

Systems

**OS/VS2 Planning Guide
for Release 2**

IBM

First Edition (March, 1973)

This edition applies to Release 2 of OS/VS2 and to all subsequent releases until otherwise indicated in new editions or Technical Newsletters. Changes are continually made to the information herein; before using this publication in connection with the operation of IBM systems, consult the latest **IBM System/360 and System/370 Bibliography**, GA22-6822, for editions that are applicable and current.

Requests for copies of IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

A form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, Publications Development, Department D58, Building 706-2, PO Box 390, Poughkeepsie, N.Y. 12602. Comments become the property of IBM.

Preface

This publication contains planning information for OS/VS2 Release 2. It is intended for installation managers and system programmers responsible for assessing the effort required to install an OS/VS2 Release 2 system. Readers should have a background in programming and maintaining computer operating systems such as OS/MVT or OS/VS2 Release 1.

This publication is for planning purposes only. The functions and capabilities described reflect the information that is currently available. This information is subject to change before the availability of OS/VS2 Release 2. Further information will be published when OS/VS2 Release 2 becomes available.

This manual contains the following independent chapters:

- Introduction -- The Introduction highlights the major points that should be considered in planning for the installation and implementation of VS2 Release 2.
- Defining the System -- This chapter provides preliminary planning information for defining the type of system that an installation needs. For example, the chapter describes the VS2 Release 2:
 - procedures and macro instructions for system generation.
 - procedures and parameters for system initialization.
 - system libraries and data sets.

For each topic, a summary of the major changes from both MVT and VS2 Release 1 is provided.

- Directing the Use of System Resources -- This chapter describes some of the new VS2 Release 2 facilities that allow an installation to control the distribution of system resources among system users. It provides planning information that helps assess the effort involved in implementing these facilities.
- Reducing I/O Device Allocation Contention Among Job Steps -- This chapter contains recommendations for using the new, unserialized path through the I/O device allocation routines.
- System Integrity -- This chapter discusses integrity in a multi-user computer system. It includes recommendations for maintaining, in any control program extensions or modifications, the same level of system integrity as that of the VS2 Release 2 system.
- Conversion Considerations -- This chapter presents preliminary descriptions of some of the most significant differences between VS2 Release 2 and MVT or VS2 Release 1. It provides the information to assess the extent of the modifications required for existing data sets, cataloged procedures, programs, data reduction routines, and operational procedures, to install VS2 Release 2.
- Appendix A: JES2 Differences from MVT/HASP and VS2 Release 1/HASP -- This appendix is a table showing the differences between JES2 and HASP (on either MVT or VS2 Release 1 systems).
- Appendix B: Virtual Storage Map -- This appendix shows the layout of virtual storage in VS2 Release 2.

- Glossary -- The glossary describes VS2 terms.

The following items are described in this publication strictly for advance planning purposes and will not be available with the initial release of VS2 Release 2:

- JES3 job entry subsystem.
- Virtual telecommunications access method (VTAM).

Additional information about these items will be available later. Availability dates for the support of the items may be obtained from the local IBM branch offices.

Prerequisite Publication

IBM System/370 Introduction to VS2 Release 2, GC28-0661

This publication describes VS2 Release 2 features and facilities, design concepts, and basic hardware requirements.

Related Publications

Introduction to Virtual Storage in System/370, GR20-4260

This publication provides an overview of the advantages and concepts of virtual storage systems.

IBM Data Processing Glossary, GC20-1699

This publication defines terms and concepts used in IBM publications.

OS/VS Virtual Storage Access Method (VSAM) Planning Guide, GC26-3799

This publication contains a detailed description of the virtual storage access method (VSAM).

IBM System/370 OS/VS1 Planning and Use Guide, GC24-5090

This publication describes OS/VS1. It can assist installation personnel in selecting and evaluating VS1 and can assist system programmers in implementing, modifying, and extending the capabilities of the VS1 control program.

IBM System/370 OS/VS2 Planning and Use Guide, GC28-0600

This publication describes OS/VS2 Release 1. It introduces VS2 concepts and provides planning information for users planning to have VS2 Release 1 installed.

Introduction to VTAM, GC27-6987

This publication describes the functions and capabilities that an installation needs to know to plan for installing the virtual telecommunications access method (VTAM).

Contents

Introduction	9
VS2 Release 2 Highlights	10
Multiprocessing	10
Job Entry Subsystems	10
Data Set Handling	10
System Resources Management	11
System Activity Measurement Facility	11
Multiple Virtual Address Spaces	11
VS2 Release 2 Planning Considerations	12
Functions Not Supported	13
Unsupported Type I Programs	15
Defining the System	17
System Generation Procedures	18
The VS2 Release 2 Starter System	19
Minimum Configuration	19
Physical Addresses for Devices	19
Stage I Macro Instructions	20
Job Entry Subsystem Generation	21
JES2 Generation	22
JES2 Generation Options	22
RMTGEN	23
RMTGEN Options	23
System Generation Changes	24
Initializing the New System Data Sets	26
Adding Installation-Written Members to System Libraries	26
Unsupported VS2 Release 1 Stage I Macro Instructions	27
System Initialization	28
VS2 Release 2 System Parameters	28
Changes to System Parameters for VS2 Release 2	34
Changes to VS2 Release 1 System Parameters	34
Unsupported VS2 Release 1 System Parameters	35
Unsupported MVT System Parameters	35
System Parameter Specification	36
System Parameter Library (SYS1.PARMLIB)	37
Changes to SYS1.PARMLIB Lists	38
IBM-Created Lists	39
Installation-Created Lists	43
System Libraries and Data Sets	45
Required Libraries and Data Sets	45
Optional Libraries and Data Sets	46
Changes to VS2 Release 1 System Libraries and Data Sets	47
Changes to MVT System Libraries and Data Sets	48
Adding Paging Space to the System	48
Directing the Use of System Resources	49
System Resources Management	50
Workload Management Routines	50
Resource-Use Routines	50
Control Function	51
The Installation Performance Specification (IPS) and the Service Concept	52
Service	52
Performance Groups	53
Performance Objectives	53
Associating a Transaction With a Performance Objective	56
The Relationship Between Workload Levels and Service Rates	57
System Resources Manager Parameters	60
IPS Parameters (IEAIPSxx)	60
Tuning Parameters (IEAOPTxx)	62
System Activity Measurement Facility (MF/1)	63
MF/1 Operation	63
MF/1 Reports and Tracing Records	65
CPU Activity Report	65
Paging Activity Report	65
Workload Activity Reports	66

Performance Group Period Report	66
Performance Group Report	66
System Summary Report	67
Channel Activity Report	67
I/O Device Activity Reports	67
Tracing	67
Tuning with MF/1	68
Tuning Example	68
Reducing I/O Device Allocation Contention Among Job Steps	73
Defining A Reserved Mount Characteristic	74
Defining Device Groupings	74
System Integrity	75
Installation Responsibility	76
Areas of Concern	77
User-Supplied Addresses for User Storage Areas	77
User-Supplied Addresses for Protected Control Blocks	77
Resource Serialization	78
Resource Identification	79
SVC Routines Calling SVC Routines	80
Control Program and User Data Accessibility	80
Control Program Extensions	81
Installation Procedural Responsibilities	82
Conversion Considerations	85
SMF Considerations	86
Differences From VS2 Release 1 System Management Facilities	86
Differences From MVT System Management Facilities	86
SMF Record Differences for VS2 Release 2	86
New System Resources Manager Data	86
SMF Recording Replaced by MF/1	87
New Storage Data	90
Job Entry Subsystem Statistics	90
Multiprocessing Data	90
Virtual Storage Access Method Data	90
Time Sharing Record Elimination	90
Changes to SMF Exits for VS2 Release 2	91
Job Control Language	92
Changes to VS2 Release 1 JCL	92
Performance Group Specification	92
Assigning Job Selection Priority	92
Specifying Data Sets in Paging Space (Virtual I/O Data Sets)	94
Requesting Data Set Freeing at Data Set Close Time	94
Specifying the Number of Resources That May Be Held for Reuse	94
Requesting Multiple Copies of an Output Data Set	95
Directing the Destination of an Output Data Set	95
Format Record Identification for the 3886 OCR Device	95
Output Limit Specification	95
Increased Data Set Definition Capability	95
Increased Volume Sharing Capability	95
Interpretation of the REGION Parameter	95
Job Step Dispatching Priority Default Value	96
Changes to MVT JCL	96
Requesting Storage for Program Execution	96
Specifying the BTAM Online Terminal Test Option	96
JES2 JCL Considerations	96
JES2 JCL for non-HASP Users	96
JES2 Control Cards	97
Assigning Job Selection Priority	97
JES2 JCL for HASP Users	97
/*JOBPARM Control Card	97
/*OUTPUT Control Card	97
Operator Command Language	99
Changes to Operator Commands	99
Multiprocessing Support	102
VTAM Support	102
Time Sharing (TSO) Support	103
System Resources Management Support	103
Job Entry Subsystem Support	103
Job Related Functions	103

Device List Functions	104
System Functions	104
Miscellaneous Display Function	105
Remote Job Entry Function	105
Miscellaneous Support	105
Displays	105
Dumps	105
Messages	105
LSQA	105
System Log	105
Time Sharing Considerations	106
System Resources Management Support	106
VSAM Support	106
Allocation Support	106
Job Entry Subsystem Support	110
Edit Support	110
Accounting Support	110
Time Sharing Integration Support	110
Data Set Conversion	111
Catalog Conversion	111
Catalog Utility and Macro Changes	113
Creating or Updating the VSAM Catalog	115
Updating an OS Catalog From a VSAM Catalog	116
Reformatting the UADS	116
Formatting the Broadcast Data Set	116
Program Conversion	117
TCAM Conversion	117
Multiprocessing Considerations	118
Hardware Configuration Considerations	118
Operator Control of Configuration	118
Programmer Considerations	119
Appendix A: JES2 Differences From MVT/HASP and VS2 Release 1 HASP	121
Appendix B: Virtual Storage Map	123
Glossary	125
Index	131

Figures

Figure 1.	Unsupported MVT Functions (Part 1 of 2)	13
Figure 1.	Unsupported MVT Functions (Part 2 of 2)	14
Figure 2.	Unsupported VS2 Release 1 Functions	14
Figure 3.	Required Device Physical Addresses for the Starter System	20
Figure 4.	Generating a VS2 Release 2 System	22
Figure 5.	Changes in Stage I Macro Instructions	25
Figure 6.	System Libraries to Which Installation-Written Members May Be Added	27
Figure 7.	System Initialization Parameter Summary	29
Figure 8.	Changes to SYS1.PARMLIB Lists	38
Figure 9.	Performance Objective Definition	55
Figure 10.	Associating a Transaction With a Performance Objective	56
Figure 11.	Performance Objectives	57
Figure 12.	System Workload Levels	58
Figure 13.	Service Rates for Different Performance Objectives at Various System Workload Levels	59
Figure 14.	Tuning Example	69
Figure 15.	Changing a Performance Objective	70
Figure 16.	Effects of a Performance Objective Change	71
Figure 17.	Changes to SMF Records for VS2 Release 2 (Part 1 of 2)	88
Figure 17.	Changes to SMF Records for VS2 Release 2 (Part 2 of 2)	89
Figure 18.	Changes to SMF Exits for VS2 Release 2	91
Figure 19.	JCL Changes for VS2 Release 2	93
Figure 20.	Changes to Operator Commands for VS2 Release 2 (Part 1 of 3)	100
Figure 20.	Changes to Operator Commands for VS2 Release 2 (Part 2 of 3)	101
Figure 20.	Changes to Operator Commands for VS2 Release 2 (Part 3 of 3)	102
Figure 21.	Changes to Time Sharing Commands for VS2 Release 2 (Part 1 of 3)	107
Figure 21.	Changes to Time Sharing Commands for VS2 Release 2 (Part 2 of 3)	108
Figure 21.	Changes to Time Sharing Commands for VS2 Release 2 (Part 3 of 3)	109
Figure 22.	Simplified Catalog Structure - OS vs. VSAM	112
Figure 23.	Some OS/VSAM Catalog Relationships	112
Figure 24.	Functions of the Access Method Services	115
Figure 25.	JES2 Differences From MVT/HASP and VS2 Release 1/HASP	121
Figure 26.	Virtual Storage Map	123

Introduction

OS/VS2 Release 2 is a virtual storage operating system with multiprogramming, multiprocessing, time sharing (TSO), and job entry subsystems. In addition to providing significant new features, it also enhances existing OS/MVT and OS/VS2 facilities. OS/VS2 Release 2 is generally upward compatible with OS/MFT, OS/MVT, OS/VS1 and OS/VS2 Release 1.

