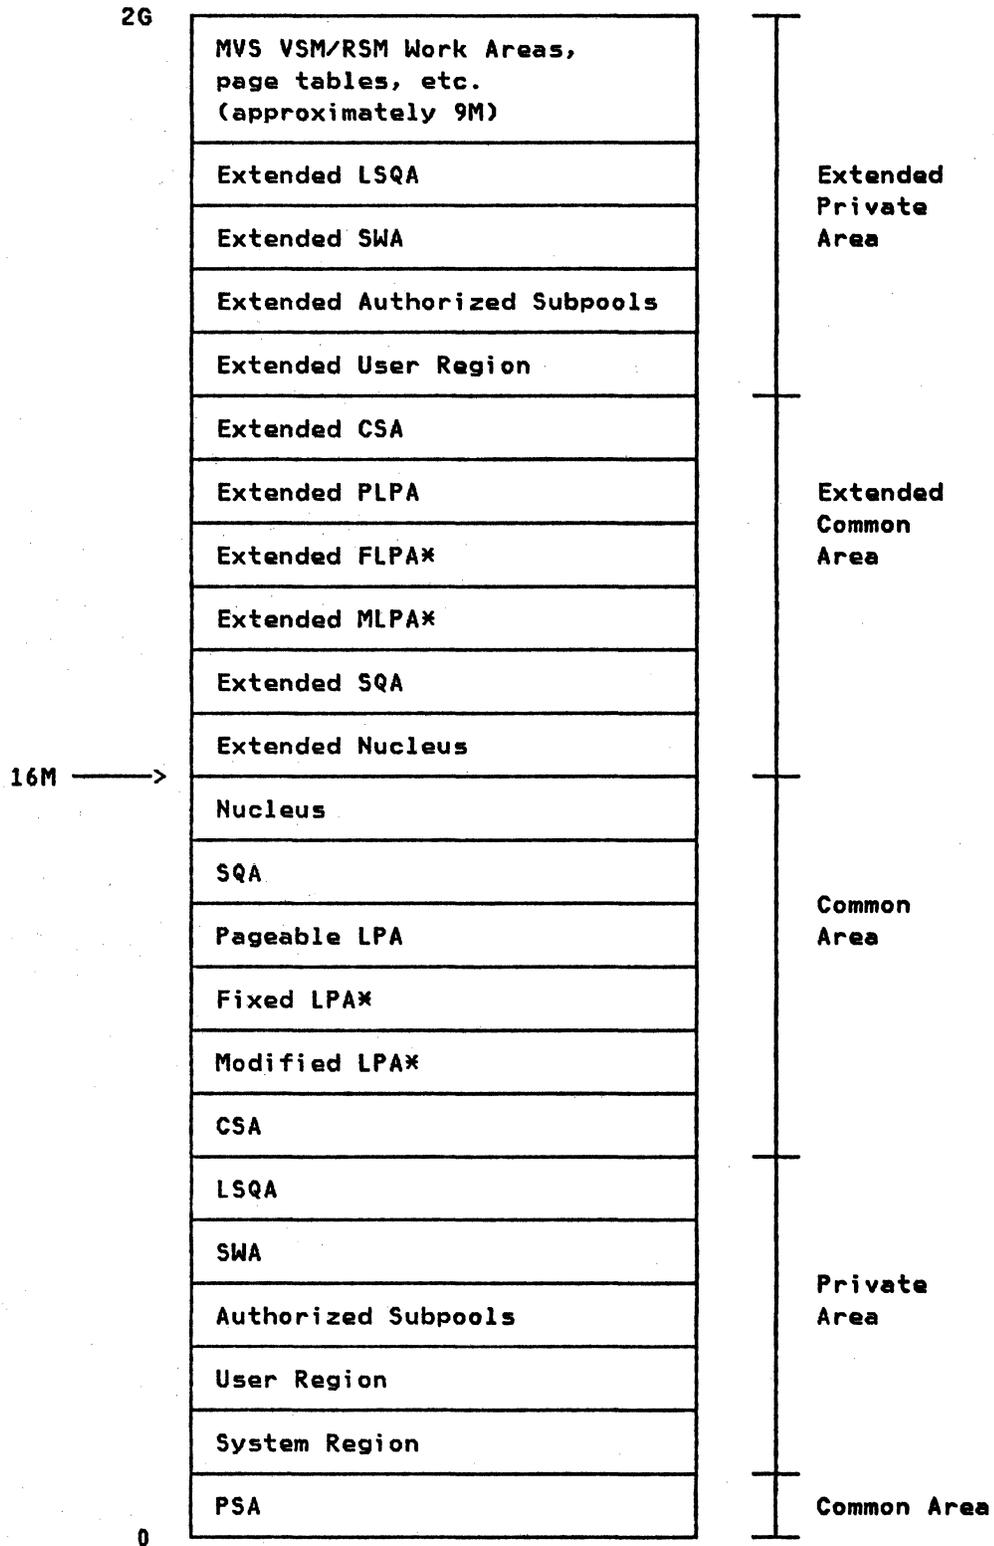
8.5 MVS/XA VIRTUAL STORAGE AREAS - CHANGES FROM MVS/370

In MVS/XA, virtual storage is divided into similar functional areas as for MVS/370 (see "MVS/370 Virtual Storage Areas" on page 103). However, to exploit the additional virtual storage offered by 31-bit addressing, most of these functional areas have an 'extension' above the 16 megabyte boundary. For example, there is 'CSA' below the boundary and 'Extended CSA' above. Although both areas provide identical functions, only programs executing in 31-bit addressing mode will be able to use the 'extended' area.

Figure 49 on page 109 shows the virtual storage layout for MVS/XA.

In ACF/VTAM V3, all VTAM data areas remain in the same functional areas as in previous releases. However, VTAM will use the 'extended' portion of each functional area in preference to the portion below the boundary (provided there are no 24-bit addressing dependencies). For example, while most of VTAM's buffers will move above the boundary, the IOBUFs will remain below as these are used by channel programs that are still 24-bit dependent.

The concepts presented in "MVS/370 Virtual Storage Areas" on page 103 will still apply to MVS/XA. However, the following changes are significant.



* Optional areas.

Figure 49. MVS/XA Virtual Storage Layout

8.5.1 Quantity of Virtual Storage

Under MVS/XA, each address space can address up to 2 gigabytes of virtual storage instead of the 16 megabytes available under MVS/370.

This alleviates most of the constraints in CSA, SP229/230 and the User Region experienced by VTAM users under MVS/370.

8.5.2 Paging

Virtual Storage requires that every block of virtual storage ('page') is backed up by a block of auxiliary storage ('slot'). This reduces the quantity of real storage blocks ('frames') needed. The process of transferring data between frames and slots is called paging.

The Virtual Storage Constraint Relief offered in ACF/VTAM V3 provides the opportunity to define networks with an extremely large number of nodes. This will lead to a significant increase in the use of virtual storage by VTAM. Therefore, network expansion planning should include the following considerations on paging characteristics:

- Extra auxiliary storage could be required to back up the additional pages that would result from exploitation of additional virtual storage.

The quantity of slots may have to be increased.

- As more pages are used, the probability of having an insufficient quantity of real storage frames will be increased (all things being equal). This shortage will delay the paging process which will, in turn, delay VTAM's responsiveness in network activation and deactivation, session establishment and termination, sending and receiving data, and error recovery.

This 'shortage' could be alleviated by:

- installing more real storage; and/or
- reducing the paging time with more and/or faster devices; and/or
- using VTAM control parameters (e.g., ITLIM) to reduce the concurrency of VTAM's activities and hence, the paging demand.

8.5.3 Extensions to Virtual Storage Areas

Each of the virtual storage areas has a counterpart that resides above the 16 megabyte boundary. Exceptions are the PSA and BLDL.

However, only programs executing in 31-bit addressing mode will be able to use the extended areas.

Note that if a dump of a virtual storage area is requested, both the base and its extension will be dumped.

8.5.4 Segment Boundaries

The segment size in MVS/XA is 1 megabyte. MVS/370 uses a segment size of 64K. Therefore, where storage is allocated in segments, it will be allocated in 1 megabyte blocks under MVS/XA.

This is only significant in its effect on storage estimation. The amount of storage lost may increase when storage allocation is 'rounded off' to 1 megabyte instead of 64K boundaries.

8.5.5 REGION, IEALIMIT and IEFUSI

In MVS/370, the size of the User Region would be controlled by the 'REGION=' JCL parameter and the IEALIMIT exit (as discussed in "User Region" on page 107). If this JCL parameter was not specified, the JES default REGION size would be used.

In MVS/XA, the function of checking and modifying the User Region size may be performed by the MVS/SMF IEFUSI exit. This is only significant to understanding what may affect the size of VTAM's User Region.

Prior to MVS/XA 2.1.2, the 'REGION=' JCL parameter controlled only the User Region below the 16 megabyte boundary. It had no effect on the Extended User Region.

Since MVS/XA 2.1.2, this parameter works as follows:

- The REGION size can be expressed in kilobytes or megabytes.
- If the REGION size exceeds 16M, MVS will attempt to obtain extended virtual storage. In this case, the default REGION size will still apply below the boundary. This should be specified for installations requiring VSCR.
- If the REGION size is less than 16M, MVS will obtain storage below the boundary. The default size for the extended REGION will be 32M.
- IEALIMIT may be used to override the User Region size below the boundary.
- IEFUSI may be used to override the User Region size on either side of the boundary.

8.5.6 Relocation of Virtual Storage Areas

In MVS/XA, the Nucleus and its extension are located 'side by side' on the 16 megabyte boundary.

The fixed LPA and fixed BLDL are not part of the System Area under MVS/XA. The LPAs (i.e., PLPA, FLPA, MLPA) and BLDL are in contiguous storage above CSA. Their extensions are located below the Extended CSA.

Note that the Linklist Look Aside (LLA) function in MVS/SP V2R1.1 removes the need for BLDL.

8.6 ALLOCATION OF EXTENDED VIRTUAL STORAGE

In MVS, the GETMAIN macro instruction is used to obtain storage. GETMAIN allows the program to specify whether the storage should be obtained from above or below the 16 megabyte boundary. This is specified through the LOC parameter:

LOC=BELOW The storage must be obtained from below the boundary.

LOC=ANY The storage may be obtained on either side of the boundary. Note that this does not guarantee allocation of extended storage. In general, storage will be allocated above the boundary. However, other factors may preclude this (e.g., current allocations, REGION specification, etc.).

LOC=RES The storage is to be obtained according to the residency mode of the requestor. If the program resides below the boundary, the storage must be obtained from below the boundary. Otherwise, the storage may be obtained from anywhere.

8.7 FIXING VIRTUAL STORAGE

Data areas may be moved in and out of real storage (see "Paging" on page 110). This can result in undesirable paging activity and page faulting (i.e., a program's execution is interrupted because required data is out of real storage).

To avoid this, a page may be fixed. This can be done through the MVS system PGSER, PGFIX or PGFIXA macro instructions.

These macro instructions are similar in function. They make a page ineligible for a 'page out' while the address space is swapped in. VTAM uses this facility.

PGSER was introduced in MVS/XA to consolidate all paging services. PGFIX and PGFIXA are supported by MVS/XA to maintain compatibility with MVS/370.

8.8 STORAGE PROTECTION

In any system capable of executing several programs concurrently, a program must be prevented from accessing areas of virtual storage 'owned' by another program. MVS allows storage protection through the 16 storage protection keys allowed by the System/370 PSW and hardware.

In MVS, keys 0-7 are assigned to system programs and keys 8-15 are reserved for user programs.

- 0 MVS Supervisor (and other system functions). Authorized user programs executing in supervisor state may obtain this key.
- 1 Job scheduler and job entry subsystem (i.e., JES2, JES3).
- 2 VSPC
- 3-4 Reserved.
- 5 Data Management (e.g., IOS, ASM, OPEN/CLOSE/EQV).
- 6 VTAM and TCAM.
- 7 IMS
- 8 All user programs using virtual storage (i.e., V=V storage - this is storage as described in "Virtual Storage" on page 102). As each user program occupies a separate address space and each address space has its own segment and page tables (used for translating virtual addresses to real addresses), a program cannot access virtual pages allocated to another address space.
- 9-15 All user programs using virtual=real storage (i.e., V=R storage - this is storage where virtual addresses are the same as real addresses). Note that V=R programs are not protected by address translation (described above) and, therefore, must each have a unique key (i.e., 9-15).

Storage protect keys are determined by either:

- the job key in the TCB, or
- the user key in the PSW, or
- the key predefined for that subpool.

VTAM uses keys 0 and 6.

SUB- POOL	STORAGE TYPE	OWNER (KEY)	FETCH PRO- TECT	PAGE- ABLE	REAL FRAME LOCN.	COMMENTS
0- 127	PVT/EPVT (Low)	Job (TCB)	Yes	Yes	ANY	These are the only SP numbers requestable by user programs.
128- 225						Reserved or undefined.
226	SQA	None (0)		No	BELOW	
227	CSA/ECSA	None (PSW)	Yes	No	ANY2	
228	CSA/ECSA	None (PSW)	No	No	ANY2	
229	PVT/EPVT (High)	Job (PSW)	Yes	Yes	ANY	
230	PVT/EPVT (High)	Job (PSW)	No	Yes	ANY	
231	CSA/ECSA	None (PSW)	Yes	Yes	ANY	
232						Reserved.
233	LSQA/ELSQA	Job (0)	No	No	ANY	Swappable. Allocated from SP255.
234	LSQA/ELSQA	Step (0)	No	No	ANY	Swappable. Allocated from SP255.
235	LSQA/ELSQA	Memory (0)	No	No	ANY	Swappable. Allocated from SP255.
236- 237	SWA/ESWA	Job (1)	No		ANY	System use only.
238						Reserved.

Figure 50. MVS Virtual Storage Subpools 0-238

SUB- POOL	STORAGE TYPE	OWNER (KEY)	FETCH PRO- TECT	PAGE- ABLE	REAL FRAME LOCN.	COMMENTS
239	SQA/ESQA	None (0)	Yes		ANY	
240	PVT/EPVT (Low)	Job (TCB)	Yes	Yes	ANY	Allocated from SP250.
241	CSA/ECSA	None (PSW)	No	Yes	ANY	
242- 244						Reserved.
245	SQA/ESQA	None (0)	No	No	ANY	
246- 249						Reserved.
250	PVT/EPVT (Low)	Job (TCB)	Yes	Yes	ANY	Allocated from SP0 in SUP/KEY0.
251	PVT/EPVT (Low)	Step (TCB)	Yes		ANY	Used for modules in low storage but not in SP252.
252	PVT/EPVT (Low)	Step (0)	No		ANY	Used for modules re-enterable and authorized.
253	LSQA/ELSQA	Job (0)	No	No	ANY	Swappable. Allocated from SP255.
254	LSQA/ELSQA	Step (0)	No	No	ANY	Swappable. Allocated from SP255.
255	LSQA/ELSQA	Memory (0)	No	No	ANY	Swappable.
	NUC/ENUC PLPA/EPLPA FLPA/EFLPA MLPA/EMLPA		No No No No	No Yes No Yes	ANY ANY ANY2 ANY	

Figure 51. MVS Virtual Storage Subpools 239-255

8.9 SUBPOOLS

A subpool is a group of logically related storage blocks sharing some common characteristics. Each subpool is identified by a unique number (0-255). In MVS/XA, all subpools (except SP226) exist above and below the 16 megabyte virtual storage boundary. Therefore, a GETMAIN request can return virtual storage on either side of the boundary. However, if LOC=BELOW is specified on the request, storage will be allocated from below the boundary.

Figure 50 on page 114 and Figure 51 on page 115 show the virtual storage subpools that exist for MVS/370 and MVS/XA. This table should be read as follows:

SUBPOOL	The number assigned as the subpool identifier.
STORAGE TYPE	Indicates the area of virtual storage where this subpool resides. Figure 52 on page 118 depicts this.
OWNER (KEY)	OWNER designates whether the storage allocated from this subpool is placed under the ownership of: JOB the task that is currently executing, or STEP the current step within the task, or MEMORY the address space containing the task. (KEY) indicates how the storage protect key is assigned. 0 Key 0 is used. This is for the MVS system control program and authorized tasks. 1 Key 1 is used. This is for the MVS Job Scheduler and JES2/JES3. PSW The key is taken from the PSW at the time of GETMAIN or can be specified on the GETMAIN request. TCB The key is taken from the TCB in use at the time of the first GETMAIN request. All subsequent GETMAIN requests will use this key (even if the TCB key is different).
FETCH PROTECT	Determines whether this subpool can be fetched by tasks with different storage protect keys.
PAGEABLE	Indicates whether this subpool is pageable or fixed.
REAL FRAME LOCN.	Indicates whether the real frames required to back up virtual storage pages are allocated from above or below 16 megabytes (in systems with more than 16M of real storage installed). ANY The frames can be anywhere in real storage. Fixed pages allocated with LOC=BELOW will be backed up below 16M. ANY2 If the virtual address is greater than 16M, the backing frame can be anywhere. Else, the backing frame will be below 16M unless GETMAIN LOC=ANY was requested. BELOW The backing frames will always be below 16M.