VS2 Release 2 Highlights

Some of the significant new features provided in VS2 Release 2 include multiprocessing, job entry subsystems, new data set handling facilities, a centralized resource management facility, a system activity measurement facility, a virtual I/O paging mechanism for temporary data sets, and multiple virtual address spaces. These are described in the following sections. For a more detailed description of these features and of the VS2 Release 2 system, see the **Introduction to VS2 Release 2**.

Multiprocessing

The incorporation of multiprocessing into the VS2 Release 2 control program provides for growth to larger and more complex systems. CPUs can be combined in two types of combinations: loosely coupled or tightly coupled. An installation can use combinations of loosely coupled and tightly coupled multiprocessing to tailor configurations to meet individual requirements.

In a tightly coupled multiprocessing system, two CPUs share real storage, operate under a single control program, and communicate with each other by means of an interrupt capability. The system control program treats the CPUs as resources, assigning them to process tasks.

Loosely coupled multiprocessing systems connect VS2 Release 2 systems by means of channel to channel adapters that pass information between the processors. One processor controls job selection for all of the processors in the system. The JES3 job entry subsystem, a successor of ASP, provides the control for loosely coupled multiprocessing systems. One or more of the JES3 processors can, itself, be a tightly coupled multiprocessor.

Job Entry Subsystems

JES2 and JES3 provide the job processing functions of:

- Reading local and remote jobs into the system.
- Scheduling jobs.
- Maintaining all data submitted with jobs.
- Supporting the system management facilities.
- Handling all output from batch jobs and time-sharing users.

Either JES2 or JES3 may be specified during system generation to run as a component of the system control program.

Data Set Handling

The virtual storage access method (VSAM) is designed to offer function and flexibility to online and data base environments. In VS2 Release 2, the key-sequenced VSAM system catalog structure is designed to reduce access time to very large catalogs. It permits the use of multiple private catalogs and the master catalog in one step, job, or session.

Virtual I/O (VIO), is a new way of processing temporary data sets. It is accomplished on a page basis in virtual storage.

The virtual telecommunications access method (VTAM) is a new direct-control teleprocessing access method with facilities available to applications programs, including those using TCAM.

System Resources Management

A new facility, called the system resources manager, is a centralized decision-maker that monitors a wide range of data about the condition of the system. It attempts to provide system resources to jobs in a manner specified by the installation. This allows an installation to specify that certain jobs get better response time than other jobs.

System Activity Measurement Facility

The system activity measurement facility (MF/1) is both a new and standard feature in VS2 Release 2. It collects information about system activities and produces trace records and reports. MF/1 is basically system oriented (as opposed to SMF which is basically job oriented). It is designed to aid in analyzing both system performance and system trends and in anticipating future system requirements.

Multiple Virtual Address Spaces

In VS2, virtual storage is an address range of 16,777,216 bytes. Release 1 of VS2 provided one virtual address range; jobs were assigned regions from this address range, just as jobs were assigned regions in real storage in MVT. VS2 Release 2 provides multiple virtual address spaces; each user (i.e., batch job, time-sharing job, and certain system components) receives its own copy of virtual storage, minus the space used for certain system functions (e.g., the nucleus, the link pack area, the system queue area).

The new virtual storage design provides the potential for greater multiprogramming, including the concurrent execution of more large jobs and a greater mix of time-sharing and batch jobs. Also, since each job is isolated in a private address space, inter-region virtual storage fragmentation is eliminated and protection features are extended.

VS2 Release 2 Planning Considerations

To make use of some of the new or changed VS2 Release 2 facilities, it will be necessary to modify some existing procedures, libraries, data sets, and programs. The following list highlights examples of some of the more significant Release 2 changes that may require modifications to be made by an installation. Neither the list nor the examples are meant to be exhaustive; rather they call attention to areas of change that should be considered when planning to install a VS2 Release 2 system.

The actual details of the changes are described in this publication in succeeding chapters.

- While the selection and specification of **Stage I system generation** options has been simplified, it is now necessary to select and specify job entry subsystem options.
- **VS2 Release 2 system initialization** is designed to enable minimum operator intervention by extending the options that can be specified in `SYS1.PARMLIB`.
- The **system catalog** structure is now a VSAM data set that can handle both VSAM and other data set types. Existing OS system catalogs (`SYSCTLG`s) must be converted to the new VSAM structure.
- The new **multi-function access method services** are used to convert an OS system catalog to a VSAM master catalog. Access method services also replace most of the catalog functions of system utilities such as `IEHMOVE`, `IEHLIST`, `IEHPROGM`.
- **Paging space** is contained in VSAM data sets in VS2 Release 2. The page data sets must be defined and preformatted during system generation using access method services.
- **System resources management** routines allow the specification of relative rates for system resource scheduling for various categories of batch and time sharing jobs. The installation identifies the categories by assigning each of them a performance group number to be specified in `JCL`. The installation specifies the distribution rates and other tuning options in `SYS1.PARMLIB` members.
- The new **system activity measurement facility (MF/1)** collects and records information about system activities in either the SMF data set or on formatted and printed reports. Some SMF data, such as CPU or paging activity is now recorded only by MF/1 on new SMF records.
- **System management facilities (SMF)** records have been extended to record data provided by new features such as the system resources manager, the virtual storage access method, or job entry subsystems. Data reduction routines may have to be revised to process new records and new or changed record fields.
- **Job control language** parameters have been added to specify items such as performance groups. Job entry subsystem control cards may be used to specify JES2 or JES3 job processing options.
- The **operator commands** have been extended to support multiprocessing and changed for communication with the job entry subsystems.

Functions Not Supported

Figure 1 shows significant MVT functions that are not supported in VS2 Release 2. Figure 2 shows unsupported VS2 Release 1 functions.

Unsupported MVT Function	Comparable VS2 Release 2 Function
Automatic SYSIN batching (ASB) reader BACKSPACE, NOTE, POINT, XDAP, and EXCP addressed to SYSIN/SYSOUT DCBs	Job entry subsystem (JES2 or JES3) ----
Conversational remote job entry (CRJE) CVOL catalog Dedicated work files	Time sharing (TSO) VSAM user catalog Virtual I/O (VIO)
Direct system output (DSO) Graphic job processor (GJP)	---- ----
IEBUPDAT utility program IEHIOSUP utility program IEHMOVE, IEHLIST, IEHPROGM (most catalog functions)	IEBUPDTE utility program ---- Access method services
IMBMDMAP IMCJQDMP service aid IMCOSJQD service aid	AMBLIST Job entry subsystem and SWA dump routines ----
Main storage hierarchy support OS readers and writers Queued telecommunications access method (QTAM)	---- Job entry subsystem (JES2 or JES3) Telecommunications access method (TCAM)
QTAM message processing programs Remote job entry (RJE) Rollout/rollin Satellite graphic job processor (SGJP)	TCAM application programs JES2 or JES3 remote job entry Paging ----
Scatter load System environment recording routines (SER0 and SER1) OS system catalog (SYSCTLG)	Paging Recovery management support (RMS) VSAM master catalog

Figure 1. Unsupported MVT Functions (Part 1 of 2)

Unsupported MVT Function	Comparable VS2 Release 2 Function
SYS1.ACCT data set	System management facilities (SMF) data sets
SYS1.SYSJOBQE	Job entry subsystem input and output queues and the scheduler work area (SWA)
Tape output for SMF	----
TESTRAN program	----
Time slicing	System resources manager
Trace	Generalized trace facility (GTF)
Transient areas	Pageable link pack area
TSO driver	System resources manager
TSO trace	Generalized trace facility (GTF)

Figure 1. Unsupported MVT functions (Part 2 of 2)

Unsupported VS2 Release 1 Function	Comparable VS2 Release 2 Function
BACKSPACE, NOTE, POINT, XDAP, and EXCP addressed to SYSIN/SYSOUT DCBs	----
Directed swap for TSO	----
IEHMOVE, IEHLIST, IEHPROGM (most catalog functions)	Access methods services
IMCOSJQD service aid	Job entry subsystem and SWA dump routines
Nucleus-only SYSGEN	----
OS readers and writers	Job entry subsystem (JES2 or JES3)
OS system catalog (SYSCTLG)	VSAM master catalog
Priority I/O queueing and ordered seek queueing	----
QTAM message processing programs	TCAM application programs
SYS1.SYSJOBQE	Job entry subsystem input and output queues and the scheduler work area (SWA)
Time slicing	System resources manager
Trace	Generalized trace facility (GTF)
TSO driver	System resources manager
TSO link pack area	----
TSO trace	Generalized trace facility (GTF)

Figure 2. Unsupported VS2 Release 1 Functions

Unsupported Type I Programs

The following Type I programs which were supported by MVT and VS2 Release 1 are not supported by VS2 Release 2 (i.e., APARs written against these programs running on VS2 Release 2 will not be accepted). Corresponding program products are supported in their place. (Contact your local IBM branch office for information about these program products).

Even though the following Type I language compilers are not supported under OS/VS2 Release 2, generated object code for an installation's programs that use any of these higher-level Type I languages, together with their corresponding library modules, will operate on this release of OS/VS2 subject to the constraints outlined in the "Program Conversion" section of this publication.

Program Name	Program Number
COBOL F	360S-CB-524
COBOL F Library	360S-LM-525
Full ANS Cobol V2	360S-CB-545
Full ANS Cobol V2 Library	360S-LM-546
FORTTRAN G	360S-FO-520
FORTTRAN H	360S-FO-500
FORTTRAN G and H Library	360S-LM-501
PL/1 F	360S-NL-511
PL/1 F Library	360S-LM-512
PL/1 F Syntax Checker	360S-PL-552
SORT/MERGE	360S-SM-023

Defining the System

During system generation and system initialization, the installation can define or change the operation of a standard or optional feature of the system control program to meet specific needs. This chapter provides planning information for defining the system. It is divided into the following sections:

- System Generation Procedures
- System Initialization
- System Libraries and Data Sets

System Generation Procedures

System generation is the process of selecting data from IBM-distributed libraries and tailoring it to create an operating system for an installation.

An I/O device-only system generation can also be performed. It is performed when only I/O devices and channels are to be added, deleted, or modified. The processor-only generation available with MVT and the nucleus-only generation available with both MVT and VS2 Release 1 are not available in VS2 Release 2.

The VS2 Release 2 system generation process is simplified over the MVT and VS2 Release 1 system generation processes in the following ways:

- Many previously optional facilities are now standard. Therefore, fewer macro instructions and parameters need to be specified.
- Several user-specified macro instructions have been clarified by consolidating parameters for related options under a single macro instruction.
- Additional system initialization lists are available to be preformatted and taken from the SYS1.PARMLIB data set for use at IPL time.

For VS2 Release 2, IBM will generate (from installation-provided system generation statements) and install one system control program per uniprocessing or multiprocessing configuration. The IBM installation process also includes performing the installation verification procedure (IVP).

There are six major steps in generating a new operating system.

Planning

The distribution libraries containing the VS2 Release 2 modules must be ordered from IBM. The options must be selected that, with the standard features, will constitute the desired system. The macro instructions needed to specify the options and standard features of the new system must be selected and coded.

Volume and data set initialization

The direct access volumes that will contain the distribution libraries (and the starter system) must be initialized. The volumes on which the new system will be generated must also be initialized. The data sets on the system volumes must be allocated (unless they will be specified by SYSGEN macros to be assembled in Stage I for allocation in Stage II).

Stage I

The macro instructions are assembled and expanded into job control language, utility control statements, assembler language statements, and linkage editor control statements.

Stage II (Jobstream Execution)

The modules from the distribution libraries are assembled, link-edited, and copied to the data sets that have been allocated on the new system volume(s).

Job entry subsystem generation

The job entry subsystem modules are assembled (JES2 only) and link-edited to reflect the options and limits specified for them. Job entry subsystem

generation may be performed concurrently with Stage II. Additionally, remote station programs may optionally be generated (RMTGEN).

Testing

After the new system is generated, the IBM Program Systems Representative uses the installation verification procedure (IVP) to verify that the new system is operational on the hardware configuration.

The VS2 Release 2 Starter System

The first time that VS2 Release 2 is generated, the IBM starter system is required to drive the system generation processing. The starter system consists of a VS2 Release 2 operating system that uses JES2 as the job entry subsystem. After the first system generation, either the starter system or the new VS2 Release 2 control program can be used to drive subsequent system generations.

Minimum Configuration

Use of the starter system requires the following minimum configuration:

- One IBM System/370 Model 145, 155II, 158, 165II, or 168. (The clock comparator, the CPU timer feature (#2001) and the advanced control program support feature (#1001) are required on the Model 145.)
- 768K bytes of real storage.
- One multiplexer channel.
- One selector or block multiplexer channel.
- Three IBM 3330 spindles, or four IBM 2314/2319 Direct Access Storage Facility devices.
- One card reader.
- One hardcopy console or one hardcopy device with a display console.
- One 9-track magnetic tape.
- One 1403 or 3211 printer.

Physical Addresses for Devices

The starter system is configured to support the devices at the addresses specified in Figure 3.

Function	Device	Multiplexer Channel	Channel 1	Channel 2	Channel 3	Channel 4
Console	158 Console	014, 016				
	3213	015, 017				
	3210/3215	009, 01F		209, 21F		
	3066	040, 060 0C0, 0D0 0E0, 0F0	140	240	340	440
Reader	2540	00C		20C		
	3505	012		20A		
Punch	2540	00D		20D		
	3525	013		20B		
Tape	2400 (7-TR-DC)		180, 181	280, 281	380, 381	480, 481
	2400 (9-TR) 3400		182, 183 184	282, 283 284	382, 383	482, 483
Printer	3211	002, 004				
	1403	00E, 00F		20E		
Direct Access Storage Device	2305-1		1F0	2F0		
	2305-2		1D0	2D0		
	2314/2319		130, 131 132, 133 134, 135	230, 231 232, 233 234, 235	330, 331 332, 333	
	3330		150, 151 152, 153	250, 251 252, 253	350, 351 352, 353	

Figure 3. Required Device Physical Addresses for the Starter System

Stage I Macro Instructions

The Stage I macro instructions are used to tailor the new system. They describe the machine configuration, the control program, the data management routines, and the user-written routines. They are also used to specify the program options that are to be included in the system.

The macro instructions that can be specified in Stage I of VS2 Release 2 are:

AFFINITY

specifies program names (e.g., emulators) and the identification of the CPU on which each program is required to execute.

CHANNEL

specifies the channel characteristics. A CHANNEL macro instruction is required for each channel in the installation's computing system.

CKPTREST

specifies standard system completion codes not eligible for automatic restart, and user completion codes that are eligible for automatic restart. If the CKPTREST macro instruction is not specified, the standard set of codes will be used.

CONSOLE

specifies the console configuration for the system.

CTRLPROG

specifies control program options. If the CTRLPROG macro instruction is not specified, the default values are assumed.

DATAMGT

specifies optional access methods (BTAM, ISAM, GAM, TCAM and VTAM) to be included in the system.

DATASET

specifies system data sets and the volumes on which they reside. It is required for each defined data set that is not located on the system residence volume. It is also used to specify user-written macros, routines, and procedures to be added to system libraries.

EDIT

specifies the physical characteristics and processing attributes of the data sets to be processed by the time sharing EDIT command.

GENERATE

specifies the data sets, volumes, and I/O devices required for the SYSGEN process, the SYSGEN output options, and the type of generation being performed.

IODEVICE

specifies characteristics of an I/O device and its operating system requirements. An IODEVICE macro instruction is required for each uniquely addressable I/O device.

SCHEDULR

specifies job and master scheduler options.

SVCTABLE

specifies the number, type, APF authorization, and entry status of the user-written SVC routines to be added to the new system.

TSO

specifies the inclusion of IBM-supplied time sharing command processors.

UNITNAME

specifies a name for a group of I/O devices and optionally makes the name available for VIO use. A UNITNAME macro instruction is required for each named group of I/O devices in the system, except for unique device types.

Job Entry Subsystem Generation

Job entry subsystems (JES2 or JES3) are added to the system through the process of job entry subsystem generation, which may be executed concurrently with Stage II. Figure 4 summarizes the steps that are involved in performing a job entry subsystem generation and relates them to the steps that are required to perform a system generation (SYSGEN).

Step	SYSGEN	JES2 Generation
Planning (selecting options)	x	x
Volume and data set initialization	x	x
Stage I	x	
Stage II (jobstream execution)	x	x
Stage II and JES2 generation may be executed concurrently.		
RMTGEN (optional)		x
Testing	x	x

Figure 4. Generating a VS2 Release 2 System

JES2 Generation

The three distribution libraries needed to perform a JES2 generation are distributed by IBM along with the VS2 Release 2 distribution libraries. They contain:

- The jobstream for performing the JES2 generation.
- The JES2 source code.
- The utilities used by the JES2 jobstream.

During JES2 generation, the JES2 control program is assembled to reflect the options and limits that the installation has chosen. The installation takes the following steps to perform a JES2 generation.

Planning

The JES2 and optional RMTGEN parameters must be chosen and coded.

Volume and data set initialization

The space required for JES2 data sets must be allocated on the new system volumes.

Jobstream Execution

The jobstream on the distribution library must be executed to apply options, assemble, link-edit, and copy the JES2 modules from the distribution library to the data sets that have been allocated on the new system volumes. The jobstream also invokes RMTGEN if any remote terminals were specified by the JES2 generation parameters.

JES2 Generation Options: The following characteristics of the JES2 control program may be specified:

- The number of card readers, printers, and card punches to be used by the job control manager.
- The maximum number of initiators that may be started at any one time.
- The default limits for a job's printed or punched output and the actions to be taken when the specified limits are exceeded.

- The JES2 job classes.
- The type of processing for SYSOUT data sets.
- The number of direct access storage volumes that may be mounted concurrently as spooling devices.

RMTGEN

In addition to generating a job entry subsystem, one or more remote station programs may also be generated. The process of generating job entry subsystem multileaving remote terminal programs is called RMTGEN. The parameters used during RMTGEN specify hardware configuration and software options. Remote terminal programs that can be generated by RMTGEN include:

- System/360 binary synchronous communication (BSC) remote terminal program.
- System/360 model 20 remote terminal program.
- 1130 remote terminal program.
- 1130 remote terminal program load program.
- System/3 remote terminal program.

Note: Remote terminals can be supported only by 2701 and 2703 emulators on the 3704 and 3705.

RMTGEN Options: The following remote terminal program characteristics may be specified:

- The type and degree of text compression.
- The amount of main storage available to the program.
- The number of buffers to be used by the program.
- The inclusion or exclusion of unit record devices.
- The unit address of the BSC adapter and devices.
- The speed of the communication line used.
- The number of buffer areas to be used for restoring compressed lines to their original form.
- The type of assembly listing to be produced during RMTGEN.

System Generation Changes

Some of the major changes to the system generation processes for MVT and VS2 Release 1 include:

- Changes to the process of initializing new system data sets.
- The reduction of the number of macro instructions that needs to be specified during Stage I (see Figure 5). This reduction is due both to the fact that previously optional facilities are now standard and to the consolidation of related options under single macro instructions in VS2 Release 2. For example, it is no longer necessary to specify two different SYSGEN macro instructions (RESMODS and LINKLIB) to add installation-written routines to the nucleus and the link library. The parameters on these macro instructions have been consolidated under the DATASET macro instruction which is now used to specify the addition of installation-written routines to any system library. (See the topic "Adding Installation-Written Members to System Libraries.")
- The addition of a job entry subsystem generation.
- The addition of parameters for specifying authorized appendages.

MVT	VS2 Release 1	VS2 Release 2
—	—	AFFINITY
ALGLIB	***	***
ALGOL	***	***
ASSEMBLR	***	***
CENPROCS	CENPROCS	—
CHANNEL	CHANNEL	CHANNEL
CHECKER	***	***
CKPTREST	CKPTREST	CKPTREST
CMDLIB	DATASET	DATASET
COBLIB	***	***
COBOL	***	***
CTRLPROG	CTRLPROG	CTRLPROG
DATAMGT	DATAMGT	DATAMGT
DCMLIB	DATASET	DATASET
EDIT	EDIT	EDIT
EDITOR	EDITOR	—
EMULATOR	—	—
FORTLIB	***	***
FORTRAN	***	***
GENERATE	GENERATE	GENERATE
GENTSO	—	—
GJOBCTL	—	—
GRAPHICS	GRAPHICS	DATAMGT
HELP	DATASET	DATASET
IMAGELIB	DATASET	DATASET
IOCONTRL	—	—
IODEVICE	IODEVICE	IODEVICE
LINKLIB	LINKLIB	DATASET
LOADER	LOADER	—
MACLIB	MACLIB	—
OUTPUT	TSO	TSO
PARMLIB	PARMLIB	DATASET
PL1	***	***
PL1LIB	***	***
PROCLIB	DATASET	DATASET
PTOP	***	***
RESMODS	RESMODS	DATASET
RPG	***	***
SCHEDULR	SCHEDULR	SCHEDULR
SECMODS	CENPROCS	—
SECONSLE	SECONSLE	CONSOLE
SORTLIB	***	***
SORTMERG	***	***
SUPRVSOR	—	—
SVCLIB	DATASET	DATASET
SVCTABLE	SVCTABLE	SVCTABLE
SYSUTILS	***	***
TELCMLIB	DATASET	DATASET
TSOPTION	TSO	TSO
UCS	UCS	—
UNITNAME	UNITNAME	UNITNAME
—	PAGE	DATASET
—	no longer specified during system generation	
***	the function specified by this macro instruction may be added to the system after system generation	

Figure 5. Changes in Stage I Macro Instructions

Initializing the New System Data Sets

New system data sets must be initialized before Stage II is executed. Initialization includes:

- allocating space for system data sets on the new system volumes.
- building the volume index of the system catalog.
- cataloging the data sets in the new system catalog.

In VS2 Release 2, the new system data sets can be initialized either during system generation by using the DATASET macro instruction or at any other time before Stage II by using the access method services.

The Stage I DATASET macro instruction did not exist in MFT or MVT. Its use for initializing new system data sets during SYSGEN remains unchanged from VS1 Release 1 and VS2 Release 1. A DATASET macro instruction is specified with the SPACE parameter in the Stage I input deck for every system data set that is to be allocated. The new system data sets are then allocated and cataloged during Stage II execution.

In VS2 Release 2, the IEHPROGM utility program can be used to catalog only non-VSAM data sets without CVOL pointers. In order to preallocate all new system data sets (VSAM and non-VSAM data sets) the access method services should be used. In this case a DATASET macro instruction should be specified without the SPACE parameter for any system data set used during system generation that is not on the system residence volume. This enables the catalog entry to be created during Stage II.

Adding Installation-Written Members to System Libraries

The DATASET SYSGEN macro instruction is used to include installation-written members in system libraries. The parameters:

```
PDS=SYS1.name  
MEMBERS=( name [, name] . . . )
```

are functionally equivalent to their counterparts on the RESMODS and LINKLIB SYSGEN macro instructions in MVT and VS2 Release 1. A new feature is provided for adding installation-written routines to the nucleus, however. If only the PDS parameter is specified, the members in the data set specified by PDS will be copied into the nucleus as individual members. If both the PDS and MEMBERS parameters are specified, the members in the data set specified by PDS will be copied into the nucleus and link-edited together with the nucleus member IEANUC01. Figure 6 indicates the system libraries to which installation-written members may be added and the form in which they must be added.

System Library	Form
SYS1.CMDLIB	load module
SYS1.IMAGELIB	load module
SYS1.LINKLIB	load module
SYS1.LPALIB	load module
SYS1.MACLIB	system macro
SYS1.NUCLEUS	load module
SYS1.PARMLIB	system parameter
SYS1.PROCLIB	system procedure
SYS1.SVCLIB	load module
SYS1.TELCMLIB	load module
SYS1.VTAMLIB	load module
SYS1.VTAMLST	system parameter

Figure 6. System Libraries to which Installation-Written Members May Be Added

Unsupported VS2 Release 1 Stage I Macro Instructions

The following VS2 Release 1 macro instructions are not supported in VS2 Release 2.

CENPROCS

Model dependent information is no longer needed during system generation processing.

EDITOR

All functions provided by this macro instruction are available in SYS1.PROCLIB cataloged procedures.

GRAPHICS

The DATAMGT macro instruction now performs the function.

LINKLIB

The DATASET macro instruction now performs the function.

LOADER

Options specified by this macro instruction are now available in SYS1.PROCLIB cataloged procedures.

MACLIB

The DATASET macro instruction now performs the function.

PAGE

The DATASET macro instruction now performs the function.

PARMLIB

The DATASET macro instruction now performs the function.

RESMODS

The DATASET macro instruction now performs the function.

SECONSLE

The CONSOLE macro instruction now performs the function.

UCS

The options specified by this macro instruction are standard in the VS2 Release 2 system.

System Initialization

Before any job can be processed, the control program, its associated control blocks and work areas must be loaded into storage and prepared for operation. Loading these control program modules is the function of the initial program loader (IPL).

After the IPL program completes its loading functions, control passes to the nucleus initialization program (NIP), which performs more functions necessary for the operation of the control program. (For example, NIP defines storage areas and initializes certain tables, work areas, and control blocks.) NIP also loads and initializes routines (such as fixed LPA routines) selected by the user. The nucleus initialization program also initializes real storage. It gives the installation control over virtual storage definition by processing installation-supplied system parameters that specify virtual storage options.

Master scheduler initialization completes the initialization process by establishing system tasks such as the console communications task, the log task, job entry subsystems, system management facilities, the recovery/termination management task and the missing interruption checker task.