Figure 52 shows the subpools according to the virtual storage area in which they reside. This table will assist in locating the VTAM data areas.

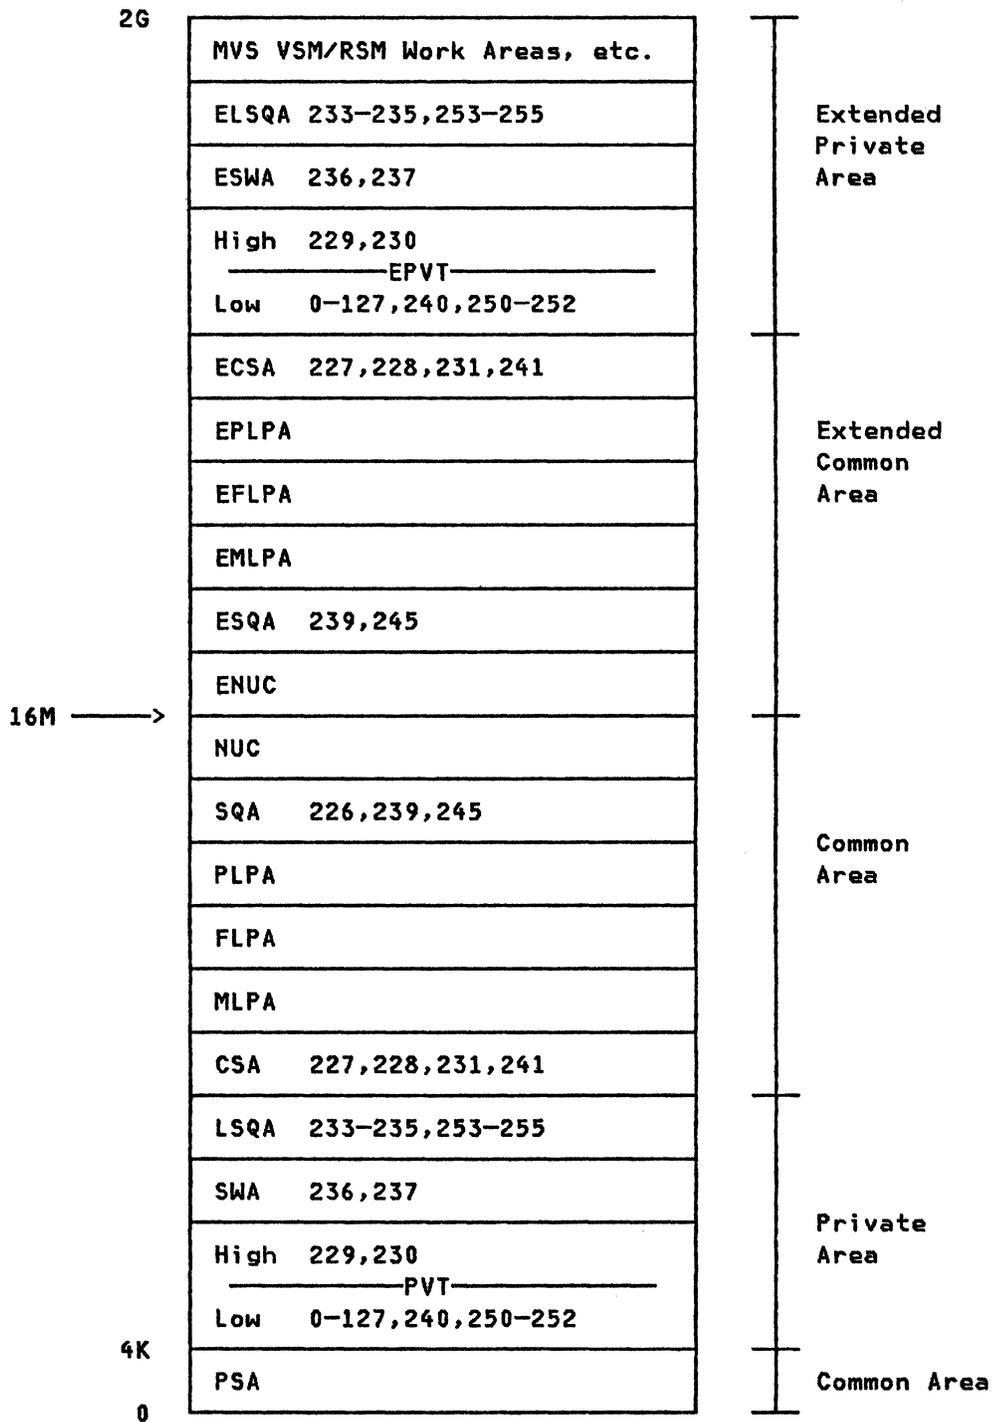


Figure 52. Virtual Storage Area - Subpool Relationships

9.0 STORAGE MANAGEMENT SERVICES

Storage Management Services (SMS) is the VTAM component that interfaces with the MVS Virtual Storage Manager (VSM) to provide other VTAM components with storage for control blocks and buffers. SMS is responsible for the following services:

- Building and deleting buffer pools (i.e., pools containing fixed-length data areas) from MVS storage subpools.
- Obtaining and releasing buffers from these pools.
- Expanding and contracting these pools.
- Getting and freeing variable length storage areas from MVS storage subpools.
- Managing storage during abnormal termination conditions. When VTAM terminates, SMS will release any CSA storage that is still allocated to VTAM (see "VTALLOC/VTFREE" on page 124).

VTAM components request services from SMS through an internal macro interface. The macros are:

BLDPOOL	Invoked during VTAM initialization to build a pool of fixed length data areas (from an MVS storage subpool) and create a directory of pool entries.
DELPPOOL	Invoked during VTAM termination to release storage allocated through BLDPOOL.
REQSTORE	Obtains a fixed-length block of storage from a pool.
RELSTORE	Returns storage acquired through REQSTORE.
VTALLOC	Obtains variable-length storage from an MVS storage subpool.
VTFREE	Returns storage obtained by VTALLOC.
GETBLK	Obtains storage from a storage pool of control blocks (preallocated by VTALLOC).
FREEBLK	Returns storage obtained by GETBLK.

To provide its services, SMS uses the following control blocks:

BPDY	Buffer pool directory.
BPENT	Buffer pool entry.
BPCB	Buffer pool control block.
PXB	Buffer pool extension block.
BFPFX/BFHDR	Buffer prefix/buffer header.
SMHDR	Prefix for common storage obtained through VTALLOC.
SCHDR	Prefix for other storage obtained through VTALLOC.
SPANC	Anchor block for a storage pool of control blocks.
PAGTB	Page table for groups of pages within a SPANC.

9.1 BUILDING AND DELETING BUFFER POOLS

Buffer pools are built primarily from specifications supplied through VTAMLST (or as overridden on the START NET command). These are supplied through the Buffer Pool Start Options in ATCSTRxx (see VTAM Installation and Resource Definition, SC23-0111):

poolname=(baseno,bufsize,slowpt,F,xpanno,xpanpt)

These options are specified for the CRPLBUF, IOBUF, LFBUF, SFBUF, SPBUF and WPBUF pools.

As VTAM initializes, BLDPPOOL macro instructions are issued to create buffer pools. BLDPPOOL calls the buffer pool build routine - ISTORFPO.

On the first invocation of ISTORFPO, the buffer pool directory (BPDTY) and a pool of work elements (i.e., SMS1 pool) are built in CSA/ECSA Fixed SP231. The BPDTY address is stored in the ATCVT (VTAM's main control block).

After that, the following buffer pools are built:

IOBUF	For input/output data.
PPBUF	Reserved.
LPBUF	For CRAs and NSCPLs.
WPBUF	For FMCBs.
NPBUF	Reserved.
LFBUF	For LULBs.
CRPLBUF	For copies of RPLs.
SFBUF	For LUCBs.
SPBUF	For LMPCBs.
APBUF	Reserved.

These pools are built in the following manner:

- SMS will issue GETMAINS to obtain storage from an MVS storage subpool. The amount of storage required for the pool will be determined from the product of **baseno** and **bufsize** as specified in the Buffer Pool Start Option.
- If the buffer pool is to be fetch protected (VTAM makes this decision internally), then CSA/ECSA SP231 is used. Else, CSA/ECSA SP241 is used.
- The buffer pool will be fixed (i.e., made non-pageable) if the **F** operand is specified on the Buffer Pool Start Option. If the pool should be fixed, SMS will issue PGSERs to fix the pages allocated.
- For each pool, an entry is made in the BPDTY and a buffer pool control block (BPCB) created.
- Each buffer in the pool is initialized with a buffer prefix/buffer header (BFPFX/BFHDR). These precede the data portion of each buffer and are used to chain buffers together.

When all pools have been created, ISTOREFO builds the CBID Table (CBIDT) within the BPDY. On each REQSTORE request, the CBIDT is searched to select the pool that will satisfy the request.

Figure 53 illustrates the relationships between the data areas after the BLDPOOLs have completed.

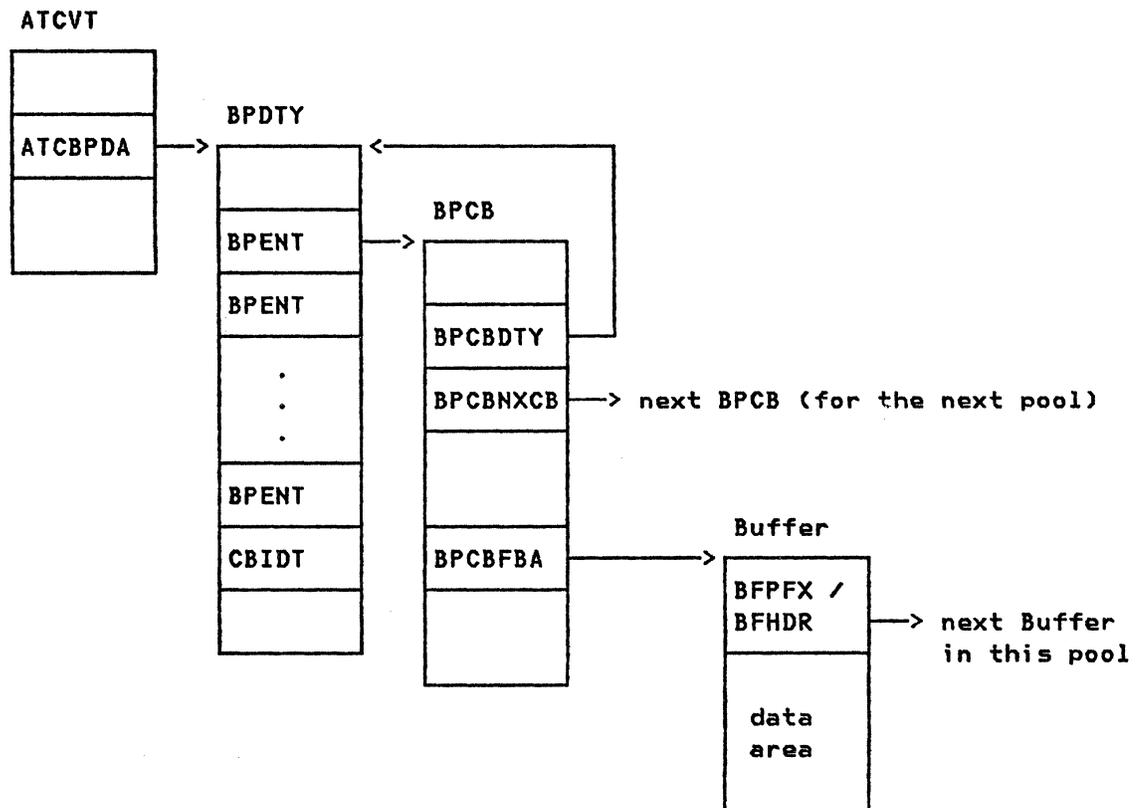


Figure 53. SMS Data Area Relationships after BLDPOOL

During VTAM termination, DELPOOL macro instructions are issued to delete buffer pools created through BLDPOOL. This is processed as follows:

- If the pages in the pool were fixed, SMS will issue PGSERs to free them.
- SMS will issue FREEMAINS to return the pages to the MVS storage subpool.
- The BPCB and BPDY entries for the buffer pool are deleted.

When all buffer pools are deleted (i.e., all BPDY entries removed), ISTOREFO deletes the SMS1 pool, BPDY and the BPDY pointer in ATCVT.

9.2 OBTAINING AND RELEASING BUFFERS

REQSTORE requests invoke the module ISTORMBA to obtain buffers from the pools. The request may be for:

- a particular control block - in this case, one buffer from the appropriate pool is allocated (e.g., a REQSTORE for a CRA will result in the allocation of a LPBUF); or
- a number of buffers from the IOBUF pool.

Most buffers are obtained by dequeuing off the 'available buffer' queue anchors in the BPCBs (i.e., BPCBFBA) and PXBs (i.e., PXBFBA).

However, buffers to hold CRAs and CRPLs (i.e., LPBUFs and CRPLBUFs respectively) are dequeued off 'quick cell' queue anchors (i.e., ATCCRA, ATCCRPL) in the ATCVT. These buffers are added to 'quick cell' queues as they are returned to the pool for reuse.

A REQSTORE request may be issued in non-priority or priority mode. For example, a request for IOBUFs to hold inbound data from a channel-attached NCP will be a priority request. SMS will treat each request according to its priority.

- A non-priority request will not be honored if the remaining number of buffers is below (or will be below) the slowdown threshold (i.e., the **slowpt** value in the Buffer Pool Start Options). The request is then queued off the BPCB (via BPCBRPHB) and will be honored later when buffers become available.
- A priority request will be honored whenever buffers are available. This type of request is not affected by 'slowdowns'. In fact, the main purpose of the **slowpt** operand in the Buffer Pool Start Options is to reserve that amount of buffers for priority requests. However, if there are no buffers available, the priority request will be queued off the BPCB (via BPCBRPHA).

Buffers are returned through RELSTORE requests.

9.3 EXPANDING AND CONTRACTING BUFFER POOLS

VTAM's buffer requirements vary over time. For example, the demand for NCSPs and CRA/RPHs (i.e., LPBUFs) should peak during network activation. The demand for copied RPLs and FMCBs (i.e., CRPLBUFs and WPBUFs respectively) will grow as LU-LU sessions are established with application programs. CRPLBUFs and IOBUFs will be required when application programs SEND data to and RECEIVE data along these LU-LU sessions. The buffer pools are expanded and contracted to meet changing demands. This process is influenced by the **xpanno** and **xpanpt** operands of the Buffer Pool Start Options.

Buffer pools are expanded in this way:

- When the number of available buffers in a pool reaches or falls below the **xpanpt** value, module ISTORFPX is invoked to expand the pool.

- ISTOREPFX will allocate a number that is dependent on the xpanno value. The storage is obtained on a page boundary (i.e., if the bufsize is less than a page, a value of 1 will result in the allocation of a whole page of buffers).
- The additional buffers allocated are collectively referred to as an extent. A pool extension block (PXB) is created for each extent. It is used to anchor a queue of free buffers within the extent.

Figure 54 shows the data area relationships after expansion has occurred.