In VS2 Release 2, changes have been made to the system initialization process to provide the installation with greater flexibility in specifying system options as well as the opportunity to simplify the IPL process. These changes include:

- The reduction of informational IPL messages.
- The elimination of the time-of-day (TOD) clock messages and associated replies, unless the installation requests that they be issued or unless there is an error in TOD clock processing.
- The implementation of multiple SYS1.PARMLIB members for SMF, volume attributes, commands, and the link library.
- The creation of a SYS1.PARMLIB member that contains installation-specified operator commands that are issued by the system control program at the completion of system initialization processing.

VS2 Release 2 System Parameters

The installation uses system parameters to control the system initialization process. System parameters may be specified through the master console or in the IEASYSxx lists in the system parameter library, SYS1.PARMLIB. Each parameter in a SYS1.PARMLIB list is specified in the same format as that required by the "SPECIFY SYSTEM PARAMETERS" message. Figure 7 provides a summary of these parameters. More detailed descriptions of the parameters follow.

Parameter	Value Specified		SYS1.PARMLIB List Read
APF	Authorized library name	⊙	IEAAPFxx
APG	Automatic priority group for system resources manager	•	
BLDL	Pageable directory for SYS1.LINKLIB		IEABLDxx
BLDLF	Nonpageable directory for SYS1.LINKLIB	•	IEABLDxx
CLPA	New link pack area to be created	•	IEALOD00 IEAPAK00
CMD	Command to be issued internally	⊙	COMMNDxx
CSA	Size of the common service area	⊙	
CVIO	Delete all VIO data sets from paging space	⊙	
DUMP	Data sets for SYS1.DUMP		
FIX	Reenterable routines for nonpageable LPA	•	IEAFIXxx
HARDCPY	Hard copy log		
IPS	Installation performance specification	⊙	IEAIPSxx
LNK	Names of data sets concatenated to SYS1.LINKLIB	⊙	LNKLISTxx
LOGCLS	Output class for log data set	⊙	
LOGLMT	WTL limit for log data set	⊙	
MAXUSER	Maximum number of virtual address spaces	⊙	
MLPA	Modifications to pageable LPA	•	IEALPAxx
NUCMAP	New nucleus map for DSS	⊙	
OPI	SYS1.PARMLIB operator intervention restrictions	•	
OPT	System resources manager tuning parameters	⊙	IEAOPTxx
PAGE	Page data set names	•	
REAL	V = R address area size	•	
SMF	SMF parameters	⊙	SMFPRMxx
SQA	Size of the system queue area		
SYSP	System parameter list to be merged with IEASYS00	•	IEASYSxx
VAL	Volume characteristics	⊙	VATLSTxx
VRREGN	Default region size for a V=R request	⊙	
WTOBFRS	Number of buffers for WTO routine use	⊙	
WTORPLY	Number of operator reply elements for WTOR routine use	⊙	
<ul style="list-style-type: none"> • — Not Included in MVT Systems ⊙ — Not Included in VS2 Release 1 Systems and MVT Systems 			

Figure 7. VS2 Release 2 System Initialization Parameter Summary

APF

specifies the value that is appended to IEAAPF to form the name of the SYS1.PARMLIB list that contains data set names with their corresponding volume identifications of authorized program libraries.

The SYS1.LINKLIB and SYS1.SVCLIB libraries are automatically included as authorized libraries.

APG

specifies an automatic priority group for use by the system resource management routines. If the APG parameter is not specified, a default value of 7 will be used.

BLDL

specifies a value that is appended to IEABLD to form the name of a SYS1.PARMLIB list that contains the names of modules for which entries are to be built in the resident BLDL table.

The directory created in response to the BLDL parameter will reside in pageable storage.

If the BLDL and BLDLF parameters are both specified, the BLDL parameter will be ignored. A warning message will be issued to identify this condition.

The IEABLD00 list is the default list.

BLDLF

specifies a value that is appended to IEABLD to form the name of a SYS1.PARMLIB list that contains the names of modules for which entries are to be built in the resident BLDL table.

The directory created in response to the BLDLF parameter will reside in fixed storage.

If the BLDLF and BLDL parameters are both specified, the BLDLF parameter will be accepted. A warning message will be issued to identify this condition.

The IEABLD00 list is the default list.

CLPA

specifies that a new pageable link pack area is to be created.

If a previously created link pack area exists, it is deleted and the space it occupied is made available for system paging use.

CMD

specifies a value that is appended to COMMND to form the name of the SYS1.PARMLIB list that contains the commands that the installation selects to be issued internally as soon as the system initialization process is complete.

If not specified in a system parameter list of by the operator, the COMMND00 list is used.

CSA

specifies the common service area.

The value specifies the size, in 1K byte multiples, of the common service area.

If CSA is not specified, the default value of 100K bytes is used.

CVIO

for system restart, specifies that all virtual I/O (VIO) data sets are to be deleted from the paging space.

DUMP

specifies tape volumes and/or the direct access storage devices to be used for the SYS1.DUMP data sets.

FIX

specifies a value that is appended to IEAFIX to form the name of the SYS1.PARMLIB list that contains the names of reenterable routines from the SYS1.LINKLIB, SYS1.LPALIB, and SYS1.SVCLIB data sets to be made resident as a fixed extension to the link pack area (fixed link pack area).

If not specified, a fixed link pack area is not created.

HARDCPY

specifies that the hard copy log is desired, whether or not the system log is to be used as the hardcopy log, and whether or not messages are to be recorded on the hardcopy log.

If the console configuration contains an active graphic console or more than one active console, a hardcopy log is required. In this case, if routing codes 1,2,3,4,7,8, or 10 are desired, they need not be specified since all of these codes are automatically assigned by the system control program.

IPS

specifies a value that is appended to IEAIPS to form the name of the SYS1.PARMLIB member that contains the installation performance specification (IPS).

The IEAIPS00 is the default list.

LNK

specifies a value that is appended to LNKLST to form the name of the SYS.PARMLIB list that contains the names of the data sets to be concatenated to SYS1.LINKLIB.

Any number of SYS1.PARMLIB lists may be specified in this manner until the number of data sets that they contain reaches 15. (The maximum number of data set names other than SYS1.LINKLIB is 15. This allows 16 data sets including SYS1.LINKLIB to be concatenated.) SYS1.LINKLIB is opened first. Then each data set that is specified is opened and concatenated in the order in which it appears (i.e., left to right).

The LNKLST00 list is the default list.

LOGCLS

specifies the JES2 output class for the log data sets.

If the LOGCLS parameter is not specified, a default class of A is used.

LOGLMT

specifies the number of write-to-log (WTL) instructions that may be issued before the log data set is scheduled for output processing.

If the LOGLMT parameter is not specified, a default value of 500 WTL instructions is used.

MAXUSER

specifies the maximum number of independent virtual address spaces available for use by the operating system.

If the MAXUSER parameter is not specified, a default value of 256 is used. The maximum number that may be specified is 1,536.

MLPA

specifies a value that is appended to IEALPA to form the name of the SYS1.PARMLIB list that contains the reenterable routines from SYS1.LINKLIB, SYS1.LPALIB, and SYS1.SVCLIB data sets to be made resident as a pageable extension to the link pack area (modified link pack area).

If the MLPA parameter is not specified, a modified link pack area is not created.

NUCMAP

specifies the creation of a new map of the nucleus in the SYS1.DSSVM data set and the deletion of the existing nucleus map.

The parameter must be specified if the nucleus for which a map exists is modified and an IPL is performed using the modified nucleus. The NUCMAP parameter need not be specified again until another modification is made to the nucleus.

A new nucleus map will be created without the specification of the NUCMAP parameter, if the name of the nucleus load module changes between IPL procedures or if a nucleus map does not exist for the first IPL after system generation.

OPI

specifies operator intervention restrictions of the system parameters in the SYS1.PARMLIB lists.

The OPI parameter can only be included in SYS1.PARMLIB lists. It may be specified for individual system parameters or for the entire set of system parameters.

OPT

specifies a value that is appended to IEAOPT to form the name of the SYS1.PARMLIB list that contains tuning parameters for the system resources management routines.

If not specified in a system parameter list or by the operator, default values provided by the system resources management routines are used.

PAGE

specifies the names for the page data sets.

The values specified include identification of the VSAM data sets for which space was previously allocated and formatted (and, optionally, the unit on which the volume containing the specified data sets should reside).

REAL

specifies nonpageable storage (virtual equals real address area).

The value specified is the number of 1K byte blocks to be reserved for nonpageable (V=R) storage.

If the value specified is 0, there will be no V=R storage in the system.

If the REAL parameter is not specified, a default of 76K will be used.

SMF

specifies a value that is appended to SMFPRM to form the name of the SYS1.PARMLIB list that contains system management facility (SMF) values.

If not specified in a system parameter list or by the operator, the SMFPRM00 list is used.

SQA

specifies the system queue area.

The value specified is the number of 64K byte segments to be added to the minimum size of the system queue area.

If not specified in a system parameter list or by the operator, a default value of one (64K byte segment) is added to the minimum size of the system queue area.

If the specified amount of virtual storage is not available, a warning message will be issued. The SQA parameter may be respecified at that time.

The size of the system queue area cannot be increased during a restart that reuses the previously initialized link pack area. If the SQA parameter is specified on the restart, and the size is less than that specified on the previous system start, a warning message will be issued and the size will be increased to that of the previous system start. If the SQA parameter is specified on the restart, and the size is greater than that specified on the previous system start, a message will be issued and the size for the previous system start will be used.

SYSP

specifies the SYS1.PARMLIB list of system parameters that is to be merged with the default list of system parameters, IEASYS00.

The value specified is appended to IEASYS to form the name of a SYS1.PARMLIB list that contains the system parameters to be used for the current IPL. Any number of IEASYSxx lists may be specified.

If not modified by this parameter, only the IEASYS00 list is used. The IEASYS00 list contains the values selected during system generation for various system parameters unless the contents of the list have been modified by the system programmer at the installation.

When multiple IEASYSxx lists are specified by the SYSP parameter, they are treated in ascending priority sequence (i.e., parameters defined in a member specified near the beginning of the list will be replaced by the same parameters defined in a member specified near the end of the list). In the process of merging parameters with the IEASYS00 list, IEASYS00 is always read, and always assumes the lowest priority.

For example, if the operator responds:

```
R 00, 'SYSP=(01,02)'
```

and these two IEASYSxx lists contain the following system parameters:

```
IEASYS01: MLPA=(00,01),BLDL=00
```

```
IEASYS02: MLPA=02,SQA=10
```

then the effective system parameters will be:

```
MLPA=02, BLDL=00, and SQA=10
```

VAL

specifies a value that is appended to VATLST to form the name of the SYS1.PARMLIB list that contains mount and allocation characteristics for any or all of the direct access volumes used by the installation.

If not specified in a system parameter list or by the operator, the VATLST00 list is used.

VRREGN

specifies nonpageable storage (virtual equals real address area) for a specific job.

The value specified becomes the default value for the real main storage size (in 1K multiptes) to be allocated to a job requiring V=R address space.

If the VRREGN parameter is not specified, a default value of 64K is used.

The VRREGN system parameter should not be confused with the REAL system parameter. REAL specifies the total amount of nonpageable (V=R) storage available to the system. VRREGN specifies the subset of the total space that is available for an individual job.

WTOBFRS

specifies WTO buffers.

The value specifies the number of buffers to be used by the write-to-operator (WTO) routines.

If the WTOBFRS parameter is not specified, a default value of 20 is used.

WTORPLY

specifies WTOR elements.

The value specifies the number of operator reply elements to be used by the write-to-operator-with-reply (WTOR) routines.

If the WTORPLY parameter is not specified, a default value of 5 is used.

Changes to System Parameters for VS2 Release 2

Figure 7 shows new system parameters for VS2 Release 2. Some VS2 Release 1 and some MVT system parameters are not supported in VS2 Release 2. Of these, some have been replaced by new parameters and others have been eliminated, either because their functions are not supported in VS2 Release 2 or because their functions are performed by the system control program in VS2 Release 2.

Changes to VS2 Release 1 System Parameters: The following VS2 Release 1 system parameters have changed for VS2 Release 2.

APG

The only value that can be specified for the APG parameter in VS2 Release 2 is the priority of the automatic priority group.

DUMP

Ten SYS1.DUMP data sets can be specified in VS2 Release 2.

PAGE

Values specified by the PAGE parameter have changed for VS2 Release 2 because paging space is made up of preformatted VSAM data sets, two of which must be defined and formatted during system generation, (See the topic "Adding Paging Space to the System" in this chapter.) The PAGE parameter also no longer specifies a particular volume to contain the link pack area.

REAL

In VS2 Release 2, the value specified by the REAL parameter represents the total amount of V=R storage in 1K byte multiples, as opposed to the number of 4K byte blocks added to a 64K byte minimum value in VS2 Release 1.

Unsupported VS2 Release 1 System Parameters: The following VS2 Release 1 system parameters are not supported in VS2 Release 2 but are provided for by comparable VS2 Release 2 functions.

LSQACEL

This parameter is not needed in VS2 Release 2 because quickcells for the local system queue area are controlled by virtual storage management routines.

PAL

The values specified by this parameter are supplied by the address space supervision routines in VS2 Release 2.

SQACEL

This parameter is not needed because quickcells for the system queue area are controlled by virtual storage management routines.

TMSL

This parameter is not needed because the system resources management routines determine the acceptable time slices.

TRACE

The system trace option cannot be specified in VS2 Release 2. It is activated by the system control program for the duration of system initialization processing, after which generalized trace facility (GTF) processing may be started.

The following VS2 Release 1 system parameters are not supported in VS2 Release 2 and have no comparable VS2 Release 2 functions.

AUXLIST

This parameter is not needed because in VS2 Release 2 it is not meaningful to distinguish between the auxiliary storage available for background and time-sharing jobs.

CPQE

The number of channel programs to be used by the address space supervision routines is no longer an installation option.

MPA

It is not necessary to specify the size of the master scheduler region because in VS2 Release 2, the master scheduler task has its own independent virtual address space.

TSOAUX

It is not necessary to specify back-up auxiliary storage space for time sharing jobs in VS2 Release 2, because each time sharing user is assigned to an individual virtual address space.

Unsupported MVT System Parameters: The following MVT system parameters are not supported in VS2 Release 2, but are provided for by comparable VS2 Release 2 functions.

MIN

This parameter is not needed in VS2 because the initiator modules are contained in the pageable link pack area.

MOD

The CPU model identification is determined in VS2 by the STIDP instruction.

RAM

Routines from SYS1.SVCLIB and SYS1.LINKLIB can be added to the VS2 link pack area by means of the SYS1.LPALIB data set or the MLPA or FIX system parameters.

RERP

In VS2, error recovery procedure (ERP) routines are always resident in the link pack area.

RSVC

In VS2, SVC routines are always resident in the link pack area.

SQS

This parameter is replaced by the SQA system parameter.

TMSL

This parameter is not needed because the system resource management routines determine the acceptable time slices.

The following MVT system parameters are not supported in VS2 Release 2 and have no comparable VS2 Release 2 functions.

ALTSYS

This parameter is not needed because the system residence devices are treated by dynamic device reconfiguration (DDR) as any other swappable device.

HRAM

MVT hierarchy support is not needed in VS2.

HSVC

MVT hierarchy support is not needed in VS2.

MPS

It is not necessary to specify the size of the master scheduler region because in VS2 Release 2, the master scheduler task has its own independent virtual address space.

QBF

This parameter is not needed in VS2 Release 2 because the job queue is replaced by the scheduler work area (SWA) and job entry subsystem data sets.

System Parameter Specification

An installation can tailor a system through the specification of system parameters, either during system generation (SYSGEN) processing or in the IEASYS00 or IEASYSxx lists of SYS1.PARMLIB. The system operator specifies system parameters through the master console in response to the "SPECIFY SYSTEM PARAMETERS" message.

The IEASYS00 (system parameter) list is always created during the SYSGEN process, and the system parameters HARDCPY, PAGE and either BLDL or BLDLF are included in it. The installation may also specify that the system parameters CSA, REAL, SQA, VRREGN, APF=00, and FIX=01 should be included in the IEASYS00 list during SYSGEN processing.

The remainder of the system parameters are specified in two ways:

- By using a system utility program to add them to the IEASYS00 or any IEASYSxx list.
- By entering them through the operator's console.

The IEASYS00 list is read for every IPL. The installation is not only able to add parameters to this list, but may change parameter values that were included in it during SYSGEN processing. Operator interaction will be minimized during system initialization processing if the installation ensures that the IEASYS00 list contains all of the system parameters required to establish the desired operating environment. The system operator will then be able to respond to the "SPECIFY SYSTEM PARAMETERS" message with end-of-block (EOB), and will not be required to respond further unless error conditions are detected and prompting occurs.

The IEASYS00 member should generally contain installation default values, or parameters for which values can be defined that do not change from IPL to IPL. The IEASYSxx members, on the other hand, should be used to contain parameters that are subject to change from IPL to IPL, depending on the user applications. (See the SYSP parameter description for additional information.)

System Parameter Library (SYS1.PARMLIB)

The SYS1.PARMLIB data set contains both IBM-created and installation-created lists of parameter values. The formats for the lists have been revised from MVT to allow better space utilization within parameter records, and to afford more flexibility in their definition. All MVT and VS2 Release 1 parameter lists that are still valid lists in VS2 Release 2 need not have their formats changed.

Changes to SYS1.PARMLIB Lists

Figure 8 shows the new or changed VS2 Release 2 SYS1.PARMLIB lists.

List	Change from MVT	Change from VS2 Release 1
COMMNDxx	New	New
IEAABD00	New	New
IEAAPFxx	New	New
IEAAPP00	New	New
IEABLDxx	None	None
IEADMP00	New	New
IEAFIXxx	New	None
IEAIPSxx	New	New
IEALOD00	New	None
IEALPAxx	New	None
IEAOPTxx	New	New
IEAPAK00	New	None
IEASYSxx	New	None
IQAORDER	Not Applicable	No longer-supported
IRBMF1xx	New	New
LNKLSTxx	No longer a single member	No longer a single member
PARMTZ	New	None
SMFPRMxx	No longer a single member named SMFDEFLT	No longer a single member named SMFDEFLT
VATLSTxx	No longer a single member named PRESRES	No longer a single member named PRESRES

Figure 8. VS2 Release 2 Changes to SYS1.PARMLIB Lists

IBM-Created Lists

The IBM-created lists that are named in column 1 are standard in SYS1.PARMLIB for VS2 Release 2.

Standard List	Alternate List	Contents
IEAABD00		Address of areas to be dumped by SYSABEND
IEABLD00	IEABLDxx	Names of SYS1.LINKLIB load modules whose directory entries are to be made in the BLDL table
IEADMP00		Addresses of areas to be dumped by SYSUDUMP
IEAIPS00	IEAIPSxx	Installation performance specification
IEALOD00		Names of SYS1.LPALIB modules whose directory control block information is to be maintained in nonpageable storage
IEAPAK00		Groups of names of SYS1.LPALIB modules that are to be packed together in the link pack area
IEASYS00	IEASYSxx	Lists of system parameters
IRBMF100	IRBMF1xx	MF/1 parameters
LNKLST00	LNKLSTxx	Names of additional data sets that may be concatenated with SYS1.LINKLIB data sets
SMFPRM00	SMFPRMxx	SMF parameters

The following lists are optionally created during the system generation process if specified by the installation.

Standard List	Alternate List	Contents
IEAAPF00	IEAAPFxx	Names of authorized libraries
IEAAPP00		Names of installation-written EXCP appendages
IEAFIX01	IEAFIXxx	Names of SYS1.SVCLIB, SYS1.LPALIB, and SYS1.LINKLIB modules that are to be made resident as a nonpageable extension to the link pack area

Of the IBM-created lists, multiple lists can be created and maintained for those that have alternate names. These lists may therefore be respecified by the installation to continue the system tailoring begun with system generation processing. The desired list may then be selected from the multiple lists by system parameters specified from the operator's console or in the IEASYS00 or IEASYSxx lists.

Authorized Appendage List (IEAAPP00)

This list contains the names of installation-written I/O appendage routines and a description of the function of each routine. Any I/O appendage that is to be used by an unauthorized routine must be named in the IEAAPP00 member of SYS1.PARMLIB. (Authorized routines are able to use any installation-written I/O appendage.) This system integrity requirement prevents unauthorized routines from incorrectly using I/O appendages.

It is not necessary to include the names of control program (e.g., system access methods) appendages in an authorized appendage list.

If the installation chooses to specify the creation of IEAAPP00 during system generation, appendage names may be entered, even if the appendages have not yet been written. This capability allows the installation to "prime" the list with the names of appendages that will be added to the system later.

See the chapter "System Integrity" for a description of authorized and unauthorized routines.

Authorized Program Library List (IEAAPFxx)

This list contains the data set names and corresponding volume indentifications for libraries that may contain programs requiring APF authorization.

The ability for installation-specified libraries to contain APF-authorized programs extends the VS2 Release 1 APF support that allowed authorized programs to reside only in SYS1.LINKLIB, SYS1.SVCLIB, and SYS1.LPALIB. The specification of additional libraries for authorized programs in the IEAAPP00 or IEAAPFxx list in VS2 Release 2 eliminates the need to concatenate libraries containing programs requiring APF-authorization.

It is not necessary to include SYS1.LINKLIB or SYS1.SVCLIB in an authorized program library list.

Dump Default Lists (IEAABD00 and IEADMP00)

The dump default lists, IEAABD00 and IEADMP00, contain the IBM-supplied defaults for the storage areas to be dumped to the data sets defined by SYSABEND and SYSUDUMP DD statements, respectively. The installation may change the defaults in these lists to satisfy individual requirements. Some of the defaults that may be specified include the system queue area, the local system queue area, the scheduler work area for the task, the task's control blocks, the GTF or supervisor trace table, the PSW at entry to ABEND, general register contents at entry to ABEND, the contents of the task's link pack area, the contents of the task's job pack area.

Fixed LPA List (IEAFIXxx)

The fixed LPA list indicates those modules from SYS1.SVCLIB, SYS1.LPALIB, and/or SYS1.LINKLIB that should be included in the nonpageable extension to the link pack area. The list should contain modules that significantly increase system performance when they are fixed rather than paged.

The fixed LPA list is created during system generation processing, if so specified by the installation.

Installation Performance Specification (IEAIPSxx)

The installation performance specification (IPS) consists of control information used by the workload management routines that are part of the system resources management routines. It includes performance group definitions, performance objectives, and service definition coefficients that are used to establish service rates.

For a more detailed description of the contents of the IPS, see the chapter: "Directing the Use of System Resources".

Link Library List (LNKLSTxx)

The link library list enables concatenating up to 15 data sets, on multiple volumes, with SYS1.LINKLIB. It is generated during the system generation process but may be modified via the IEBUPDTE utility program. Each data set that is concatenated with SYS1.LINKLIB may have up to 16 extents. (Note: After additions or changes are made that extend data sets concatenated with SYS1.LINKLIB, the system should be re-IPLed to update the description of the SYS1.LINKLIB to the system.)

In VS2 Release 2, multiple link library lists may be maintained in SYS1.PARMLIB. The list to be used is selected at IPL time by the LNK system parameter.

LPA Directory Load List (IEALOD00)

The link pack area (LPA) directory load list indicates those modules residing on SYS1.LPALIB or in the pageable LPA that are to have the contents supervision directory control block information initialized by NIP in nonpageable storage. When this list is used, contents supervision routines do not have to search to find the module in the LPA directory. The number of page faults that may occur is therefore reduced. Candidates for this list are modules that are frequently used.

LPA Packing List (IEAPAK00)

The link pack area (LPA) packing list contains the names of modules residing on the SYS1.LPALIB data set that may be grouped together by packing the LPA. Grouping these modules may sharply reduce page faults as well as eliminate wasted space within the LPA.

Modules that frequently pass control to one another should be grouped together to reduce page faults.

The total size of all modules in a group should not exceed 4K (page size).

The proper loading of modules into the LPA is done at NIP time, based upon the data in IEAPAK00. Modules that are not listed in IEAPAK00 are then loaded into the LPA in the order that they appear in the SYS1.LPALIB directory.

There are no alternate lists for IEAPAK00. If IEAPAK00 is not found as a member of SYS1.PARMLIB, the modules will be loaded in the order that they appear in the SYS1.LPALIB directory.

MF/1 Parameter List (IRBMF1xx)

This list contains the parameters that control MF/1 operation. The options that may be specified are described in the topic: "MF/1 Operation" in the chapter: "Directing the Use of System Resources".

Resident BLDL List (IEABLDxx)

The resident BLDL list specifies the names of modules from SYS1.LINKLIB, or any data set concatenated to SYS1.LINKLIB, that are to have directory entries built by NIP. This eliminates the directory search required when a load module is requested from SYS1.LINKLIB.

The directory will reside either in the fixed or pageable area of the system, depending upon whether the BLDLF or BLDL reply is specified by the operator

in response to the "SPECIFY SYSTEM PARAMETERS" message. The BLDL and BLDLF parameters are mutually exclusive.

The member name for the standard list is IEABLD00. The load module names must be listed in the same order as they appear in the directory; that is, they must appear in ascending collating sequence. (Note: Directory entries in the resident BLDL table are not updated as a result of updating the load module in the library. The old version of the load module is used until an IPL process updates the resident BLDL table.)

The IEABLD00 list indicated by BLDL=00 or BLDLF=00 is created during system generation processing. The list will be null if no entries are specified for it at SYSGEN.