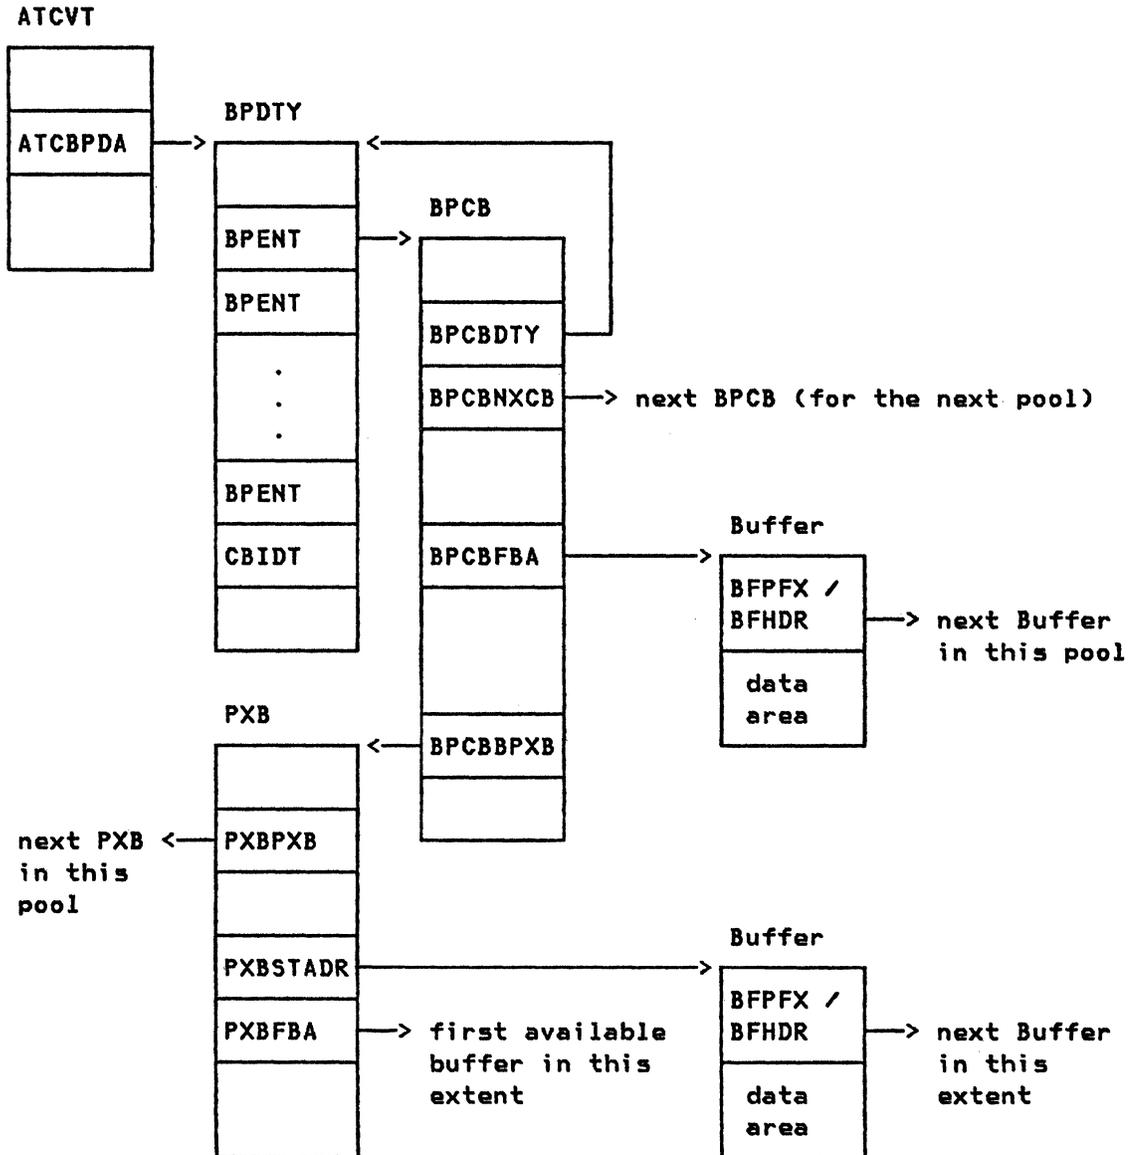


Figure 54. SMS Data Area Relationships after Pool Expansion

An expanded pool will be contracted when demand has fallen:

- SMS calculates the contraction threshold value as:

$$\text{xpanpt} + (2 \times \text{no. of buffers in an extent})$$
- When the number of available buffers in a pool reaches or exceeds this value, ISTORPFX is scheduled to free all unused extents.

9.4 GETTING AND FREEING VARIABLE LENGTH STORAGE

When a VTAM component requires storage for a data area that is not contained within a buffer, it issues a VTALLOC or GETBLK request. Generally, VTALLOC is issued to obtain storage for a control block pool or for a data area that is not kept within a pool. GETBLK is used to obtain storage from a control block pool.

VTALLOC/VTFREE: When VTALLOC is issued, the module ISTORMAF is invoked to GET-MAIN the storage from MVS. Storage obtained from CSA or SQA will be prefixed with an SMHDR. If it is from the VTAM private area, the SCHDR is used. These headers contain information on area length, storage protection key and subpool number.

The SMHDR is also used to chain storage areas together. CSA and SQA storage is chained off the ATCVT 'obtained storage' anchor - ATCOROBT.

When a VTFREE is issued to free CSA and SQA storage, the storage area is chained off the 'storage to be freed' anchor - ATCORTBF - in the ATCVT. The module ISTORMMG is scheduled (via an SRB) to issue FREEMAINS for these areas. Private storage is freed directly by ISTORMAF.

GETBLK/FREEBLK: Many of VTAM's frequently used control blocks are kept within pools. These are packed into pages to improve page reference patterns (i.e., to reduce page faults). The pools are:

- 0 Private RUPE Pool
- 1 Common RUPE Pool
- 2 Private SIB Pool
- 3 SSCP FMCB Pool
- 4 NAB Pool
- 5 DVT/EPT Pool
- 6 CDRSC Pool
- 7 ACDEB Pool
- 8 HSQH Pool
- 9 ERTE Pool
- 10 WREEID Pool
- 11 FMCB Extension Pool
- 12 SIB Extension Pool
- 13 RSQE Pool
- 14 UECB/VRPL Pool
- 15 IOBLOCK Pool
- 16 SRTE Pool
- 17 NLDM Trace Pool
- 18 PHIB Pool
- 19 Message Timer Pool

The GETBLK request can be issued for one or more control blocks from these pools. ISTCPCSM is invoked to obtain the storage:

- The anchor for the pool (SPANC) is located by taking the SPANC pointer in the ATCVT and then indexing from that address to the desired SPANC.
- The first PAGTB pointed to by that SPANC is located. If its SPTE entries indicate that there is sufficient storage available to satisfy the request, the page is located and made available.

Otherwise, the next PAGTB is examined.

- If all PAGTBs indicate that the request cannot be satisfied, then VTALLOC is issued to obtain more storage for the pool.

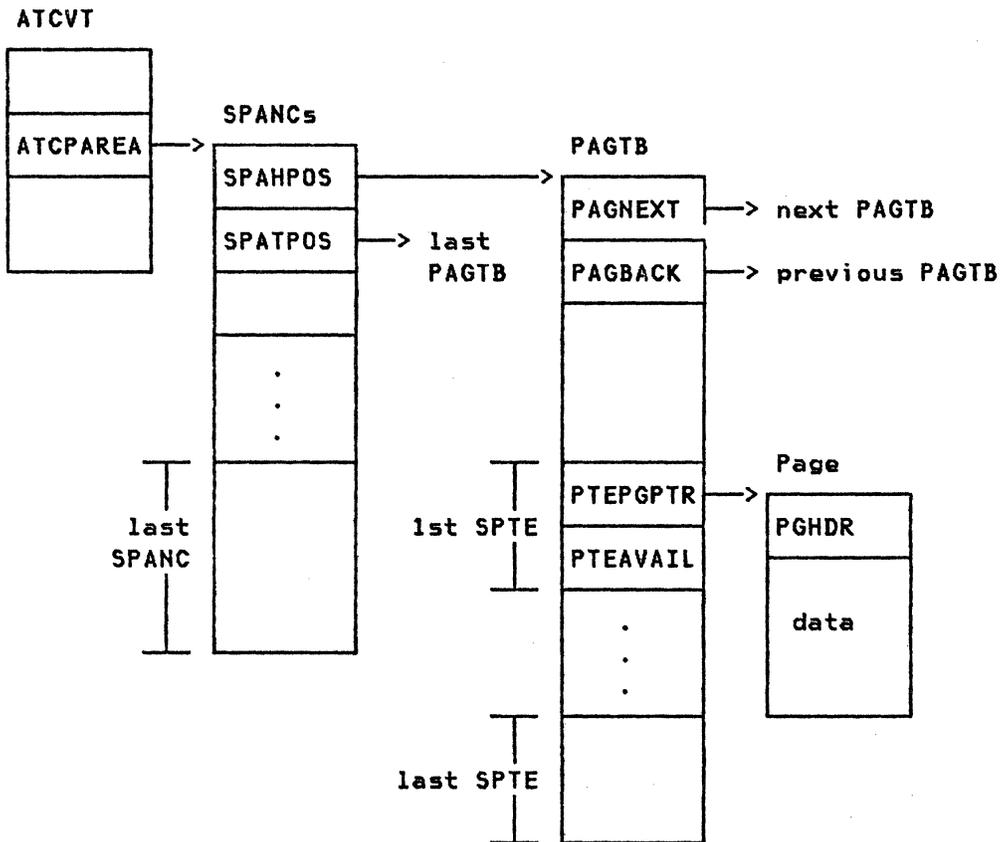


Figure 55. SMS Variable Length Storage Pools

The FREEBLOCK request is used to return storage to these pools.

10.0 VIRTUAL STORAGE IN VTAM

This chapter describes how VTAM allocates and uses virtual storage in a MVS/XA system - a topic which has received considerable attention in very large networks.

Although Virtual Storage Constraint Relief is a key feature of ACF/VTAM vs, it is important to understand the degree to which this is offered and the effects of this relief.

Virtual storage should not be viewed as an isolated entity but as part of the total network storage system. This excerpt from Network Program Products Planning, SC23-0110 provides a valuable perspective on network storage.

The elements of a network can be regarded as a series of related storage spaces...a host processor provides storage for VTAM and application programs...communications controllers contain storage for NCPs and related programs...cluster controllers and programmable terminals contain storage space.

You should be concerned with balancing storage needs among network elements. If increased storage requirements exceed the storage capacity of VTAM or the host processor, VTAM must begin rejecting requests. Similarly, an NCP may exceed its storage capacity if it receives more data...than it can send out....

Hence, not only does virtual storage influence the potential for network growth and expansion; it also influences the speed and success of data transfer between the end points in an application-terminal or application-application session across the network.

10.1 OVERVIEW OF VTAM DATA AREAS

This section provides an overview of the major data areas in ACF/VTAM V3.

The intention is to provide the basic understanding necessary for discussion on VTAM's use of virtual storage.

For the purpose of this overview, ACF/VTAM Data Areas may be broken into two groups:

- **Buffers:** These are defined by the user in ATCSTRxx and allocated at VTAM initialization.
- **Control Blocks:** These are defined and allocated internally by VTAM. ACF/VTAM Control Blocks may be sub-grouped by function:

CONFIG Configuration Control Blocks: These define the components and structure of the domain.

API ACF/VTAM to Application Program Interface Control Blocks: These permit application programs to use ACF/VTAM facilities.

SESSION Session Control Blocks: These are used to control SSCP-SSCP, SSCP-PU, SSCP-LU and LU-LU sessions.

POI Program Operator Control Blocks: These are used for the Program Operator Interface between ACF/VTAM and an application program.

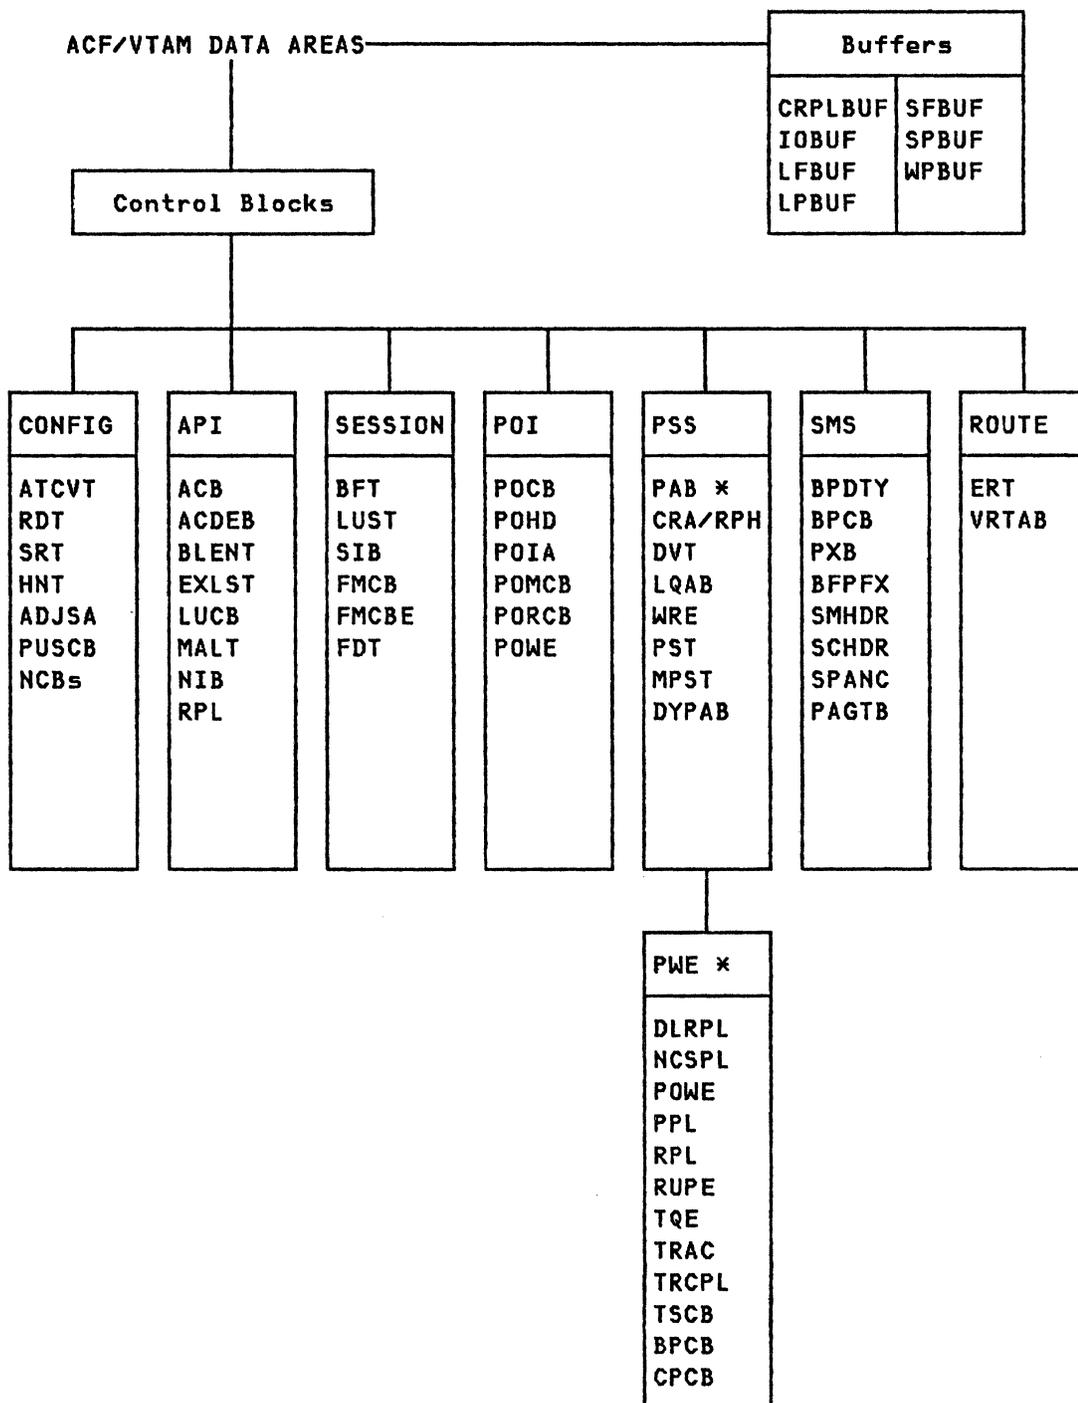
PSS Process Scheduling Control Blocks: These are used in scheduling and dispatching ACF/VTAM activities. These control blocks are used by VTAM's Process Scheduling Services.

There are also **Process Work Elements (PWE)** that represent entities of work for the dispatcher. These are chained off the PAB control block.

SMS Storage Management Control Blocks: These are used to control VTAM's storage allocation and de-allocation.

ROUTE Route Management Control Blocks: These are used to control explicit and virtual routes.

Figure 56 on page 129 below shows major ACF/VTAM Data Areas by these categories.



* See "PAB (Process Anchor Block)" on page 142 for an explanation of the relationships between PABs and PWEs.

Figure 56. ACF/VTAM Data Areas

10.2 VTAM DATA AREA DICTIONARY

This section provides information on major VTAM data areas. Those areas which are internal to VTAM routines or overlay other data areas have been omitted.

The descriptions provided are brief. They may provide an initial understanding of the function and volume of these data areas. However, it is not intended as a means of quantifying virtual storage from 'first principles'.

For more details, see VTAM Diagnosis Reference, SC23-0117 and VTAM Data Areas, SC23-0121.