If specified during system generation and not modified through operator communication or a list of system parameters (IEASYSxx) the default IEABLD00 is used.

SMF Parameter List (SMFPRMxx)

The SMF parameter list contains the parameters that control SMF operations. Multiple SMF parameter lists may be maintained in SYS1.PARMLIB. The list to be used is selected at IPL time by the SMF system parameter.

The single SYS1.PARMLIB member (SMFDEFLT) supplied by IBM in MVT and VS2 Release 1 is not valid. In VS2 Release 2, the IBM-supplied SYS1.PARMLIB member that contains the default SMF parameters is called SMFPRM00. The format of the SMFPRM00 and SMFPRMxx lists is different from the MVT and VS2 Release 1 SMFDEFLT list. The ALT and PRM keywords are not supported because the SMF data sets, SYS1.MANX and SYS1.MANY must be cataloged in VS2 Release 2. The MDL keyword has also been deleted. The SID keyword has been expanded to allow the specification of a System/370 model number.

System Parameter List (IEASYSxx)

The system parameter list may include all system parameters that are valid input to the "SPECIFY SYSTEM PARAMETERS" message. The standard list, IEASYS00, is created during Stage II of the system generation process. This parameter list enables an installation to specify all parameters needed during a particular IPL process without requiring that they be specified from an operator's console.

Operator intervention of an individual parameter list may be restricted by use of the 'OPI=YES/NO' parameter.

Installation-Created Lists

The installation is also able to create certain lists, that are not supplied by IBM, to further tailor the system to installation needs. Most of these may also be maintained as multiple lists. The following lists may be created by the installation in VS2 Release 2:

Default List	Alternate List	Contents
COMMND00	COMMNDxx	Commands to be issued by the system control program at the completion of system initialization processing.
IEALPA00	IEALPAxx	Names of SYS1.SVCLIB, SYS1.LPALIB, and SYS1.LINKLIB modules that are to be made resident as a pageable extension to the link pack area.
IEAOPT00	IEAOPTxx	System resources management tuning parameters.
PARMTZ		Local time constant
VATLST00	VATLSTxx	Volume attributes

Command List (COMMND00)

This list contains commands (e.g., START) that are to be issued by the system control program at the completion of system initialization processing. It may also be used to specify whether or not the time-of-day (TOD) clock initialization messages and associated responses should be issued.

The COMMNDxx list is specified by the CMD system parameter.

Local Time Constant (PARMTZ)

This SYS1.PARMLIB member may be used to contain the value (in hours, minutes, and seconds) that the local time differs from Greenwich Mean Time (GMT), and the direction (east or west) from the GMT.

This constant is used to set the time-of-day clock and overrides the local time value that is specified at system generation for inclusion in the CVT.

Modified Link Pack Area List (IEALPAxx)

This list may be used to indicate reenterable routines from the SYS1.LINKLIB, SYS1.LPALIB, and SYS1.SVCLIB data sets that are to be contained in the pageable extension to the link pack area. The routines in this area (called the modified link pack area) are resident only for the current IPL. If a routine is searched for, the modified link pack area is searched first, followed by the pageable link pack area. If a system restart is initiated and a previously created link pack area is used, the routines in the modified link pack area must be respecified if desired.

The IEALPAxx list is specified by the MLPA system parameter.

System Resources Management Tuning Parameter List (IEAOPTxx)

This parameter list contains the values that are used by the system resources management algorithms to influence when jobs should be swapped in or out of

real storage. The values are described in the chapter: "Directing the Use of System Resources".

Volume Attribute List (VATLSTxx)

The volume attribute list may be used to define the mount characteristics (permanently resident, reserved) and allocation characteristics (storage, public, private) for any, or all, direct access volumes used by an installation.

This list was a single SYS1.PARMLIB member called PRESRES in MVT and VS2 Release 1. Multiple members of the volume attribute list may be maintained in VS2 Release 2.

The VATLSTxx list is specified by the VAL system parameter.

System Libraries and Data Sets

Some system data sets are basic to VS2 Release 2 and are required for every operating system. Other system data sets are optional and are required only when selected options are included in the system.

Required Libraries and Data Sets

The following libraries and data sets are required for every VS2 Release 2 operating system and must be on the system residence volume.

- **SYS1.LOGREC** (error recording data set) -- The machine error recording data set is used by recovery management routines to record statistical data about machine errors and software errors.
- **SYS1.NUCLEUS** (nucleus library) -- The nucleus library contains the resident portion (nucleus) of the control program and may contain modules selected and link-edited during system generation.
- **SYS1.SVCLIB** (SVC library) -- The SVC library contains some OLTEP modules and some appendage modules.

The following libraries and data sets are required for every VS2 Release 2 operating system and must be on direct access volumes. It is not necessary that they be on the system residence volume.

- **Job Entry Subsystem Data Set** -- This data set contains spooled and checkpointed records. It may be a multivolume data set.
- **Master Catalog** -- The master catalog is a key-sequenced VSAM data set that contains pointers to cataloged data sets and user catalogs. The name can be specified at system generation time on the PDS= parameter of the Stage I DATASET macro instruction. This catalog replaces the MVT and VS2 Release 1 system catalog (SYSCTLG).
- **Page Data Sets** -- The page data sets are known collectively as paging space. Paging space contains the “paged out” portions of address spaces and the common system area, as well as the data written to virtual I/O data sets.
- **SYS1.DSSVM** -- This data set contains the command language modules used by the dynamic support system (DSS).
- **SYS1.LINKLIB** (link library) -- The link library contains programs and routines that are not in the link pack area and are referred to by ATTACH, LINK, LOAD, or XCTL macro instructions.
- **SYS1.LPALIB** (link pack area library) -- The link pack area library contains all modules that are loaded into the pageable link pack area. In VS2 Release 2, the pageable link pack area is a read-only area. Therefore, modules contained in it must be refreshable. Modules that are modified during execution (non-refreshable) should be included in the modified link pack area. Because all of the modules in this library are brought into the LPA at IPL time, this library may be placed on a demountable volume that may be removed after IPL.
- **SYS1.MACLIB** (macro library) -- This library contains macro definitions for the system macro instructions used by the assembler-language processor.
- **SYS1.PARMLIB** (parameter library) -- The parameter library contains default parameters for the initialization of system functions.

- **SYS1.PROCLIB** (procedure library) -- The procedure library contains cataloged procedures used for job definition.
- **SYS1.SAMPLIB** -- This library contains independent utility programs, the IPL text, and the SMF exit routines.
- **SYS1.STGINDEX** -- This VSAM data set is used as an internal catalog for virtual I/O data sets that must be saved across job steps and IPLs.

The following libraries and data sets are required if the TSO Stage I system generation macro instruction is coded or if installation-written time-sharing command processing routines are to be used. They must be on direct access volumes and may be on the system residence volume.

- **SYS1.BROADCAST** (broadcast data set) -- the broadcast data set contains two types of time-sharing messages: notices and mail.
- **SYS1.CMDLIB** (command library) -- This library contains time-sharing command processor programs, service routines, and utility programs.
- **SYS1.UADS** (user attribute data set) -- This data set contains a list of terminal users who are authorized to use time-sharing and information about the users.

Optional Libraries and Data Sets

The following data set is optional and, if selected, must be on the system residence volume.

- **PASSWORD** -- This data set contains records that associate the names of protected non-VSAM data sets with the passwords assigned to the data sets.

The following libraries and data sets are optional and, if selected, must be on direct access volumes unless otherwise stated.

- **SYS1.DCMLIB** -- This library contains program function key (PFK) definitions for display consoles and is required if a display console with the optional PFK feature is included in the system.
- **SYS1.DUMP** (dump data set) -- This data set contains core image dumps recorded in the event of abnormal termination of a critical task. The dump data set may reside on tape.
- **SYS1.HELP** -- This data set is required if the time-sharing **HELP** command is used. Each member of this data set contains information regarding the syntax, subcommands, and functions for each time-sharing command.
- **SYS1.IMAGELIB** -- This data set contains the 1403 and 3211 universal character set (UCS) and forms control buffers (FCB) image modules. It is required if a 1403 printer with the UCS feature, 3211 printer, 3890 document processor device, or 3886 optical character recognition device is included in the system.
- **SYS1.MANX/SYS1.MANY** (SMF data sets) -- The system management facility (SMF) data sets contain the information collected by the SMF routines. They may also contain measurement information collected by the system activity measurement facility (MF/1). Although the SMF data sets are optional they must be cataloged.
- **SYS1.TELCMLIB** (telecommunications library) -- This library contains the telecommunications subroutines in load module format and is required if any

of the TP access methods other than the virtual telecommunications access method (VTAM) are selected.

- SYS1.VTAMLST -- This data set contains lists of internal parameters in card image format used during VTAM initialization.
- SYS1.VTAMLIB -- This library contains the VTAM executable code in load module form.
- SYS1.VTAMOBJ -- This data set contains the 3705 network control program (NCP) that is loaded by VTAM during VTAM initialization.

Changes to VS2 Release 1 System Libraries and Data Sets

System data sets and libraries in VS2 Release 2 are the same as their counterparts in VS2 Release 1 with the following exceptions:

- The page data set, SYS1.PAGE, has been replaced by two or more uniquely named VSAM data sets. This paging space must be defined before the first IPL. (See the topic "Adding Paging Space to the System.") The page data sets are then named by means of the PAGE parameter during system initialization.
- The system catalog, SYSCTLG, has been replaced by a key-sequenced VSAM data set. Conversion aids are discussed in the topic "Catalog Conversion" in the chapter "Conversion Considerations".
- The data sets that support job entry subsystems are new.
- The format of the SYS1.BROADCAST data set has been modified to include a change level indicator in the header record and a new field for each user identification in the mail directory records. This field contains the address of the last message in the chain of messages associated with each user identification. A VS2 Release 2 format for the broadcast data set is obtained by using a new subcommand (SYNC) of the time-sharing ACCOUNT command. (See the topic "Formatting the Broadcast Data Set" in the "Conversion Considerations" chapter.)
- The load module size of modules in the SYS1.DSSVM data set has been increased from 1K bytes to 2K bytes.
- The SYS1.SYSJOBQE data set does not exist. It has been replaced by a scheduler work area (SWA) in virtual storage for initiator records and by a job entry subsystem data set for input, queued jobs, and output.
- The content of the SYS1.LOGREC data set has been expanded to include records for software error recording, missing interrupts, and dynamic device reconfiguration(DDR) routines.
- The SMF data sets, SYS1.MANX and SYS1.MANY, may now be used for system activity measurement facility (MF/1) recording. They must be cataloged in VS2 Release 2.
- The log data sets are no longer named SYS1.SYSVLOGX and SYS1.SYSVLOGY by IBM. They are dynamically allocated by the job entry subsystem.
- The SYS1.PARMLIB library contains new members to control the initialization of the system. See the topic "SYS1.PARMLIB" in the section "Defining the System" for descriptions of these members.
- The SYS1.STGINDEX data set is new.

- The format of the SYS1.UADS data set has been modified to include entries for authorized job entry subsystem users as well as time-sharing users, and to contain information for new keywords on ACCOUNT subcommands. An existing user attributes data set (UADS) is converted to the new VS2 Release 2 format by using an Account Facility utility program (UADSREFM). (See the topic "Reformatting the UADS" in the "Conversion Considerations" chapter.)
- The SYS1.VTAMLIB, SYS1.VTAMLST and SYS1.VTAMOBJ data sets are new.

Changes to MVT System Libraries and Data Sets

System data sets and libraries in VS2 Release 2 are also the same as their counterparts in MVT, with the exceptions previously noted and the following:

- The SYS1.ACCT data set no longer exists. The SMF data sets SYS1.MANX and SYS1.MANY should be used in its place.
- The SYS1.DSSVM data set is new.
- The SYS1.LPALIB data set is new. It contains all of the SVCs and most of the modules that, in MVT were contained in the SYS1.LINKLIB and SYS1.SVCLIB libraries. All of the modules contained in this new data set appear in the fixed or pageable link pack area in VS2.
- The page data sets are new.
- The SMF data sets, SYS1.MANX and SYS1.MANY, may now be used for VSAM recording.

Adding Paging Space to the System

Before the first system IPL, the installation must define and format paging space. Two or more VSAM data sets on two or more separate devices should be used for the page data sets. Two data sets are required because two copies of certain critical pages are maintained on external storage. During system generation, the installation defines and formats page data sets in the same manner as other system data sets. (See the topic "Initializing New System Data Sets").

The installation may include more paging space than the two required page data sets, by specifying additional page data sets during system generation.

If it is necessary to add paging space to the system after system generation processing is complete, the installation must first use the access method services to define and format the page data sets, and then re-IPL the system and issue the PAGE system parameter to add the data sets to the existing paging space.

No more than 25 page data sets can be specified during system generation. The number of page data sets must never exceed 32.

Page data sets can be added to the system in even or odd numbers; there is no need to "pair" them.

Directing the Use of System Resources

This chapter describes some of the VS2 Release 2 facilities that allow installations to direct the use of system resources.

System resources management routines allow an installation to specify, in measurable terms, the relative rate at which system resources should be provided to various categories of background and terminal jobs.

The system activity measurement facility (MF/1) provides reports on selected areas of system activity. The installation uses the information in these reports to determine how to specify performance objectives for the system resources management routines.

System Resources Management

System resources management routines are new in VS2 Release 2. They provide:

- The capability for an installation to specify, in measurable terms, the rate at which system resources should be provided to various users of the system.
- A means by which the system control program is able to monitor the use of system resources and attempt to maintain their most efficient use at all system workload levels.

These routines monitor the use of system resources to determine when background and time sharing jobs should be swapped in or out to meet the installation-specified objectives.

The resources management routines are divided into three primary categories:

- Workload management routines.
- Resource-use routines.
- A control function.

Workload Management Routines

The function of the workload management routines is to ensure that the relative rate at which processing resources are provided to individual jobs or time sharing users corresponds with installation performance specifications. Their design enables installations to specify a certain service level for any member of a subset of users under any level of system workload, without the necessity of relying solely on the operator's job mix selection. Workload management routines control the service rate (rate of use of system resources) of a background or terminal job by causing address spaces to be swapped into or out of real storage.

Resource-Use Routines

The resource-use routines consist of a set of algorithms that provide recommendation values for managing system resources on an absolute basis without regard to the installation's job performance specifications. These algorithms include:

I/O Load Balancing

This algorithm attempts to maintain a dispatchable job mix that has a balanced use of logical channels. The algorithm monitors the I/O load and produces recommendation values that indicate when swapping is necessary to correct a detected I/O subsystem imbalance. The I/O load balancing algorithm is used only when SMF data set activity recording is being performed.

CPU Load Balancing

This algorithm produces recommendation values for maintaining a job mix in which jobs that make heavy use of the CPU are dispatched at reasonable intervals while keeping CPU utilization at a high level.

Main Storage Occupancy

This algorithm determines the amount of swapping necessary to maintain real storage occupancy within certain bounds. It also determines whether or not the bounds are adequate, and adjusts them if they are not.

Page Replacement

This algorithm determines candidates for page stealing. Page stealing takes place when a certain page within a user's private address space has gone unreferenced by that address space for a certain amount of execution time, and when pages within the pageable link pack area have gone unreferenced for a certain amount of time. It results in the page frames associated with that page being freed for other users.

Device Allocation

This algorithm determines the unit to be used for data sets for which an allocation choice must be made from a list of more than one device candidate. The goal is to allocate from lightly-utilized channels, and to collect allocations for the same address space on the same logical channel without collecting them on the same device, so that swapping that address space will have a predictable effect on the system I/O load.

Automatic Priority Group (APG)

This algorithm reorders dispatchable address spaces within the automatic priority group based on the degree to which they are I/O bound; within the APG, I/O-bound jobs are given higher dispatching priorities than CPU-bound jobs.

To fully realize the benefits of this algorithm, the installation should place all possible users within the APG. In VS2 Release 2, this is facilitated because the APG priority is used as the default step dispatching priority. (See the topic "Changes to VS2 Release 1 JCL" in the chapter "Conversion Considerations.")

ENQ/DEQ Algorithm

This algorithm determines when a batch or terminal job is in control of a system resource and is delaying other jobs that are enqueued on the same resource. When such a determination is made, the job that is enqueued upon the system resource is made non-swappable for an installation-specified time interval. The interval is specified by the ERV parameter in the IEAOPTxx member of SYS1.PARMLIB.

Control Function

The control function gathers the results (recommendation values for potential swap decisions) produced by the workload management routines and the resource-use algorithms. It uses a control algorithm to combine and evaluate the individual recommendation values to produce a composite recommendation value for each potential swap decision. The control function then uses these values to determine the amount of swapping that should be performed in an attempt to satisfy the installation-specified performance objectives and efficiently use system resources.

The control algorithm uses optional installation-specified parameters called resource factor coefficients to weight the recommendation values produced by the I/O and CPU load balancing algorithms. The use of these coefficients (which are contained in the IEAOPTxx member of SYS1.PARMLIB) provides one method of system tuning. For example, if the installation increases the CPU coefficient, the recommendation value produced by the CPU load balancing algorithm will be more heavily weighted when it is evaluated by the control algorithm in conjunction with other recommendation values.

The Installation Performance Specification (IPS) and the Service Concept

In VS2 Release 2, an installation can specify, in measureable terms, the performance that a member of any subset of its users is to receive, under any system workload conditions, during any period in the life of a background or terminal job. The installation accomplishes this specification by setting up an installation performance specification (IPS) in SYS1.PARMLIB. (Optionally, the installation can define more than one IPS, although only one can ever be used at a given time). In the IPS, an installation defines:

- Performance groups -- subsets of users that should be managed in distinguishable ways.
- Performance objectives -- specification of distinct rates, called service rates, at which CPU, I/O, and real storage resources are to be provided to jobs in a performance group at different system workload levels.

A service rate is the amount of service units per second that a background or terminal job receives. A service unit is a measure of processing resources.

The installation uses service rates to specify the relative priority of the different subsets of users. The workload management routines then attempt to control these service rates and thereby meet the installation performance specification by causing the swapping of address spaces to take place.

The average applications programmer is not concerned with the activity of the system resources management routines or with the IPS. He takes advantage of the facility by identifying the performance group in which his jobs are to be included. (See the topic, "Job Control Language" for a description of the new JCL parameter that is to be used.) A TSO user will generally have his performance group number specified in the user attribute data set (UADS). The installation prescribes the performance groups that are to be applied to the specific jobs.

The following paragraphs describe, in greater detail, the concepts of service, performance groups, and performance objectives.

Service

Service units are used to measure the amount of processing resources provided to each address space. They are computed as a combination of three basic processing resources:

- CPU execution (i.e., one unit equals the execution of 10,000 instructions).
- I/O measure (i.e., the SMF I/O event count unit).
- real storage occupancy (i.e., one frame occupied for a multiple of CPU execution units).

The installation can specify coefficients by which to multiply each of these factors, as illustrated by A, B, and C in the following formula:

$$\text{service units} = A(\text{CPU}) + B(\text{I/O}) + C(\text{frames})$$

When an installation specifies performance objectives, it specifies one or more service rates, (i.e., how many service units per second a job should receive). The installation does not specify any particular amount of the individual resources that a job is to receive; it is assumed that different jobs will use CPU, I/O, and real storage resources in different proportions. However, the installation is able to

influence the relative importance of CPU, I/O, or real storage resources within the service definition by supplying coefficients to be multiplied by each resource. For example, if real storage is a critical resource, the installation can increase the value of C, the real storage service definition coefficient, in the service formula. Therefore, a job that uses a great deal of real storage will accumulate more service units than a job that uses less real storage. Likewise, if real storage is not a critical resource, C can be assigned the value 0, and a job will not accumulate service units at all by using real storage -- the real storage factor will drop out of the service definition. IBM supplies defaults for each of the coefficients.

Once the installation has determined the service definition (i.e., supplied the coefficients or accepted the defaults), the installation should normally be interested only in the number of service units and the service rates each job receives, not in particular amounts of CPU, I/O, or real storage used. The system activity measurement facility (MF/1) and the system management facilities (SMF) will record service unit measurements. (See the topic "MF/1 Reports and Tracing Records" in this chapter.) From this information, the installation can interpret service rate measurements in terms of response or turnaround time.

Performance Groups

The purpose of performance groups is to group together user transactions that the installation considers to have similar performance requirements. Basically, a user transaction is associated with an address space. It is considered to be a job or a job step for batch and remote job entry. For time sharing (TSO) jobs, a transaction is considered to be an interaction. A teleprocessing applications program that has multiple users and executes in a single virtual address space is considered to be a batch job and its terminal users are not distinguished from each other by the system resources management routines.

The installation can define as many as 255 performance groups, each identified by a distinct performance group number. Each performance group definition includes as many as eight performance group periods. By dividing a performance group into periods, an installation can associate different performance objectives with different periods in the life of a transaction. The duration of a period can be specified either as a number of real-time seconds or as a number of accumulated service units. For example, one performance group might be defined for short compile-load-go jobs that should have a very rapid turnaround time. The installation divides the performance group into two periods. The first period is associated with a high service rate and lasts until N number of service units are accumulated, where N has previously been determined to be a sufficient number of service units to complete a short compile-load-go job. If a job in the group is not complete after N service units are accumulated, it enters the second period, for which a lower service rate has been specified, and remains in this period for the duration of the job. service units are accumulated, it enters the second period, for which a lower service rate has been specified, and remains in this period for the duration of

Performance Objectives

A performance objective specifies service rates; how many service units per second an associated transaction should receive under different system workload conditions.

The motive for specifying different performance objectives is to give certain users better service at the expense of other users. However, the relative importance of this motive is dependent on the size of the system workload (i.e., the degree to which the demand for system resources exceeds the supply).

The installation specifies the treatment of different subsets of jobs by fixing arbitrary points, called workload levels, at which it defines target service rates for each subset of jobs. Conceptually, increasing workload levels can be thought of as increasing competition or demand for system resources. Thus, in order to preserve a high service rate for high priority jobs as demand (workload level) increases, an installation will reduce the target service rate for lower priority jobs.

In other words, if the system workload is average, all jobs should receive satisfactory service rates. However, if the system workload is heavy, the installation will want to provide an acceptable service rate to high priority jobs, and lower the service rate for lower priority jobs. Under very heavy workload conditions the installation might completely sacrifice the service rate of lower priority jobs (i.e., assign a service rate of 0) in order for the system to continue to provide acceptable service to the higher priority jobs. By defining workload levels, the installation can express the varying service relationships between groups of users at different system workload levels.

For example, an installation defines three performance objectives, numbered 2, 5, and 8. When the system workload is light, each performance objective should specify a better than satisfactory service rate for the transactions associated with it; the installation assigns the number 2 to this workload level and assigns corresponding service rates to each performance objective, as illustrated in the chart at the top of Figure 9. Under very heavy workload conditions, the installation wants to completely sacrifice the service rates for objective 5 and to lower the service rate for objective 8 but still provide acceptable service to jobs associated with objective 2. This workload level is assigned the number 8; the corresponding service rates associated with each objective are illustrated in Figure 9. The two intermediate workload levels shown in Figure 9 are established to reflect the changing relationships between the performance objectives as the workload level increases from 2 to 8.

Performance Objectives

Performance Objective Number	Service Rate for Workload Level 2	Service Rate for Workload Level 4	Service Rate for Workload Level 6	Service Rate for Workload Level 8
2	44	40	36	32
5	60	10	0	0
8	100	20	18	6