Each description is preceded by a title (i.e., the data area's acronym and its full name) and the following information:

Category The category of data area (e.g., API, BUFFER, CONFIG, POI, PSS, PWE, ROUTE, SESSION, SMS, OTHER). See Figure 56 on page 129.

Base Size For fixed length areas, the size of the entire data area. For variable length areas, the size of the base portion. (ACF/VTAM V3 specifications.)

Entry Size For variable length areas only, the size of each entry. (ACF/VTAM V3 specifications.)

Old Location The MVS virtual storage area used in the last release of VTAM that did not exploit 31-bit addressing (i.e., ACF/VTAM V2R2).

New Location The MVS virtual storage area used in the release of VTAM that exploits 31-bit addressing in the MVS/XA environment (i.e., ACF/VTAM V3).

Note: If a data area is shown in an extended location (e.g., EPVT SP250), it does not mean that it will be always allocated in the extended location. It may be located below the 16 megabyte virtual storage boundary (see "Allocation of Extended Virtual Storage" on page 112 for more details).

These are presented by acronym, in alphabetical order.

WARNING: The size, location and other characteristics of these areas may change with release or maintenance levels.

DATA AREA(s)	loc	SQA	LPA	227	231	241	LSQ	SWA	229	230	250	251
ACDEB	a					p						
ADJSA	a			f								
APWA	a				p							
ATCVT	b					p						
AVT	b					p						
BFT	a											
BPCB	a				f							
BPDY (LFBUF)	a				f							
CONFT	b				p							
CRA/RPH (LPBUF)	a				p							
CRPLBUF	a				p							
CSCB	a				p							
DLRPL	b					p						
DVT	a								p			
DWA	b			f								
DYPAB	a					p					p	
EPT	a					p						
ERT	b										p	
FDT	a									p		
FMCB	a				p						p	
FMCBE	a					p					p	
HNT	a			f								
IOBUF	b				f							
LMPCB (SPBUF)	a				p							
LQAB	a										p	
LUCB (LFBUF/SFBUF)	a				f							
MLCA	a										p	
MPST	b			f								
NCBs	b			f								
NCSPL	a					p						
PAGTB	a					p	f				p	
PST	a			f								
QAB	a				p							
RDT	a										p	
RUPE	a					p					p	
SIB, SIBX	a										p	
SPANC	a					p						
SRT directory	a					p						
SRT entries	a										p	
UECB	b					p						
VRBLK	a			f								
WRE	a										p	

a - anywhere, b - below, f - fixed, p - pageable

Figure 57. VTAM Data Area - MVS Subpool Cross Reference

ACB (Access Method Control Block)

Category.....API
Base Size.....108 bytes
Location.....Application program storage

The ACB is located in the application program's own storage and defines the interface between VTAM API code in the application program and VTAM API routines (e.g., OPNDST, SEND, RECEIVE, CLSDST). These routines can only be invoked from an application program when the ACB has been initialized with an OPEN ACB request. It remains allocated until the application program issues a CLOSE ACB and frees the storage.

The ACB is required by every VTAM application program. Some single-address space application programs use more than one ACB (e.g., NCCF). TSO requires an ACB for each user address space.

ACDEB (ACF/VTAM Data Extent Block)

Category.....API
Base Size.....154 bytes
Old Location...CSA Pageable SP241
New Location...ECSA Pageable SP241

The ACDEB is created when an OPEN ACB is issued and remains allocated and used until the application program closes its ACB. An ACDEB is also created for the VTAM SSCP.

Basically, it contains information from the ACB and keeps track of the application program's status and session counts. In addition, it points to the application program's LUCB which, in turn, point to the FMCBs for LUs with which it is in session. It is also the queueing point for all outstanding RECEIVE ANY requests issued by the application program.

The ACDEB is the control point for terminating sessions when an application program abends or issues a CLOSE ACB while sessions still exist.

ADJSA (Adjacent Subarea Table)

Category.....CONFIG
Base Size.....8 bytes
Entry Size.....4 bytes
New Location...ECSA Fixed SP227

The ADJSA is a table of host or NCP subareas adjacent to this host. It contains an entry for each adjacent subarea, which points to an entry in this host's HNT (the HNT entry represents the link station for the channel-attached host or NCP).

This was introduced in ACF/VTAM V3. See also HNT.

APWA (PSS Work Area Control Block)

Category.....PSS
Base Size.....496 bytes
Old Location...CSA Pageable SP231
New Location...ECSA Pageable SP231

The APWA is the work area used by the three dispatching queues in a PST (Process Scheduling Table). It is used to maintain information for recovery should a process fail.

Each PST requires two APWAs. There are three PSTs for VTAM and one PST for each application program subtask that also opens an ACB.

ATCVT (ACF/VTAM Communication Vector Table)

Category.....CONFIG
Base Size.....2540 bytes
Old Location...CSA Pageable SP241
New Location...unchanged

The ATCVT is the main ACF/VTAM control block and is pointed to by the address at location X'408' in the VTAM address space. There is one ATCVT for each VTAM.

It is created when VTAM initializes and is addressed by all VTAM components. This block contains the:

- PABs for major/minor node activation/deactivation
- ECBs and QABs for DISPLAY, VARY, MODIFY and HALT
- pointers to the default USS and LOGMODE tables
- pointers to the HNT, CONFT, RDT segments and SRT directory
- pointer to the active NCB queue
- pointer to the MPST queue and ACDEB chain
- pointer to the ERT and VRTAB
- pointer to the BPDY
- addresses of VTAM routines not contained in DVTs

AVT (Address Vector Table)

Category.....OTHER
Base Size.....52 bytes
Old Location...CSA Pageable SP241
New Location...unchanged

The AVT is used in debugging to locate the address space ID (ASID).

BLENT (Buffer List Entry)

Category.....API
Entry Size.....16 bytes
Location.....Application program storage

The number of BLENTs in a system will depend on application program design and activity.

The BLENT is an entry in the buffer list used when transmitting discontinuous data with a single SEND through the VTAM API. Each entry points to a separate piece of data. The list is used to recombine the data.

BPCB (Buffer Pool Control Block)

Category.....SMS
Base Size.....152 bytes
Old Location...CSA Fixed SP231
New Location...ECSA Fixed SP231

There is one BPCB for each of the eleven fixed-length buffer pools (there were twelve pools in ACF/VTAM V2R1).

It serves as the anchor block for each pool to chain free buffers and a queue of waiting buffer requests. See "Building and deleting buffer pools" on page 120 for more information.

BPDTY/BPENT (Buffer Pool Directory/Entry)

Category.....SMS
Base Size.....384 bytes
Entry Size.....16 bytes
Old Location...CSA Fixed SP231
New Location...ECSA Fixed SP231

The BPDTY is the main control block for SMS.

It contains an entry and points to the BPCB for each buffer pool. It also contains the CBIDT (an additional 320 bytes). See "Building and deleting buffer pools" on page 120 for more information.

CONFT (Configuration Table)

Category.....CONFIG
Base Size.....3028 bytes
Old Location...CSA Pageable SP231
New Location...unchanged

The CONFT contains information derived from the VTAM starter list (ATCSTRxx). There is only one in each SSCP.

CPCB (Control Point Control Block)

Category.....PWE
Base Size.....64 bytes
Old Location...CSA SP231, PVT SP229
New Location...ECSA SP231, EPVT SP229

The CPCB is used by Configuration Services when a small work element is required (e.g., to process INOP RUs).

It is also used by other VTAM components as a prefix in NCSPLs, RUPEs, PPLs and DLRPLs to contain codes that identify the process or command requested.

CRA/RPH (Component Recovery Area/Request Parameter Header)

Category.....PSS
Base Size.....1344 bytes
Location.....LPBUF

The CRA is used to indicate which ACF/VTAM locks are currently held, provide an RPH (108 bytes) for each process in progress and maintain component recovery records (CRRs).

At VTAM initialization, three CRAs are allocated. Thereafter, one is created for each application program that opens an ACB. Each process in progress (i.e., actively dispatched by PSS) requires a CRA/RPH:

- every SSCP request,
- every VTAM macro issued by an application program,
- every VTAM operator command, or
- every response required.

CRPLBUF (Copy RPL Buffer)

Category.....BUFFER
Base Size.....116 bytes
Old Location...CSA Pageable SP231
New Location...ECSA Pageable SP231

The CRPLBUF is used to copy Request Parameter Lists (RPLs).

Whenever a VTAM application program issues an RPL-based request (e.g., OPNDST, CLSDST, SEND, RECEIVE), a CRPLBUF is required.

CSCB (Command Scheduler Control Block)

Category.....PSS
Base Size.....156 bytes
Old Location...CSA Pageable SP231
New Location...ECSA Pageable SP231

The CSCB contains VTAM operator commands and is used by command scheduling routines. It identifies the POI program that issued the command and contains routing information.

DVT (Destination Vector Table)

Category.....PSS
Base Size.....24 bytes
Entry Size.....4 bytes
Old Location...PVT SP229
New Location...EPVT SP229

The DVT contains the addresses of routines in the sequence that they are to be executed. Each routine takes an entry in the table.

Each PAB points to a DVT which defines the program execution sequence to accomplish the process represented by the PAB. See PAB for more information.

DWA (Disabled Work Area)

Category.....OTHER
Base Size.....2368 bytes
Old Location...CSA Fixed SP227
New Location...unchanged

The DWA is a work area for disabled code. This area is required for each central processor (CP) in the processor complex where a VTAM operates.

The amount of storage required for DWAs by each VTAM is:

$(\text{No. of CPs} + 1) \times \text{DWA Size}$

DYPAB (Dynamic PAB)

Category.....PSS
Base Size.....32/48 bytes
Old Location...CSA Pageable SP241, PVT SP229
New Location...ECSA Pageable SP241, EPVT SP229

The DYPAB is created to contain a PAB (or extended PAB) which does not reside within a 'major control block'.

See PAB for more information.

EPT (Entry Point Table)

Category.....PSS
Base Size.....32 bytes
Entry Size.....16 bytes
Old Location...CSA Pageable SP241
New Location...ECSA Pageable SP241

The EPT contains pointers to VTAM routines.

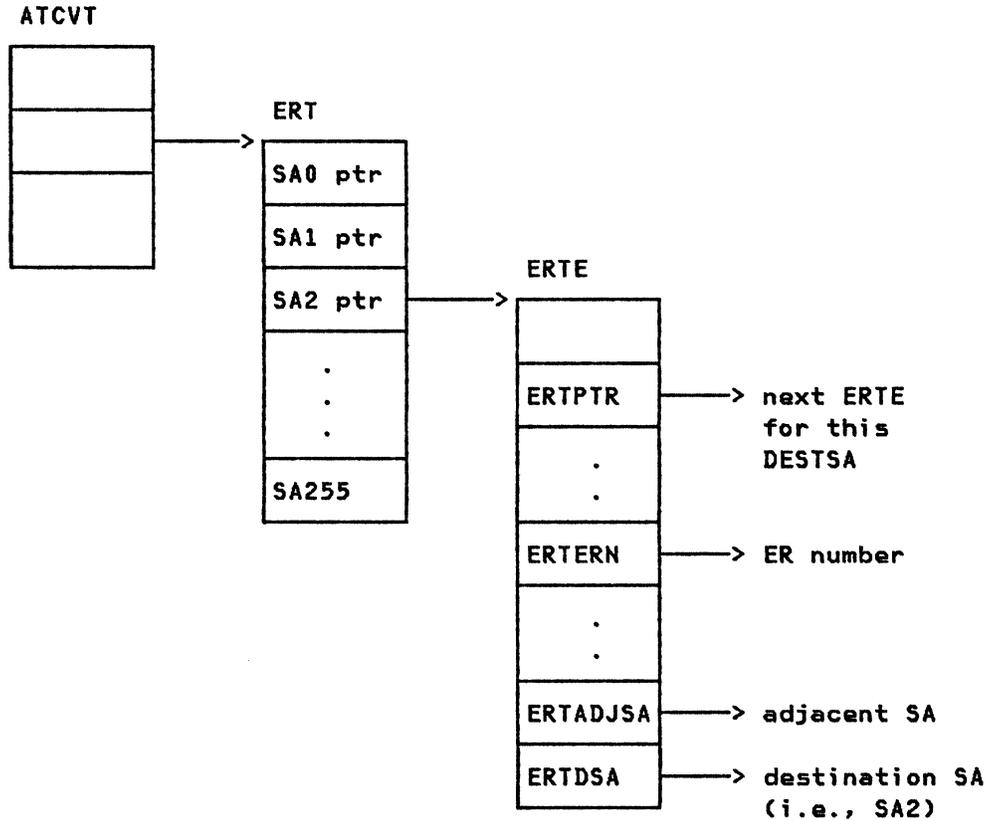


Figure 58. ERT and ERTE Relationships

ERT/ERTE (Explicit Route Table/Entry)

Category.....ROUTE
 Base Size.....see below
 Entry Size.....40 bytes
 Old Location...PVT SP229
 New Location...unchanged

The ERT is created when a path table is activated or when an ER.OP is received. This table is an array of entries for every possible subarea in the network. Its size is ((MAXSUBA+1)*4).

The ERT array points to ERTEs. An ERTE is created for each PATH DESTSA definition.

EXLST (Exit List)

Category.....API
 Base Size.....variable
 Location.....Application program storage

The EXLST is created when a VTAM application program opens its ACB.

It contains the addresses of the program's API exit routines.

FDT (FMCB Directory Table)

Category.....SESSION
Base Size.....10 bytes
Entry Size.....6 bytes
New Location...varies

The FDT was introduced in ACF/VTAM V3 to replace the FLUT. When VTAM needs to locate a session from a network address pair, VTAM uses the HNT to locate the PLU session partner. This results in an LUCB (for the application program or the SSCP).

The address is 'hashed' to derive an index into the FDT (which is pointed to from the LUCB). The index leads to an anchor for a queue of FMCBs, which are then searched sequentially for a matching network address.

The FDT is created for each LUCB allocated. If the APPL EAS value is greater than 30, the FDT is located in EPVT SP230. Otherwise, it is located in the LUCB.

FLUT (FMCB Lookup Table)

Category.....SESSION
Base Size.....10 bytes
Entry Size.....5 bytes
Location.....varies

The FLUT was replaced by the FDT in ACF/VTAM V3.

FMCB/FMCBE (Function Management Control Block and Extension)

Category.....SESSION
Base Size.....160 bytes
Old Location...PVT SP250 (SSCP FMCBs)
New Location...EPVT SP250
Location.....WPBUF (other FMCBs)

The FMCB/FMCBE is created for every session that the SSCP or an application program has. The FMCB is used by the Transmission Subsystem Component and contains the PABs for inbound and outbound data requests. It also contains queue anchors for waiting requests/responses related to the session (i.e., queued RPLs and TSCBs) and status information (e.g., VR pacing status).

The FMCB Extension contains information required by Logical Unit Services (e.g., PLU and SLU network address, name of session partner, RUPE pointers).

See FDT also.

HNT (Host Network Table)

Category.....CONFIG
Base Size.....16 bytes
Entry Size.....17 bytes
New Location...ECSA Fixed SP227

This table was introduced in ACF/VTAM V3 and with the ADJSA replaces the MNT and SNT.

The HNT keeps track of element addresses within the host subarea. An entry is built for each application, channel-attached device and adjacent subarea link station activated in the host subarea. The entry points to the node's associated control block (i.e., LUCB, NCB or RDTE).

The HNT is a chain of 2K blocks. The first block is built when VTAM initializes. Each subsequent node activation results in an entry in the HNT (until the highest element address is assigned).

See also ADJSA.

ICNCB (Intelligent Controller Node Control Block)

Category.....CONFIG
Base Size.....472 bytes
Old Location...CSA Fixed SP227
New Location...unchanged

The ICNCB is created for each activated channel-attached PU (e.g., 37x5, local 3274).

It contains channel programs and PABs for the PU. It is used by the VTAM routines that communicate with the IO Supervisor.

INNCB (Intermediate Network Node Control Block)

Category.....CONFIG
Base Size.....72 bytes
Old Location...CSA Fixed SP227
New Location...unchanged

The INNCB is created for each intermediate network node (INN). It contains PABs, queues for INN requests and information on adjacent nodes. The INNCB may be appended with a 16-byte extension when an adjacent node enters slowdown mode. This points to the NCB for the node in slowdown and a queue of PIUs waiting to be sent.

IOBUF (IO Buffer)

Category.....BUFFER
Base Size.....(UNITSZ+78)
Old Location...CSA Fixed SP231
New Location...unchanged

The IOBUF is used to hold each PIU that enters or leaves VTAM. Hence, these are required for data transmission through all session types and the VTAM Communications Network Management Interface (CNMI).

The length of each IOBUF will be the sum of the installation's UNITSZ specification and a 78-byte header. The number of IOBUFs required will vary with network activity. SMS will expand (provided there is adequate storage in the subpool) and contract the IOBUF pool according to demand.

LDNCB (Local Device Node Control Block)

Category.....CONFIG
Base Size.....472 bytes
Old Location...CSA Fixed SP227
New Location...unchanged

The LDNCB is created for each activated channel-attached non-SNA device (e.g., local 3277 or non-SNA 3278).

It contains the characteristics of, current status of and CCWs for the terminal. It is used to control and schedule I/O to the device.

LFBUF (Large Fixed Buffer)

Category.....BUFFER
Base Size.....120 bytes
Old Location...CSA Fixed SP231
New Location...ECSA Fixed SP231

The LFBUF is used to contain the LUCB for each active VTAM application program with an EAS value that is greater than 30 (see SFBUF below).

SMS will expand and contract the LFBUF pool according to demand.

LPBUF (Large Pageable Buffer)

Category.....BUFFER
Base Size.....1334 bytes
Old Location...CSA Pageable SP231
New Location...ECSA Pageable SP231

The LPBUF is used to contain the CRA/RPH. It is also used for the PST of each task with at least one ACB open. When VTAM terminates, two LPBUFs are used for DYPABs.

LQAB (Locked Queue Anchor Block)

Category.....PSS
Base Size.....16 bytes
Old Location...PVT SP250
New Location...EPVT SP250

LQABs are anchors for waiting SSCP, Logical Unit Services or Storage Management Services requests. Each of these services has a number of LQABs equal to (MAX-SUBA+1). These are used to chain WREs for each subarea.

LUCB (Logical Unit Control Block)

Category.....CONFIG
Base Size.....64 bytes
Location.....SFBUF

The LUCB is allocated for each active VTAM application program (i.e., created on OPEN ACB and deleted on CLOSE ACB). When created, its address is placed in the HNT. It contains pointers to other control blocks associated with the application program (e.g., PST, ACDEB, FDT, FMCB).

The size of the LUCB for an application program depends on the APPL EAS value specification. If the EAS value is 6 or less, the FDT will be contained within the LUCB. Otherwise, the application program's FDT will be located elsewhere and the LUCB will point to it.

MLCA (Main Line Communications Area)

Category.....OTHER
Base Size.....1436 bytes
Old Location...PVT SP250
New Location...EPVT SP250

The MLCA contains a work area for network definition routines. It is created before network activation and deleted when that has completed.

MNT (Major Node Table)

Category.....CONFIG
Entry Size.....12 bytes
Old Location...CSA Fixed SP227

This has been replaced by the HNT and ADJSA in ACF/VTAM V3.

MPST (Memory Process Scheduling Table)

Category.....PSS
Base Size.....96 bytes
Old Location...CSA Fixed SP227
New Location...unchanged

The MPST is created for each address space that OPENS one or more ACBs. All PSTs belonging to that address space are chained off its MPST.

NCSPL (Network Configuration Services Parameter List)

Category.....PWE
Base Size.....372 bytes
Old Location...CSA Pageable SP241
New Location...ECSA Pageable SP241

The NCSPL is created for each VTAM operator command in progress (i.e., DISPLAY, VARY, MODIFY or HALT). It contains the symbolic name of the resource and work area addresses.

NIB (Node Initialization Block)

Category.....API
Base Size.....64 bytes
Location.....Application program storage

The number of NIBs in a system will depend on application program activity.

The NIB describes the characteristics of a session request (i.e., OPNDST or CLSDST) and identifies the LU involved.

PAB (Process Anchor Block)

Category.....PSS
Base Size.....16/32 bytes
Location.....(within the major control block)

The PAB is a dispatching point for Process Scheduling Services (PSS). It represents a sequence of operations required to complete a VTAM process. The PAB points to a DVT which points to the VTAM modules (in sequence) required to complete the process.

When a process is required, a process work element (PWE) is queued off the appropriate PAB. The PAB is then scheduled by queuing to the PST. This identifies the routines required by that work element. The PAB is then scheduled for dispatching.

In some cases, the PAB may contain a 16-byte extension. This extension is created when the system is very busy and is used to queue additional PWEs without interrupting the chain of PWEs in process.

The PAB is always contained within another control block. However, certain PABs reside in a DYPAB (the only function of a DYPAB is to contain a PAB). Control blocks that contain PABs are called 'major control blocks', for example:

ACDEB	ATCVT	CONFT	FMCB	INNCB	LUCB
NCBs	PST	VRBLK			

PWEs queued to PABs are:

BPCB	CPCB	DLRPL	NCSPL	POWE	PPL
RPL	RUPE	TQE	TRAC	TRCPL	TSCB

PAGTB (Page Table)

Category.....SMS
Base Size.....16 bytes
Old Location...CSA Pageable SP241, PVT SP229
New Location...ECSA Pageable SP241, EPVT SP229

The PAGTB points to pages of storage that may be obtained or released by Storage Management GETBLK/FREEBLK routines.

It is contiguous with the SPTE.

POCB (Program Operator Control Block)

Category.....POI
Base Size.....311 bytes
Old Location...PVT SP250
New Location...EPVT SP250

The POCB is created for each VTAM application program that uses the program operator interface (POI).

It serves as an anchor point for POMCBs and PORCBs related to that application program.

POHD (Program Operator Message Header)

Category.....POI
Base Size.....4 bytes
Location.....Application program storage

The POHD is created for each message or command exchanged between ACF/VTAM and the application program over the POI.

It helps the application program keep track of the messages and commands.

POIA (Program Operator Interface Area)

Category.....POI
Base Size.....12 bytes
Old Location...PVT SP250
New Location...EPVT SP250

The POIA is created for the VTAM POI (i.e., only one per SSCP).

It is the anchor block for POCBs.

POMCB (Program Operator Message Control Block)

Category.....POI
Base Size.....154 bytes
Old Location...PVT SP250
New Location...EPVT SP250

The POMCB is created to hold messages destined for application programs using the VTAM POI.

PORCB (Program Operator Reply Control Block)

Category.....POI
Base Size.....154 bytes
Old Location...PVT SP250
New Location...EPVT SP250

The PORCB is created for each reply that is outstanding from an application program using the VTAM POI.

POWE (Program Operator Work Element)

Category.....PWE
Base Size.....319 bytes
Old Location...PVT SP250
New Location...EPVT SP250

The POWE is created for each message scheduled to an application program using the VTAM POI.

The POMCB and PORCB are built from the POWE.

PST (Process Scheduling Table)

Category.....PSS
Base Size.....184 bytes
Old Location...CSA Fixed SP227
New Location...ECSA Fixed SP227

PSTs are created at different times:

- When VTAM initializes, one is created for the SSCP subtask and another for the TNSTAT subtask.
- A PST is created for each application program subtask that OPENS one or more ACBs.

It serves as control point for scheduling all requests related to the subtask.

QAB (Queue Anchor Block)

Category.....OTHER
Base Size.....24 bytes
Old Location...CSA Pageable SP231
New Location...ECSA Pageable SP231

The QAB anchors standard ACF/VTAM queues (e.g., for RDT segments, MODIFY, HALT, VARY, DISPLAY).

RDT (Resource Definition Table)

Category.....CONFIG
Base Size.....variable
Old Location...PVT SP250
New Location...EPVT SP250

Each SSCP has an RDT.

The RDT is a chain of segments. An RDT segment is created for each major node activated in the domain:

- NCP
- Application program
- Local non-SNA
- Switched terminal
- Local SNA
- CDRM
- CDRSC

Each RDT segment will contain entries for each minor node within the major node.

Each entry contains status information and pointers to other relevant control blocks.

RPL (Request Parameter List)

Category.....API
Base Size.....112 bytes
Location.....Application program storage

The number of RPLs in a system will depend on application program design and activity. VTAM application programs may either dedicate an RPL for each session or share a 'pool' of RPLs among its sessions.

The RPL is used as a parameter list and feedback area for VTAM API session and IO requests.

RUPE (Request/Response Unit Processing Element)

Category.....PWE
Base Size.....132 bytes
Old Location...CSA Pageable SP241, PVT SP250
New Location...ECSA Pageable SP241, EPVT SP250

The RUPE is the work element used by VTAM's SSCP, Configuration Services, LU Services, PU Services and Session Services.

The RUPE is accompanied by a work area of variable length.

SFBUF (Small Fixed Buffer)

Category.....BUFFER
Base Size.....64 bytes
Old Location...CSA Fixed SP231
New Location...ECSA Fixed SP231

The SFBUF is used to contain the LUCB.

It is created for each active VTAM application program with an EAS value that is greater than or equal to 30 (see LFBUF above).

SIB (Session Information Block)

Category.....SESSION
Base Size.....228 bytes
Old Location...PVT SP250
New Location...EPVT SP250

The SIB is created for each half-session within a LU-LU session. It is created and deleted with the session.

These keep track of session status, for example:

- ACTIVE: OPNDST or OPNSEC completed
- PENDING ACTIVE: BIND sent/received but response pending.
- QUEUED: Application program's LOGON exit has not been scheduled.

For a cross-network session, the SIB is extended by 100 bytes.

SNT (Specific Node Table)

Category.....CONFIG
Base Size.....8 bytes
Entry Size.....8 bytes
Old Location...CSA Fixed SP227

The SNT was replaced by the HNT and ADJSA in ACF/VTAM V3.

SPANC (Storage Pool Anchor Block)

Category.....SMS
Entry Size.....28 bytes
Old Location...CSA Pageable SP241
New Location...ECSA Pageable SP241

The SPANC anchors a chain of page tables (PAGTB). See "Getting and freeing variable length storage" on page 124 for more information.

SPBUF (Small Pageable Buffer)

Category.....BUFFER
Base Size.....96 bytes
Old Location...CSA Pageable SP231
New Location...ECSA Pageable SP231

The SPBUF is allocated for each VTAM API SEND request that uses the the large message processing option (LMPE0).

SRT (Symbol Resolution Table)

Category.....CONFIG
Base Size.....8192 bytes
Entry Size.....13 or 20 bytes
Old Location...CSA Pageable SP241 (directory)
New Location...ECSA Pageable SP241
Old Location...PVT SP250 (entries)
New Location...EPVT SP250

The SRT contains a directory and an entry for each RDT entry. The SRT entry is extended by 7 bytes for resources in other SNI connected networks.

Each network has its own SRT. If the network has a cross-network session, then an SRT will be created for the other network.

The SRT is used to convert (via a hashing algorithm):

- the symbolic name of a network node, or
- the name of a table

into the address of the control block for that node (e.g., RDT) or table.

SRT entries are created for NCPs, LINES, PUs, LUs, CDRSCs, CDRMs, APPLs, shadow resources, switched resources, USS tables, LOGMODE tables, etc.

TSCB (Transmission Subsystem Control Block)

Category.....PWE
Base Size.....68 bytes
Location.....IOBUF

The TSCB is used by the TSC (Transmission Subsystem Control) to process inbound and outbound data.

It is found in IOBUFs.

VRBLK (Virtual Route Block)

Category.....ROUTE
Base Size.....152 bytes
Old Location...CSA Fixed SP227
New Location...ECSA Fixed SP227

The VRBLK is created for each virtual route (VR) to a destination subarea (when a PATH table is activated). It keeps the status of each transmission priority on the VR.

WPBUF (Working Pageable Buffer)

Category.....BUFFER
Base Size.....152 bytes
Old Location...CSA Pageable SP231
New Location...ECSA Pageable SP231

The WPBUF is used to contain the FMCB. Hence, one WPBUF is required for each LU-LU session involving a VTAM application program.

WRE (Waiting Request Element)

Category.....PSS
Base Size.....20 bytes
Old Location...PVT SP250
New Location...EPVT SP250

The WRE is created for a VTAM process that is waiting on the completion of some event (e.g., one PAB may be suspended while another is allowed to run).

XCNCB (Cross-Channel Node Control Block)

Category.....CONFIG
Base Size.....624+ bytes
Old Location...CSA Fixed SP227
New Location...unchanged

The XCNCB is created for each activated channel link to an adjacent host. It is similar to the ICNCB.

11.0 VTAM VIRTUAL STORAGE ESTIMATION

The purpose of this section is:

- to provide a structure for estimating the VTAM virtual storage requirement;
- to estimate the virtual storage required below the 16 megabyte virtual storage boundary (in MVS/XA environments); and
- to help in further understanding the use and location of virtual storage.

These tables will require more time to use than those provided in Appendix A of Network Program Products Planning, SC23-0110.

11.1 INTRODUCTION

The tables below are for ACF/VTAM V3 in the MVS environments.

The following columns have these meanings:

Storage Area The location in virtual storage which contains the corresponding data areas. This is usually the subpool number.

Amount The quantity of virtual storage required (in kilobytes).

Below 16M ? If **yes** appears in this column, the data area(s) will always be obtained from below the 16M virtual storage boundary. Otherwise, the storage may be obtained on either side of the boundary. (See "Allocation of Extended Virtual Storage" on page 112.)

Amount below If the data area must always be obtained from below the 16M virtual storage boundary, this column quantifies the storage below.

To use these tables, obtain the installation's:

VTAMLST definitions (including NCP and other major nodes), and

NLDM (if installed) definitions in the AAUPRMLP member of the NCCF library.

These definitions have to be analyzed and summarized to provide input to the tables below. If VTAM Virtual Storage Estimation will be a continuing process at the installation (and this is recommended), consider developing a program to extract these numbers from the VTAMLST and AAUPRMLP definitions. This program could also perform the required calculations.

When completed, the tables below will provide only a rough approximation of VTAM's virtual storage requirements. VTAM's actual usage should be measured and the results used to validate and/or calibrate the figures given below. See "Monitoring Virtual Storage" on page 159 for information on measuring.

11.2 VTAM VIRTUAL STORAGE ESTIMATION TABLES

TABLE A: CONSTANT/VARIABLE COMMON STORAGE		Storage Area	Amount	Below 16M ?	Amount below
A	VTAM AND TSO IN SYSTEM, NOT STARTED				
	1. NUCLEUS modules 2. LPA modules	NUCLEUS PLPA	3K 442K	yes yes	3K 104K
B	VTAM STARTED (NO MAJOR NODES ACTIVE)				
	1. PSS tables	SP227	12K	yes	12K
	2. SMS tables, etc.	SP231	12K	yes	4K
	3. IOBUF allocation	IOBUF	K	yes	K
	4. Other ATCSTRxx bufpool allocations	SP231	K		
	5. SSCP tables and SMS storage	SP241	20K	yes	4K
6. ISTRACON RACEAS value x 6	SP241	K			
	BA. IOBUF (B3)	IOBUF	K		K
	BB. OTHER CSA (B1+B2+B4+B5+B6)	CSA	K		K
C	MAJOR NODES ACTIVATED				
	1. ((No. local NCPs - z) x 656) + 12K	SP227	K	yes	K
	2. Other local devices (Table C: ZCB)	SP227	K	yes	K
	3. Other local devices (Table C: ZCA)	IOBUF	K	yes	K
	4. Sum (MAXBFRU x (UNITSZ+78)) for all local NCPs	IOBUF	K	yes	K
	5. If NLDM PIU trace activated, (no. buffers x PIU trace bufsize)	SP231	K		
6. EPT	SP241	8K			
	CA. IOBUF (C3+C4)	IOBUF	K		K
	CB. OTHER CSA (C1+C2+C5+C6)	CSA	K		K
D	NETWORK ACTIVATION (OF MINOR NODES)	SP231	8K		

(continued on next page)

TABLE A: CONSTANT/VARIABLE COMMON STORAGE (continuation)		Storage Area	Amount	Below 16M ?	Amount below
E	LU-LU SESSION ESTABLISHMENT				
	1. Constant fixed storage				
	a. SPBUFs	SP231	4K		
	b. IOBUFs	IOBUF	4K	yes	4K
	2. TSO and like sessions				
	a. No. users x 296	SP227	K	yes	K
	b. (No. TSO users x 1932) + (No. other users x 616)	SP231	K		
	c. No. users x (UNITSZ+78)	IOBUF	K	yes	K
	d. No. TSO users x 48	SP239	K		
	e. No. users x 694	SP241	K		
3. Other sessions					
a. 12K+(No. of LUs OPNDSTed x 160)	SP231	K			
b. No. of LUs OPNDSTed x 84	SP241	K			
EA. IOBUF (E1b+E2c)	IOBUF	K		K	
EB. OTHER CSA(E1a+E2a+E2b+E2d+E2e+E3)	CSA	K		K	
ZAA. NUCLEUS REQUIRED (A1)	NUCLEUS	3K		3K	
ZAB. PLPA REQUIRED (A2)	PLPA	442K		104K	
ZAC. IOBUF REQUIRED (CA+EA)	IOBUF	K		K	
ZAD. OTHER CSA REQUIRED (BB+CB+D+EB)	CSA	K		K	

NOTES (Table A)

- A This is not VTAM's total module requirement. See Table D.
- B1 Includes DWA, MPST, PST, etc.
- B2 Includes BPDY, BPCB, SMS pools, CONFT, APWA (for PSS), QABs, etc.
- B3 The default IOBUF specification results in an 8K allocation.
All IOBUFs are located below the 16 megabyte boundary.
- B4 The default specifications for the other buffer pools result in a 32K allocation.
- B5 Includes ATCVT, SRT and SPANC, packed CB pools, etc.
- B6 Required for FDT entries for SSCP-PU, SSCP-LU sessions. Use the default value of 3000. If the ISTRACON RACEAS constant has been zapped and changed (e.g., to improve performance as the network has more than 3000 of these sessions), use that value.
- C1 The constant 12K allows for 2 local NCPs (about 1K is below).
If there are two or more local NCPs, then 'z=2'. Else, 'z=0'.
Includes VRBLK, VRINDEX, TGCB.
656 = ADJSA ent(4)+SSCP-PU FMCB(160)+ICNCB(472)+HNT ent(17)
(Out of the 656 bytes, 472 bytes will be below.)
- C2 Obtain this value from TABLE C below.
- C3 Obtain this value from TABLE C below.
- C4 All IOBUFs are located below the 16 megabyte boundary.
- C5 NLDM PIU trace is invoked by the NLDM START TRACE command.
- D For CRA/RPHs (which are in LPBUFs).
- E1b All IOBUFs are located below the 16 megabyte boundary.
- E2 Include any other VTAM applications that open an ACB for each

- user (e.g., NCCF, TAF, etc.).
- E2a For MPST and PST. All storage is below.
- E2b For 1xSFBUF(LUCB), 2xWPBUF(FMCB), 1xCRPLBUF and TSO VTIOC areas.
- E2c All IOBUFs are located below the 16 megabyte boundary.
- E2d For TSO SRBs.
- E2e For ACDEB and FMCBE.
- E3a Constant requirement - CRPLBUF. Variable - FMCB(WPBUF).
- E3b For FMCBE.
- ZAC This is the constant IOBUF requirement. If ZAC is greater than BA, then IOBUF pool expansion will occur (therefore, consider increasing the initial IOBUF allocation).

TABLE B: DYNAMIC COMMON STORAGE		Storage Area	Amount	Below 16M ?	Amount below
A	NETWORK ACTIVATION (OF MINOR NODES)				
	1. (No. of LUs) x (UNITSZ+78)	IOBUF	K	yes	Y
	2. For NLDM SAW, (SAW bufsize/UNITSZ) x (UNITSZ+78)	IOBUF	K	yes	K
	3. If NLDM PIU trace is activated, (PIU bufsize/UNITSZ) x (UNITSZ+78)	IOBUF	K	yes	K
AA. IOBUF (A1+A2+A3)		IOBUF	K		K
B	LU-LU SESSION ESTABLISHMENT				
	1. (No. concurrent OPNDSTs x 348) + 512	SP241	K		
	2. No. concurrent OPNDSTs x 136	SP231	K		
	3. For NLDM SAW, (SAW bufsize/UNITSZ) x (UNITSZ+78)	IOBUF	K	yes	K
	4. If NLDM PIU trace is activated, (PIU bufsize/UNITSZ) x (UNITSZ+78)	IOBUF	K	yes	K
BA. IOBUF (B3+B4)		IOBUF	K		K
BB. OTHER CSA (B1+B2)		CSA	K		K
ZBA. IOBUF REQUIRED (see notes)		IOBUF	K		K
ZBB. OTHER CSA REQUIRED (BB)		CSA	K		K
ZBC. DYNAMIC CSA REQUIRED (see notes)		CSA	K		K

NOTES (Table B)

- A1 Use the maximum number of LUs activated concurrently.
All IOBUFs are located below the 16 megabyte boundary.
- A2 NLDM Session Awareness (SAW) is invoked for all resources when NLDM is initialized. All IOBUFs are located below.
- A3 NLDM PIU trace is invoked by the NLDM START TRACE command.
All IOBUFs are located below the 16 megabyte boundary.
- B Use the ITLIM value instead of 'No. concurrent OPNDSTs', if this is the number of OPNDSTs.
- B1 $348 = 132(\text{RUPE}) + 216(\text{RUPEWA})$. 512 for PAGTB.
- B2 For CRPLBUF (116 rounded to multiple of 8 and 16 bytes added).
- B3 All IOBUFs are located below the 16 megabyte boundary.
- B4 All IOBUFs are located below the 16 megabyte boundary.
- ZBA If Network Activation and Session Establishment occur concurrently, then $(ZBA = AA + BA)$. Else, ZBA is the larger of either AA or BA.
- ZBC If Network Activation and Session Establishment occur concurrently, then $(ZBC = ZBA + ZBB)$. Else, ZBC is the larger of either ZBA or ZBB.

TABLE C: COMMON STORAGE FOR CHANNEL-ATTACHED (LOCAL) DEVICES		Storage Area	Amount	Below 16M ?	Amount below
A	MAJOR NODES ACTIVATED				
	1. Local non-SNA devices				
	a. (No. clusters x 96) + (No. TERMINALS x 688)	SP227	K	yes	K
	b. (No. TERMINALS) x (UNITSZ+78)	IOBUF	K	yes	K
	2. Local SNA devices (non-NCP)				
	a. ((No. CTC PUs) x 808) + ((No. LOCAL PUs) x 656)	SP227	K	yes	K
b. Sum (MAXBFRU x (UNITSZ+78))	IOBUF	K	yes	K	
ZCA.	IOBUF REQUIRED (A1b+A2b)	IOBUF	K		K
ZCB.	SP227 REQUIRED (A1a+A2a)	SP227	K		K

NOTES (Table C)

- A1a For clusters, count the number of LBUILD major nodes.
96=Non-SNA 3270 FMCBE ext(16)+LUST(80)
688=ADJSA ent(4)+SSCP-PU FMCB(160)+NCB(472)+HNT ent(17)+LUST(32)
(Out of the 688 bytes, 472 bytes will be below.)
- A1b VTAM allocates a number of IOBUFs equal to the maximum of the number required for the last three reads plus one.
All IOBUFs are located below the 16 megabyte boundary.
- A2a For CTC PUs, count the number of PUs in VBUILD TYPE=CA nodes.
For LOCAL PUs, count the number of PUs in VBUILD TYPE=LOCAL nodes.
808=ADJSA ent(4)+SSCP-PU FMCB(160)+NCB(624)+HNT ent(17)
(Out of the 808 bytes, 624 bytes will be below.)
656=ADJSA ent(4)+SSCP-PU FMCB(160)+NCB(472)+HNT ent(17)
(Out of the 656 bytes, 472 bytes will be below.)
- A2b For CTC PUs, use ((No. CTC PUs) x MAXBFRU x (UNITSZ+78) x 2).
For LOCAL PUs, use ((No. LOCAL PUs) x MAXBFRU x (UNITSZ+78)).
Sum these two products for the table.
All IOBUFs are located below the 16 megabyte boundary.
- ZCA Transfer results to Table A.
- ZCB Transfer results to Table A.

TABLE D: CONSTANT/VARIABLE PRIVATE STORAGE		Storage Area	Amount	Below 16M ?	Amount below
A	VTAM AND TSO IN SYSTEM, NOT STARTED		0K		
B	VTAM STARTED (NO MAJOR NODES ACTIVE)				
	1. MVS VSM CBs, tuning statistics	SP229	80K		
	2. LSQA	SP233	28K	yes	28K
	3. SWA	SP235	56K	yes	56K
	4. Load modules	SP251	1276K	yes	28K
	5. System Region	SP252	24K	yes	24K
	BA. PRIVATE HIGH (B1)	PVTH	K		K
	BB. PRIVATE LOW (B2+B3+B4+B5)	PVTL	K		K
C	MAJOR NODES ACTIVATED				
	1. (Sum of EAS values on all APPL definitions) x 6 (for FDT entr.)	SP230	K		
	2. PAGTB, SPANCs	SP250	4K		
	3. Major Nodes (for RDT segm.)				
	(#APPL majnodes x 140)				
	(#CA majnodes x 140)				
	(#LOCAL majnodes x 148)				
	(#LBUILD majnodes x 140)				
	(#NCP majnodes x 812)				
	(#SWNET majnodes x 160)				
	(#CDRM majnodes x 140)				
	(#CDRSC majnodes x 140)	SP250	K		
	4. Minor node def's (for RDT entr.)				
	(#APPL x 208)				
	(#CDRM x 268)				
	(#CDRSC x 212)				
	(#CTC links x 544)				
	(#local PU x 288)				
	(#local LUs x 176)				
	(#local TERMINALS x 316)				
	(#GROUP x 100)				
	(#LINE x 140)				
	((#remote PU + PUDRPOOL) x 148)				
	((#remote LU + LUDRPOOL) x 176)				
	(#remote CLUSTERS x 140)				
(#remote TERMINALS x 176)					
(#PU.T4 link stations x 164)	SP250	K			
5. (#DESTSAs on PATH x 40) + 1K	SP250	K	yes	K	
	CA. PRIVATE HIGH (B1)	PVTH	K		K
	CB. PRIVATE LOW (B1+B2+B3+B4+B5)	PVTL	K		K

NOTES (Table D)

(continued on next page)

C FMCB/FMCBEs counted later (in D)

C5 ERT/ERTEs (all stor. below). Number of DESTSAs on PATH statements

TABLE D: CONSTANT/VARIABLE PRIVATE STORAGE (continuation)		Storage Area	Amount	Below 16M ?	Amount below
D	NETWORK ACTIVATION (OF MINOR NODES)				
	1. Minor node definitions in majnodes (#LINE x 48) + (#local/remote PUs x 284) + (#local/remote LUs x 268) + (#local/remote CLUSTERS x 284) + (#local/remote TERMINALS x 268)	SP250	K		
	DA. PRIVATE LOW	PVTL	K		
E	LU-LU SESSION ESTABLISHMENT				
	1. Same-domain session (No. PLU-SLU sessions) x 276	SP250	K		
	2. Cross-domain session				
	a. No. XPLU-SLU sessions x 340	SP250	K		
	b. No. PLU-XSLU sessions x 552	SP250	K		
	3. Cross-network session				
a. No. XPLU-SLU sessions x 540	SP250	K			
b. No. PLU-XSLU sessions x 752	SP250	K			
4. Cross-network intermediate session (No. XPLU-XSLU sessions) x 1068	SP250	K			
5. For NLDM SAW, (no. SAW buffers) x (SAW bufsize)	SP250	K			
	EA. PRIVATE LOW (E1+E2+E3+E4+E5)	PVTL	K		
	ZDA. PRIVATE HIGH REQUIRED (BA+CA)	PVTH	K		K
	ZDB. PRIVATE LOW REQUIRED (BB+CB+DA+EA)	PVTL	K		K

NOTES (Table D * continuation)

- D1 48 = PUSCB(16)+PLNCB(16)+PLSCB(16)
284 = FMCB(160)+FMCBE(84)+SRTE(24)+PLNCB(16)
268 = FMCB(160)+FMCBE(84)+SRTE(24) (SSCP-PU, SSCP-LU FMCBs)
- E In the context of this table:
PLU = application program within this domain
SLU = terminal (for which an LU is defined) within this domain
XPLU = application program in another domain (or network)
XSLU = terminal (LU defined) in another domain or network
- E1 276 = SIB(228)+SRTEs(48)
E2a 340 = SIB(228)+SRTEs(96)+RSQE(16)
E2b 552 = SIB(228)+SRTEs(96)+RSQE(16)+RCDRS(212)
E3a 540 = SIB(228)+SRTEs(192)+RSQE(16)+SIBX(104)
E3b 752 = SIB(228)+SRTEs(192)+RSQE(16)+RCDRS(212)+SIBX(104)
E4 1068 = SIB(228)+SRTEs(192)+RSQE(16)+RCDRSs(424)+SIBXs(208)
Only applies to hosts that act as Gateway nodes for an LU-LU session where the PLU and SLU reside in another network.

TABLE E: DYNAMIC PRIVATE STORAGE		Storage Area	Amount	Below 16M ?	Amount below
A	MAJOR NODES ACTIVATED 1. (No. LINES x 140) + (No. PUs/CLUSTERS x 148) + (No. LUs/TERMINALS x 176) + 32K	SP250	K		
B	NETWORK ACTIVATION (OF MINOR NODES) 1. RUPE Work Areas (FSM process)	SP250	8K		
C	LU-LU SESSION ESTABLISHMENT 1. Same-domain session a. (No. concurr. LOGONs-ITLIM)x372 b. (No. concurr. OPNDSTs or ITLIM) x 4096 c. (No. queued OPNDSTs/SIMLOGONs) x 2048 2. Cross-domain session (No. XPLU-SLU LOGONs x 2723) + (No. PLU-XSLU LOGONs x 5968) 3. Cross-network session (No. XPLU-SLU LOGONs x 4021) + (No. PLU-XSLU LOGONs x 7924) 4. Cross-network intermediate session (No. XPLU-XSLU sessions) x 4021	SP250 SP250 SP250 SP250 SP250 SP250	K K K K K K		
	EA. PRIVATE LOW (see notes)	PVTL	K		K
	ZEA. PRIVATE LOW REQUIRED (see notes)	PVTL	K		K

NOTES (Table E)

- A Do this calculation for the largest major node only.
- A1 32K for RDT workareas, skeleton DVTs and RUPEs.
- C In the context of this table:
 - PLU = application program within this domain
 - SLU = terminal (for which an LU is defined) within this domain
 - XPLU = application program in another domain (or network)
 - XSLU = terminal (LU defined) in another domain or network
 Most of these figures are averages taken from observation.
- C1a $372 = \text{RUPE}(132) + \text{RUPEWA}(240)$
- C4 Only applies to hosts that act as Gateway nodes for a LU-LU session where the PLU and SLU reside in another network.
- EA If all LU-LU sessions are established concurrently (ie., C1, C2, C3 and C4 are concurrent), then $EA = (C1a + C1b + C1c + C2 + C3 + C4)$. Else,
 - sum the storage requirement for each combination of events, and
 - set EA as the largest value of the combinations.
- ZEA If Major Node Activation (A), Network Activation (B) and LU-LU Session Establishment (C) occur concurrently, then $ZEA = (A + B + EA)$. Otherwise, sum the storage requirement for each combination of events, and set ZEA as the largest value of the combinations.

TABLE Z: SUMMATION		Amount	Amount below
A	NUCLEUS	K	K
B	PLPA	K	K
C	IOBUF		
	1. Constant (ZAC)	K	K
	2. Dynamic (ZBA)	K	K
	Peak (1+2)	K	K
D	OTHER CSA		
	1. Constant (ZAD)	K	K
	2. Dynamic (ZBC)	K	K
	Peak (1+2)	K	K
E	PRIVATE HIGH		
	1. Constant (ZDA)	K	K
F	PRIVATE LOW		
	1. Constant (ZDB)	K	K
	2. Dynamic (ZEA)	K	K
	Peak (1+2)	K	K
G	CSA CONSTANT TOTAL (ZAC+ZAD)	K	K
H	CSA DYNAMIC TOTAL (ZBA+ZBC)	K	K
I	CSA PEAK TOTAL (G+H)	K	K
J	PRIVATE CONSTANT TOTAL (ZDA+ZDB)	K	K
K	PRIVATE PEAK (J+ZEA)	K	K

12.1 WHY MONITOR VIRTUAL STORAGE?

Virtual storage requirements vary significantly from one installation to another. Models to estimate virtual storage have to be generalized for the sake of simplicity. For example, the model presented in "VTAM Virtual Storage Estimation" on page 149 makes a number of assumptions that may not apply to a given installation. Therefore, it is important to measure actual usage and validate the results from the model.

Monitoring virtual storage will help in identifying and analyzing problems that may occur in the network. It will also help in the recognition of trends that may forewarn of potential problems or may be used in medium or long-term planning.

12.2 AN APPROACH

Virtual storage monitoring should be a continuing activity (i.e., measurements summarized and analyzed on a daily, weekly and monthly basis). Obviously, personnel should be allocated (this need not be a full-time responsibility) within the organization to perform these tasks.

In MVS/XA systems with RMF V3.2 installed, the allocation and use of virtual storage can be measured and reported on. The RMF measurement data is written to the SMF datasets. This data can be processed by the RMF Post Processor programs or may be processed by other programs to provide usage information in other formats (for an example, see Figure 63 on page 164).

The standard reports produced by the RMF Post Processor are shown below:

- **Common Storage Summary:** This provides a summary of all common areas. See Figure 59 on page 160.
- **Common Storage Detail:** This provides a breakdown of common area usage by subpool and protection key. See Figure 60 on page 161.
- **Private Area Summary:** This provides a summary of private area usage. See Figure 61 on page 162.
- **Private Area Detail:** This provides a breakdown of private area usage by subpool. See Figure 62 on page 163.

For more information on the use of RMF, see RMF 3.2 Reference and User Guide, LC28-1138.

V I R T U A L S T O R A G E A C T I V I T Y

OS/VS2 SYSTEM ID CMC1 DATE 09/02/84 INTERVAL 10.00.010
 RELEASE 03.8 RPT VERSION 32 TIME 11.00.00 CYCLE 1.000 SECONDS

COMMON STORAGE SUMMARY

NUMBER OF SAMPLES 60

STATIC STORAGE MAPALLOCATED CSA/SQA										
AREA	ADDRESS	SIZE	BELOW 16M			EXTENDED(ABOVE 16M)				
			MAX	AVG	MIN	MAX	AVG			
EPVT	1F00000	2017MMIN								
ECSA	1AB7000	4388K	SQA 304K..	304K..	304K	448K..	448K..	448K		
EMLPA	0	0K	CSA 1672K..	1672K..	1672K	48K..	48K..	48K		
EFLPA	0	0K								
EPLPA	1980000	1244K	ALLOCATED CSA BY KEY							
ESQA	1140000	8448K	0	408K..	408K..	408K	28K..	28K..	28K	
ENUC	1000000	1278K	1	160K..	160K..	160K	4K..	4K..	4K	
-----16 MEG BOUNDARY-----			2	0K..	0K..	0K	0K..	0K	0K	
NUCLEUS	FBA000	278K	3	0K..	0K	0K	0K..	0K	0K	
SQA	F6A000	320K	4	0K..	0K	0K	0K..	0K	0K	
PLPA	C71000	3044K	5	16K..	16K..	16K	16K..	16K..	16K	
FLPA	C6F000	8K	6	1084K..	1084K..	1084K	0K..	0K	0K	
MLPA	0	0K	7	0K..	0K	0K	0K..	0K	0K	
CSA	A00000	2492K	8-F	4K..	4K..	4K	0K..	0K	0K	
PRIVATE	1000	10.0M								
PSA	0	4K	SQA	EXPANSION INTO CSA						
				0K..	0K	0K	0K..	0K	0K	

PLPA INTERMODULE SPACE - 7K IN PLPA AND 8K IN EPLPA
 PLPA SPACE REDUNDANT WITH MLPA/FLPA - 0K IN PLPA AND 0K IN EPLPA

	BELOW 16M			ABOVE 16M		
	MIN	MAX	AVG	MIN	MAX	AVG
CSA						
1FREE PAGES(BYTES)	820K..	820K..	820K	4340K..	4340K..	4340K
LARGEST FREE BLOCK	820K..	820K..	820K	4340K..	4340K..	4340K
ALLOCATED AREA SIZE	1672K..	1672K..	1672K	48K..	48K..	48K
SQA						
FREE PAGES(BYTES)	16K..	16K..	16K	8000K..	8000K..	8000K
LARGEST FREE BLOCK	16K..	16K..	16K	7996K..	7996K..	7996K
ALLOCATED AREA SIZE	320K..	320K..	320K	8448K..	8448K..	8448K

MAXIMUM POSSIBLE USER REGION - 10100K BELOW AND 2008M ABOVE

Figure 59. RMF Virtual Storage Activity Report.....Page 1.

V I R T U A L S T O R A G E A C T I V I T Y

OS/VS2 SYSTEM ID CMC1 DATE 09/02/84 INTERVAL 10.00.010
 RELEASE 03.8 RPT VERSION 32 TIME 11.00.00 CYCLE 1.000 SECONDS

COMMON STORAGE DETAIL

ALLOCATED CSA BY SUBPOOL BY KEY (BELOW 16 MEG)

SP227 SP228 SP231 SP241

	MINIMUM				SUBPOOL	MIN	MAX	AVG
0	4K..	268K..	136K..		226	52K..	52K..	52K
1	4K..	76K..	80K..		239	112K..	112K..	112K
2					245	140K..	140K..	140K
3								
4								
5								
6	108K..	4K..	916K..	56K..				
7								
8-F				4K..				
ALL	108K..	12K..	1260K..	292K..				
	MAXIMUM							
0	4K..	268K..	136K..					
1	4K..	76K..	80K..					
2								
3								
4								
5				16K..				
6	108K..	4K..	916K..	56K..				
7								
8-F				4K..				
ALL	108K..	12K..	1260K..	292K..				
	AVERAGE							
0	4K	268K	136K					
1	4K	76K	80K					
2								
3								
4								
5				16K				
6	108K	4K	916K	56K				
7								
8-F				4K				
ALL	108K	12K	1260K	292K				

Figure 60. RMF Virtual Storage Activity Report.....Page 2.

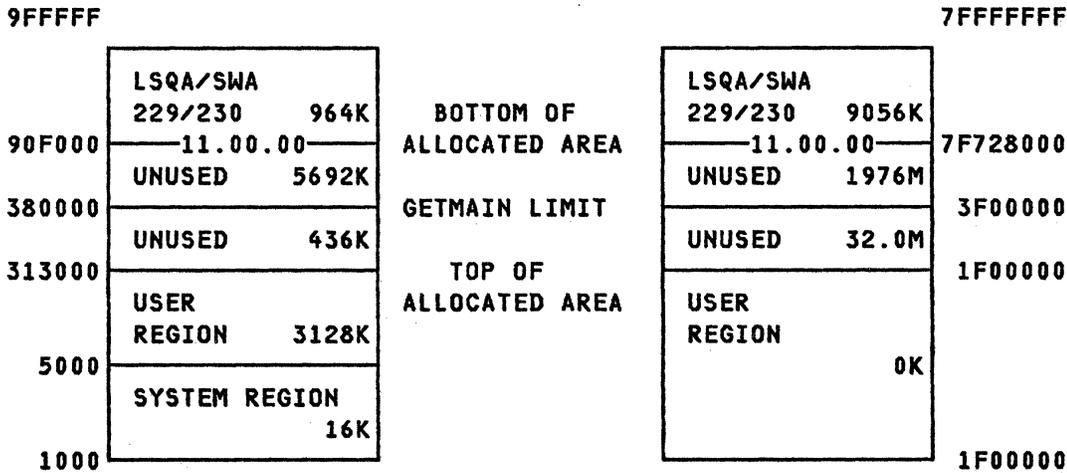
V I R T U A L S T O R A G E A C T I V I T Y

OS/VS2 SYSTEM ID CMC1 DATE 09/02/84 INTERVAL 10.00.010
 RELEASE 03.8 RPT VERSION 32 TIME 11.00.00 CYCLE 1.000 SECONDS

PRIVATE AREA SUMMARY

JOB NAME - NET REGION REQUESTED 3500K
 STEP NAME - NET REGION ASSIGNED(BELOW 16M) 3500K
 PROGRAM NAME - ISTINM01 REGION ASSIGNED(ABOVE 16M) 32.0M
 NUMBER OF SAMPLES - 60

PRIVATE STORAGE MAP



	BELOW 16M			ABOVE 16M		
	MIN	MAX	AVG	MIN	MAX	AVG
LSQA/SWA/229/230						
FREE PAGES(BYTES)	5696K..	5696K..	5696K..	1976M..	1976M..	1976M
LARGEST FREE BLOCK	5692K..	5692K..	5692K..	1976M..	1976M..	1976M
PAGES ALLOCATED (IN BYTES)	964K..	964K..	964K..	9056K..	9056K..	9056K
USER REGION						
FREE PAGES(BYTES)	792K..	792K..	792K..	32.0M..	32.0M..	32.0M
LARGEST FREE BLOCK IN GETMAIN LIMIT	436K..	436K..	436K..	32.0M..	32.0M..	32.0M
PAGES ALLOCATED (IN BYTES)	3128K..	3128K..	3128K..	0K..	0K	0K

Figure 61. RMF Virtual Storage Activity Report.....Page 3.

V I R T U A L S T O R A G E A C T I V I T Y

OS/VS2 SYSTEM ID CMC1 DATE 09/02/84 INTERVAL 10.00.010
 RELEASE 03.8 RPT VERSION 32 TIME 11.00.00 CYCLE 1.000 SECONDS

PRIVATE AREA DETAIL

JOB NAME - NET

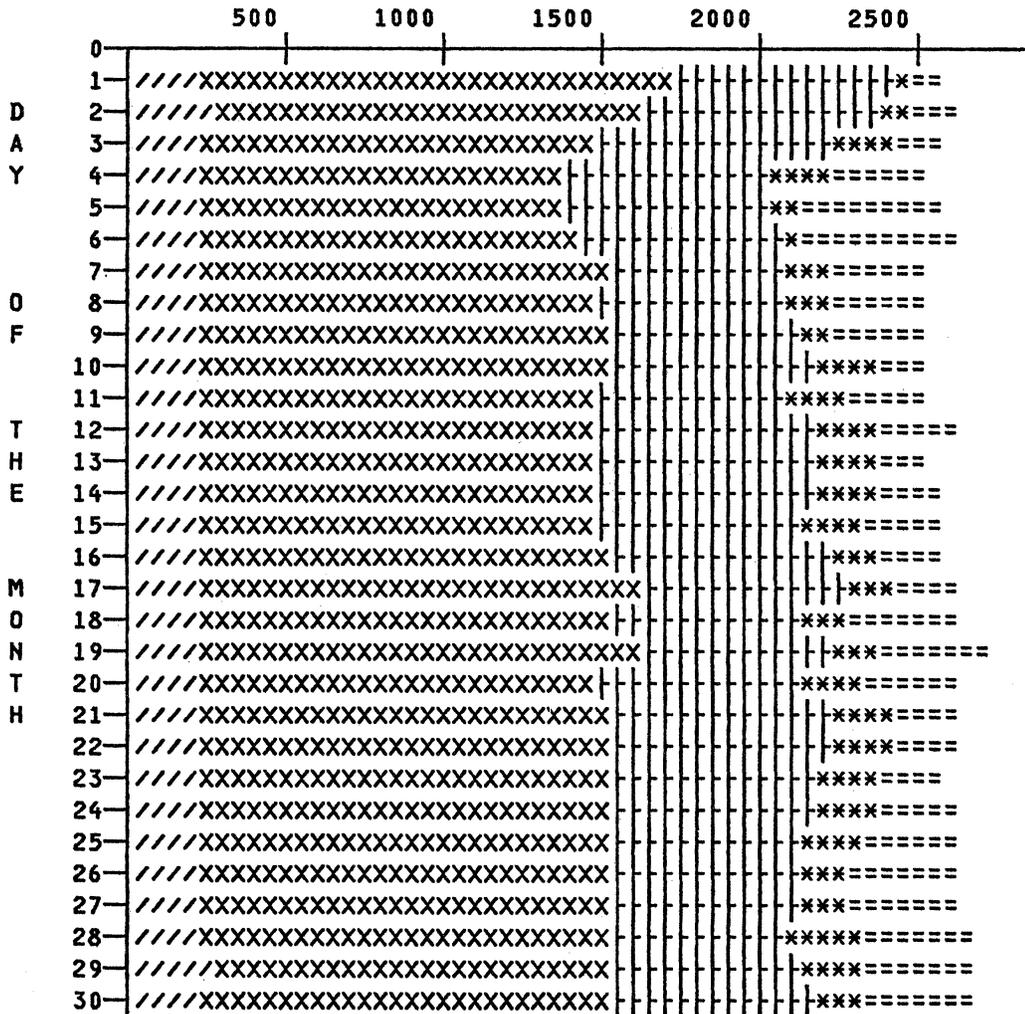
NUMBER OF BYTES OF ALLOCATED BLOCKS BY AREA (BELOW 16MEG)

SUBPOOL (AREA)	MIN	MAX	AVG
229	752K 11.00.00	752K 11.00.00	752K
230	84K 11.00.00	84K 11.00.00	84K
236 (SWA)	80K 11.00.00	80K 11.00.00	84K
237 (SWA)	12K 11.00.00	12K 11.00.00	12K
255 (LSQA)	32K 11.00.00	32K 11.00.00	32K
USER REGION			
0	1676K 11.00.00	1676K 11.00.00	1676K
251 (MODULES)	1104K 11.00.00	1104K 11.00.00	1104K
252 (REENTRANT)	8K 11.00.00	8K 11.00.00	8K

Figure 62. RMF Virtual Storage Activity Report.....Page 4.

CSA UTILIZATION BY KEY FOR HOST CMC1
OCTOBER 1984

K B Y T E S



KEY:
 / CSA KEY 1
 X CSA KEY 6
 | CSA KEY 0
 * SQA X CSA
 = CSA FREE

Figure 63. CSA Usage Activity Report (produced by user program)

The appendices are:

- A. "VTAM Parameter Changes" on page 167: Lists the changes to the VTAM start options and definitions in ACF/VTAM V3.
- B. "VTAM Application Program Interface (API) Changes" on page 171: Describes the VSCR related changes to the VTAM Application Program Interface in ACF/VTAM V3.
- C. "New VTAM Sense Codes And Messages" on page 175: Summarizes the changes to VTAM messages and sense codes in ACF/VTAM V3.
- D. "Request/Response Unit Changes" on page 177: Lists the RUs changed to accommodate the extended network address.
- E. "SNA Control Vector Changes" on page 179: Shows the changed SNA Control Vectors required for the ACTCDRM and ACTPU operations.
- F. "SSP Enhancements" on page 181: Summarizes the changes to the NCP Dump Formatter and Configuration Report Program in ACF/NCP V4.
- G. "Example of Symptom String Subset Data" on page 183: Illustrates the data that results from the function discussed in "Symptom String Subset" on page 42.

A.0 VTAM PARAMETER CHANGES

The following changes have been made to VTAM parameters in ACF/VTAM V3.

A.1 VTAM START OPTIONS

These are new or changed parameters that are specified in the VTAM start list (ATCSTRxx).

1. CSALIMIT

- Range: 0-2048M
- CSALIMIT can now be suffixed with a 'K' or 'M' to denote 'kilobytes' and 'megabytes' respectively.
- The largest specifiable CSALIMIT can be 2048M for MVS/XA systems or 16M for MVS/370 systems.
- In MVS/XA systems, the CSALIMIT value now discriminates between CSA and Extended CSA. Two values can be specified.

2. HOSTSA

- Range: 1-255
- HOSTSA may now have a value greater than MAXSUBA. An information message (IST796I) will be issued to inform of this condition.

3. ITLIM

- Range: 0-65535
- ITLIM will be used to pace requests in cross-domain and cross-network environments as well as in single-domain environments.

4. MAXSUBA

- Range: 3-255
- MAXSUBA does not have to be coded if there are no pre-ENA nodes in the network.
- MAXSUBA must be coded if any pre-ENA node exists in the network. This will allow ENA nodes to interpret and build pre-ENA addresses.

5. TRACE TYPE=VTAM (VTAM Internal Trace)

- There will be no defaults assumed if the OPTION= operand is omitted. The trace will be activated but will remain idle. Some severe error conditions will be traced.

6. MAXAPPL

- This parameter will be ignored by ACF/VTAM V3.
- It will be removed from ACF/VTAM V3 publications.

7. VTAMEAS

- This parameter will be ignored by ACF/VTAM V3. The value in ISTRACON will be used instead.
- It will be removed from ACF/VTAM V3 publications.

A.2 VTAM DEFINITION STATEMENTS

These are new or changed definitions for major and minor nodes in VTAMLST.

1. ENA Related Changes

- Those definitions that specify MAXSUBA, subarea address or element are now affected by the following changes:
 - The range of subarea and element addresses is not constrained by the MAXSUBA value.
 - The maximum subarea address is 255.
 - The maximum element address is 32767.
- BUILD (VTAM and NCP)
 - The MAXSUBA and SUBAREA operands should be reviewed.
- CDRM
 - The ELEMENT and SUBAREA operands should be reviewed.
- GWNAU (NCP only)
 - The ELEMENT operand should be reviewed.
- GWPATH
 - The ADJNETEL, ADJNETSA, ELEMENT and SUBAREA operands should be reviewed.

- HOST (VTAM and NCP)
 - The SUBAREA operand should be reviewed.
- NETWORK (VTAM and NCP)
 - The MAXSUBA and SUBAREA operands should be reviewed.
- PATH
 - The ADJSA and DESTSA operands should be reviewed.

2. APPL MAXPVT

- Range: 0-2048M
- The MAXPVT operand on the APPL definition statement can now be suffixed with a 'K' or 'M' to denote 'kilobytes' and 'megabytes' respectively.
- The largest specifiable MAXPVT can be 2048M for MVS/XA systems or 16M for MVS/370 systems,
- In MVS/XA systems, the MAXPVT value does not discriminate between the private area and the extended private area. They are treated as a single area.

3. CDRM RECOVERY

- Format: RECOVERY=YES|NO
- The new RECOVERY operand indicates whether the SSCP-SSCP session between the named CDRM and the host CDRM should be restarted automatically after a session outage.
- The default is YES.

For a complete description of VTAM parameters, refer to VTAM Installation and Resource Definition, SC23-0111.

B.0 VTAM APPLICATION PROGRAM INTERFACE (API) CHANGES

B.1 MVS/XA EXTENDED VIRTUAL STORAGE EXPLOITATION RELATED

VTAM API user exit routines and macro-based requests (e.g., OPNDST, SEND) can execute in either the 24-bit or 31-bit addressing mode. Figure 64 summarizes the addressing modes in which exits receive control from VTAM.

EXIT	ADDRESSING MODE *	COMMENTS
DFASY	1	Mode of original request if OPTCD=SYN.
LERAD	3	
LOGON	1	
LOSTERM	1	
NSEXIT	1	
RELRQ	1	
RESP	1	
RPL	2	
SCIP	1	
SYNAD	3	
TPEND	1	As for LERAD above.

- * The addressing mode will be determined by one of these factors:
1. The addressing mode used when the ACB was opened.
 2. The addressing mode used with the RPL-based request.
 3. The addressing mode used when the CHECK request was issued.

Figure 64. VTAM API Exits - Addressing Modes

The following describes changes to VTAM API macro instructions resulting from MVS/XA exploitation.

ACB Although the ACB will remain below the boundary, this macro will generate an ACB which conforms to the requirements of full 31-bit addressing. It may be accessed by programs executing in AMODE(31)/RMODE(24) or AMODE(24)/RMODE(24).

CHECK This will execute in the addressing mode used at the time the macro instruction was issued. All data areas referenced must be in a storage location compatible with the addressing mode. Otherwise, the results

will be unpredictable as programs in 24-bit addressing mode will ignore the high order address byte. It is the user's responsibility to ensure this compatibility.

- CLOSE** This can execute in AMODE(31)/RMODE(24) or AMODE(24)/RMODE(24). This macro instruction creates a parameter list which must be in 24-bit storage. However, a program in RMODE(ANY) may create its own parameter list in 24-bit storage and use the execute form of CLOSE.
- CLSDST** All RPL based macros will execute in the addressing mode used at the time the macro instruction was issued.
- EXECPPL** As for CLSDST above.
- EXLST** This will generate an EXLST with valid 31-bit pointers. The program will not require re-assembly as the expansion will be identical for both addressing modes.
- GENCB** All manipulative macro instructions can be used by programs executing in any addressing mode - RMODE(ANY)/AMODE(ANY). However, the data to be manipulated and parameter lists required by the macro instructions must be located below the boundary. These macro instructions will support the changes in size of VTAM control blocks.
- INTRPRET** As for CLSDST above.
- MODCB** As for GENCB above.
- NIB** This will generate an NIB with valid 31-bit pointers. The program will not require re-assembly as the expansion will be identical for both addressing modes.
- OPEN** As for CLOSE above.
- OPNDST** As for CLSDST above.
- OPNSEC** As for CLSDST above.
- RCVCMD** As for CLSDST above.
- REQSESS** As for CLSDST above.
- RESETR** As for CLSDST above.
- RPL** This macro instruction will generate an RPL which contains valid 31-bit pointers. The program will not require re-assembly as the expansion will be identical for both addressing modes.
- SEND** As for CLSDST above.
- SENDCMD** As for CLSDST above.
- SESSIONC** As for CLSDST above.
- SETLOGON** As for CLSDST above.

SHOWCB As for GENCB above.
SIMLOGON As for CLSDST above.
TERMSESS As for CLSDST above.
TESTCB As for GENCB above.

B.2 OTHER CHANGES

There are changes to the Communications Network Management Interface (CNMI). These result from ENA, the NMVT and IBM 3710 support.

B.3 NO CHANGE

There are no changes in ACF/VTAM V3 for the following VTAM RECORD API functions:

- Authorized Path
- Parallel Sessions
- Multi-memory Applications
- Programmed Operator Interface (POI)

New Sense Codes

1. 08120004
 - This new sense code indicates a pre-ENA address is not available.
2. 086B0000
 - This new sense code is issued if the SNA address list in the NMVT RU is incorrect.
3. 088E0000
 - This new sense code is issued if there are ENA incompatibilities (i.e., the LU has an address incompatible with MAXSUBA and GW_NCP is pre-ENA) between subareas.

New/Modified VTAM Messages

1. IST077I
 - This message now includes 'SLOWDOWN=YES' if the NCP is in slowdown.
 - The literal appears after the channel unit address.
2. IST171I
 - This message has been changed to include up to 5 digit session counts for both active sessions and session requests.
3. IST396I, IST397I
 - Two new commands describing cross-network route information.
4. IST526I
 - The message group displayed, beginning with IST526I (this includes IST527I, IST598I, IST5698I, IST569, IST570I) are summarized into one message - IST526I.
5. IST533I
 - Using the DISPLAY ROUTE command with the 'TEST=YES' option, solicits the message group beginning with IST533I.

6. IST799I

- This message is received after a DISPLAY ID= command when there is an NCP dump active. This message will be issued for DYNAMIC, MOSS, or CSP type dumps.
- Format: IST799I NCP|MOSS|CSP DUMP IN PROGRESS

7. IST815I

- This message is received after a DISPLAY ID= command.
- Issues if an SSCP-SSCP session recovery is supported for a CDRM.

8. IST838I

- A new message for a message group that is initiated by the DISPLAY TRACE command. The message group includes (IST838I, IST839I, IST840I, IST314I).

9. IST841I

- A new message for a message group that is initiated by the DISPLAY TRACE command.
- Format: IST841I NO RESOURCES ARE BEING TRACED FOR nnnnnnn

Non-Replicated VTAM Messages: The following VTAM messages are affected by the function described in "Elimination of Message Flooding" on page 38.

This means that replicates of these messages will be suppressed:

IST121I IST182I IST192I IST208I IST264I IST301I
IST342I IST343I IST344I IST345I IST346I IST348I
IST366I IST367I IST436I IST526I IST561I IST562I
IST563I IST564I IST565I IST566I IST659I IST663I
IST678I IST678I IST738I IST820I IST822I IST860I
IST999I

Refer to VTAM Messages and Codes, SC23-0114 for text.

The network address field in the following RUs has been replaced with the element address:

- X'01020F' ABCONN (Abandon Connection)
- X'010218' ABCONNOUT (Abandon Connection Out)
- X'010216' ACTCONNIN (Activate Connect In)
- X'01020A' ACTLINK (Activate Link)
- X'010302' ACTTRACE (Activate Trace)
- X'01020E' CONNOUT (Connect out)
- X'010201' CONTACT
- X'010280' CONTACTED
- X'010217' DACTCONNIN (Deactivate Connect In)
- X'01020B' DACTLINK (Deactivate Link)
- X'010303' DEACTTRACE (Deactivate Trace)
- X'010202' DISCONTACT
- X'010331' DISPSTOR (Display Storage)
- X'010208' DUMPFINAL
- X'010206' DUMPINIT (Dump Initial)
- X'010207' DUMPTXT
- X'010215' ESLOW (Enter Slowdown Mode)
- X'010301' EXECTEST (Execute Test)
- X'010215' EXSLOW (Exit Slowdown Mode)
- X'01021A' FNA (Free Network Addresses)
- X'410235' INITPROC (Initiate Procedure)
- X'010281' INOP (Inoperative)
- X'010205' IPLFINAL
- X'010203' IPLINIT (IPL Initial)

- X'010204' IPLTEXT
- X'810620' NOTIFY (SSCP→PU)
- X'410384' RECFMS (Record Formatted Maintenance Statistics)
- X'010381' RECMS (Record Maintenance Statistics)
- X'010334' RECSTOR (Record Storage)
- X'010382' RECTD (Record Test Data)
- X'410385' RECTR (Record Test Results)
- X'010383' RECTRD (Record Trace Data)
- X'41028A' REQACTLU (Request Activate Logical Unit)
- X'010284' REQCONT (Request Contact)
- X'01021B' REQFNA (Request Free Network Address)
- X'410304' REQMS (Request Maintenance Statistics)
- X'410210' RNAA (Request Network Address Assignment)
- X'410307' ROUTE_TEST
- X'010209' RPO (Remote Power Off)
- X'010211' SETCV (Set Control Vector)
- X'410305' TESTMODE

Control Vector X'06' is used during the ACTCDRM operation to pass the ENA support indicator.

0	X'06'
1	LENGTH IN BINARY
2	CDRM PROFILE
3	bit 0, name pair session support indicator bit 1, address pair session support indicator bit 2, parallel session support indicator bit 3, URC support indicator for all resources on SSCP bit 4, CDINIT (TYPE=DQ) bit 5, PCID session support indicator bit 6, CDSESEND support indicator bit 7, reserved
4	bit 0, PLU capability indicator bit 1, network-qualified address pair support indicator bit 2, INIT_OTHER_CD format 2 support indicator bit 3, INIT_OTHER_CD format 3 support indicator bit 4, format 3 and 4 of CDINIT support indicator bit 5, format 1 of CDINIT support indicator bit 6, NOTIFY NS support indicator bit 7, notification of lost session awareness indicator
5	bit 0, <u>ENA SUPPORT</u>

Figure 65. Control Vector X'06'

Control Vector X'0B' is used during the ACTPU operation to pass the ENA support indicator.

0	X'0B'	
1	LENGTH IN BINARY	
2-n	bit 0,	Lost Subarea Requirement
	bit 1,	ALS station support
	bit 2,	Gateway support
	bit 3,	Notification of other-network lost route
	bit 4,	Notification of same-network lost route
	bit 5,	CONTACTED (loaded) format
	bit 6,	reserved
	bit 7,	reserved
3	bit 0,	SSCP T2.1 node support
	bit 1,	<u>ENA Support</u>
	bits2-7,	000000

Figure 66. Control Vector X'0B'

F.1 NCP DUMP FORMATTER ENHANCEMENTS

A number of improvements have been made:

1. Control blocks are now formatted and printed with their configured group in a hierarchical fashion.
2. Format and print the Level 1, 3, 4, and 5 Save Areas for the 3705 and the Level 1, 2, 3, 4, and, 5 Save Areas for the 3725. The return address will also be printed.
3. The following tables will now be printed, in chronological order:
 - Dispatcher Trace Table (DTT) for the 3725 only
 - Address Trace Block (ATB)
 - Channel Adapter Trace Table (CA) for the 3725 only
 - Dispatch Priority Table (DPT) and QCB addresses in chain
4. The complete prefix is printed when a buffer pool is printed. This is a change from previously only printing the first byte of the prefix.
5. The following information/control blocks are now formatted and printed:
 - Queue Control Block (QCB)
 - Load map of the 3725 and HJN of each module.
 - The module and displacement in which an abend occurred.
 - HWE and HWX
 - Committed Buffers Block (CBB) for the CUB, SCB, and LCB
 - The status of the system at the time of a dump (i.e., SLOWDOWN, PSEUDO SLOWDOWN, CWALL, etc.) will be printed.
 - The Level 1 registers and external registers for the 3725.
 - The MOSS Interface Control Block and the MOSS Trace Facility
 - ABEND Control Block
 - All the storage protect keys for the 3725 static dump
6. The SNA or non-SNA control blocks can be formatted and printed separately.

F.2 CONFIGURATION REPORT PROGRAM (CRP)

CRP has been improved with the following:

1. Syntax checking is now performed on the *REPORT and *OPTION statements
2. Printing of Stage 1 deck is now optional
3. Output lines/page is user specifiable
4. Subchannel addresses of an EP line will now be displayed in the report
5. CRP will not process the DIALNO operand
6. The CUTYPE column has been deleted from the SNA report
7. The BATCH operand has been added to the SNA reports
8. The GPOLL operand for the CLUSTER macro is now printed
9. The TGN operand is now printed
10. The line control of each LINE macro will be printed
11. CRP will only compute the element address for devices in a 3725 gen. The NET ADDR heading in the report will be changed to ELMT ADDR

G.0 EXAMPLE OF SYMPTOM STRING SUBSET DATA

This extract from a SVCDUMP shows the additional information that results from ACF/VTAM's use of the MVS/XA SDWA Variable Recording Area (SDWAVRA) as discussed in "Symptom String Subset" on page 42.

The new information is shown in HIGHLIGHT.

MODULE SVCDUMP DATE ... TIME ... PAGE 000001

TITLE FROM DUMP: ISTAPCES-VTAM IRB DUMP

ERRORID FROM THIS DUMP = SEQ00023 CPU00 ASID000B TIME03.27.20.0

ACTIVE CPU'S AT TIME OF DUMP

ADDR	VERS.	SERIAL	MODEL
0000	FF	021368	3084

***** DUMP ANALYSIS AND ELIMINATION (DAE) *****

THIS DUMP WAS NOT SUPPRESSED BECAUSE
THE VRA KEY TO ALLOW SUPPRESSION OF DUPLICATE DUMPS WAS ABSENT.

CRITERIA FOR USE AS A UNIQUE DUMP IDENTIFIER BY DAE:

MINIMUM NUMBER OF SYMPTOMS: 05 FOUND: 06
MINIMUM TOTAL STRING LENGTH: 025 FOUND: 086

SYMPTOMS REQUIRED TO BE PRESENT:

MOD/ CSECT/

SYMPTOMS THAT ARE TO BE USED IF AVAILABLE, BUT ARE NOT REQUIRED:

AB/S AB/U REXN/ FI/ REGS/ HRC1/ CID1/ SUB1/

MVS SYMPTOM STRING:

CSECT/ISTCPCSB AB/S00C4 REXN/ISTAPCES FI/4070504070705870D08C41

REGS/09170 CID1/28901

RETAIN SEARCH ARGUMENT:

RIDS/ISTCPCSB AB/S00C4 RIDS/ISTAPCES#R VALU/H70D08C41 REGS/09170

VALU/C28901

SYMPTOMS PRESENT FOR USE AS A UNIQUE DUMP IDENTIFIER BY DAE:

<u>MVS KEY</u>	<u>RETAIN KEY</u>	<u>SYMPTOM DATA</u>	<u>EXPLANATION</u>
CSECT/	RIDS/	ISTCPCSB	ASSEMBLY MODULE CSECT NAME
AB/S	AB/S	S00C4	ABEND CODE-SYSTEM
REXN/	RIDS/	ISTAPCES	RECOVERY ROUTINE CSECT NAME
FI/	VALU/H	4070504070705870D08C41	FAILING INSTRUCTION AREA
REGS/	REGS/	09170	REGS/PSW DIFFERENCE
CID1/	VALU/C	28901	COMPONENT IDENTIFIER

ADDITIONAL SYMPTOM DATA NOT USED BY DAE TO IDENTIFY THIS DUMP:

<u>MVS KEY</u>	<u>RETAIN KEY</u>	<u>SYMPTOM DATA</u>	<u>EXPLANATION</u>
AMD1/	VALU/C	84#107	MODULE ASSEMBLY DATE
VRS1/	VALU/C	310	VERSION-PRODUCT/PTF ID
CDB1/	VALU/C	5662	BASE COMPONENT ID
ASID1/	VALU/H	000B	TASK RELATED ASID

**SNA ENA AND VIRTUAL
STORAGE GUIDE
GG24-1664**

**READER'S
COMMENT
FORM**

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Your comments will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

Note: Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

NOTE: Staples can cause problems with automated mail sorting equipment.
Please use pressure sensitive or other gummed tape to seal this form.

What is your occupation? _____

Cut or Fold Along Line

Reader's Comment Form

Fold and tape

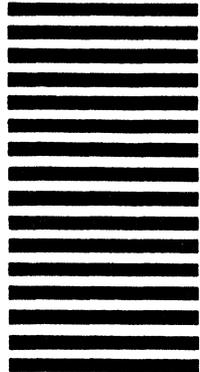
Please Do Not Staple

Fold and tape



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.



POSTAGE WILL BE PAID BY ADDRESSEE:

IBM International Systems Center
Department 986, Building 622-3
P.O. Box 12195
Research Triangle Park
Raleigh, North Carolina 27709
U.S.A.

Fold and tape

Please Do Not Staple

Fold and tape



GG24-1664-00

SNA ENA and Virtual Storage Guide

GG24-1664-00

PRINTED IN THE U.S.A.

IBM

GG24-1664-0