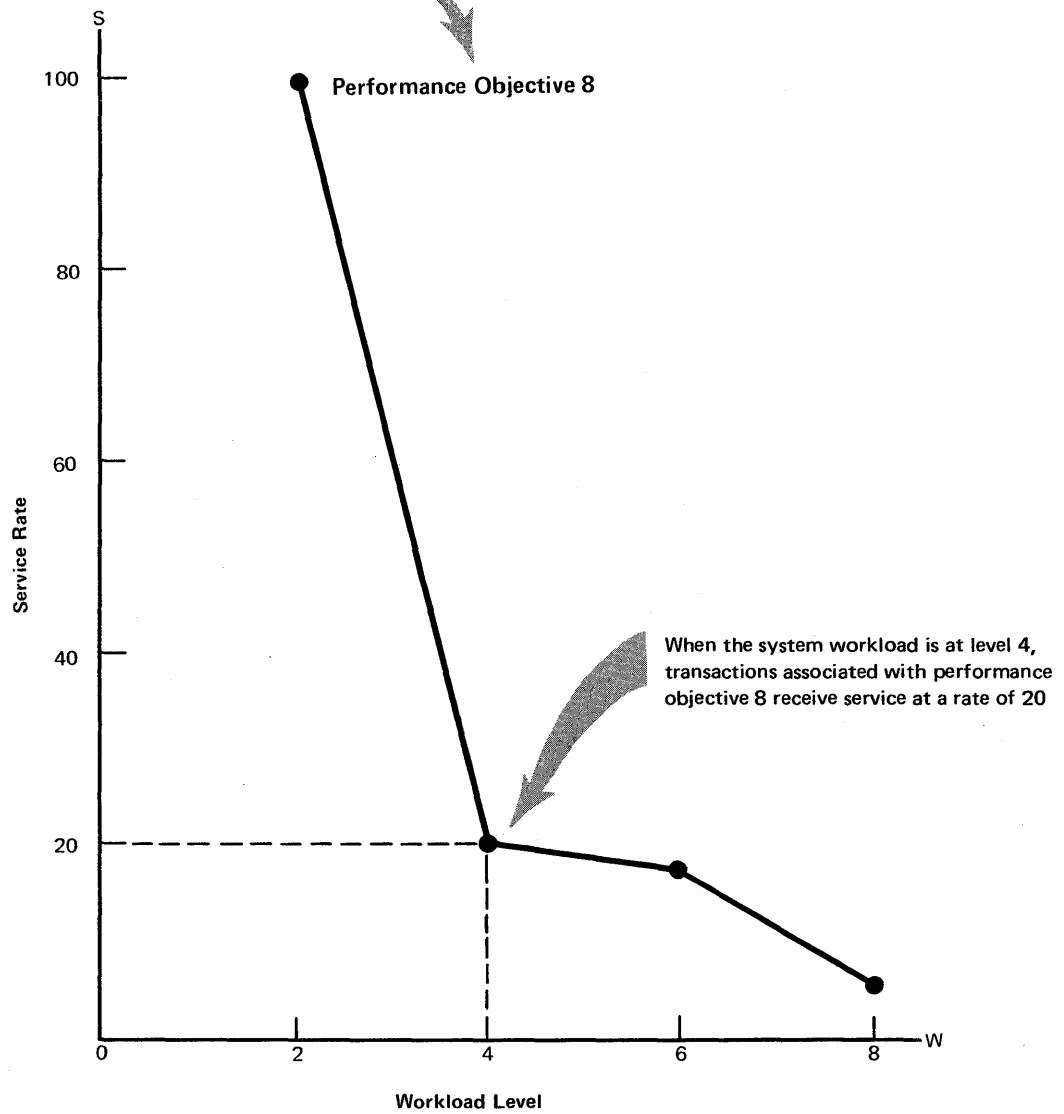


Figure 9. Performance Objective Definition

The system resources manager tracks the service rates provided to all users. As the system workload level increases or decreases, it adjusts the service rates to maintain the relationships between the performance objectives (and therefore between the transactions associated with the objectives) that the installation has defined.

Figure 9 also contains a graphic representation of one of the performance objectives (performance objective 8). This graph illustrates the service rates that the installation has defined for the system workload levels of 2, 4, 6, and 8. The workload management routines attempt to provide the service rate that applies to the performance objective for the associated job at the current system workload level. In Figure 9, when the current system workload is 4, the workload management routines recommend swapping in order to maintain a service rate of 20 for the jobs that are associated with performance objective 8. Jobs associated with other performance objectives will be targeted at service rates indicated in their performance objectives for a workload level of 4.

Associating a Transaction with a Performance Objective

Every transaction that enters the system has a performance group number. In Figure 10, transaction A is associated with performance group 9, transaction B is associated with performance group 8. For its first performance group period, transaction B will be associated with performance objective 5. This means that, for this period, the workload management routines will attempt to maintain a service rate for transaction B that equals the service rate specified in performance objective 5 for the workload level that equals the actual workload level of the system at a given time.

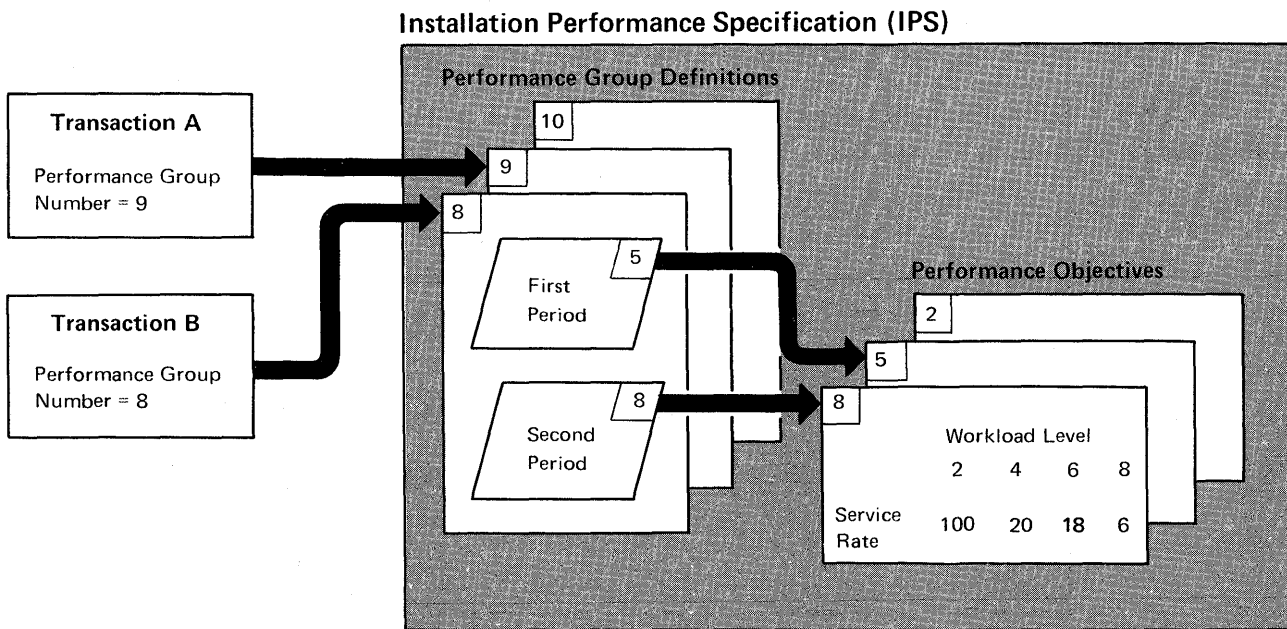


Figure 10. Associating a Transaction With a Performance Objective

The Relationship Between Workload Levels and Service Rates

To summarize the relationship between workload levels and service rates, consider a hypothetical installation whose users fall into one of three distinct subsets. Each subset can be considered to be a performance group. They are:

- Performance Group A -- A long running batch job that is run only once a day is the only job in this class and is extremely important.
- Performance Group B -- Moderate length batch jobs, of intermediate importance make up this class.
- Performance Group C -- Relatively short batch jobs, least in importance, make up this class.

Figure 11 illustrates the performance objectives that have been defined for each of the three performance groups. (Three distinct performance groups, each with only one defined period and one associated performance objective, are assumed for simplicity of illustration.)




Figure 11. Performance Objectives

Group A jobs are associated with performance objective 3, group B jobs with performance objective 8, and group C jobs with performance objective 5. Jobs associated with performance objective 3 demand a service rate of approximately 120 at workload levels 2 and 4. At workload level 6, performance objective 3 jobs demand a service rate of 100 while performance objective 5 jobs have

stopped receiving service, and performance objective 8 jobs have reduced their demands to 18 service units per unit time. At workload level 8, even performance objective 8 jobs nearly stop receiving service in deference to A jobs. As long as the A job receives a service rate of approximately 100, the installation knows that the job will complete processing in a satisfactory amount of time.

The workload level at which the system operates is a function of the number of jobs in each performance objective actually initiated at any time. Figure 12 illustrates system workload levels that result from three different possibilities:

- One job from each performance objective.
- One job from performance objective 3, four from performance objective 8, and two from performance objective 5.
- One job from performance objective 3, seven from performance objective 8, and one from performance objective 5.

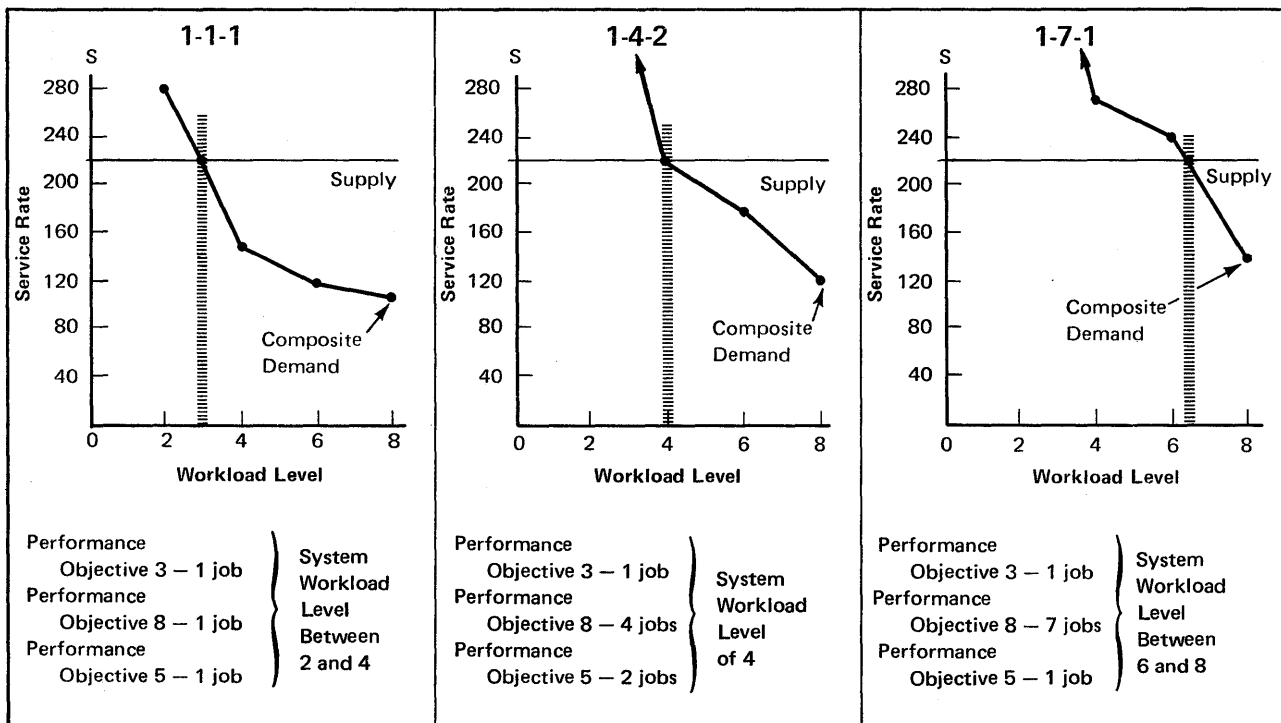


Figure 12. System Workload Levels

In these illustrations a simplified assumption is made that the system provides a composite service rate of 220 service units per unit of time to all active jobs; this is the "supply" line. The "demand" line represents the composite service demands from all jobs at each workload level. It is derived by multiplying the number of jobs in each performance group by the service rate specified for that performance objective at the particular workload level. In other words, with a job mix of 1-4-2, (the center graph on Figure 12), the composite service rate demanded by all jobs at workload level 4 is $1 \times 120 + 4 \times 20 + 2 \times 10 = 220$. (At workload level 4, a job associated with performance objective 3 demands a service rate of 120, a performance objective 8 job demands 20 and a performance objective 5 job demands 10.) The intersection of the supply and

demand lines produces the resulting system workload - that point of equilibrium at which the system can meet the composite demands of all active users.

In Figure 13, the different system workload levels that result from the three different job mixes are superimposed over the three performance objectives to illustrate the service rates that apply as the system workload level increases.




Figure 13. Service Rates for Different Performance Objectives at Various System Workload Levels

As shown in Figure 13, the following conclusions about the installation's performance objective specifications may be drawn from this hypothetical analysis:

- With 1 job initiated in each performance objective:
 - The system will be operating at a workload level between 2 and 4.
 - The performance objective 3 job will receive approximately 121 service units per unit time.
 - The performance objective 8 job will receive approximately 64 service units per unit time.
 - The performance objective 5 job will receive approximately 35 service units per unit time.
- With 1 performance objective 3 job, 4 performance objective 8 jobs, 2 performance objective 5 jobs:
 - The system will be operating at a workload level of 4.

- The performance objective 3 job will receive approximately 120 service units per unit time.
- The performance objective 8 jobs will each receive approximately 20 service units per unit time.
- The performance objective 5 jobs will each receive approximately 10 service units per unit time.
- With 1 performance objective 3 job, 7 performance objective 8 jobs, 1 performance objective 5 job:
 - The system will be operating at a workload level between 6 and 8.
 - The performance objective 3 job will receive approximately 100 service units per unit time.
 - The performance objective 17 jobs will each receive approximately 15 service units per unit time.
 - The performance objective 5 job will not be receiving service.

In each of the preceding examples, the sum of the service rates provided to all users will equal approximately 220 service units per unit time (the composite system “supply”).

System Resources Manager Parameters

The system resources manager parameters fall into two classes. Those that affect the recommendation values produced by workload management routines are contained in the SYS1.PARMLIB member IEAIPS00 (or IEAIPSxx); those that otherwise affect swapping decisions are contained in the SYS1.PARMLIB member IEAOPT00 (or IEAOPTxx).

These parameters provide the installation with a great degree of flexibility in stating performance requirements.

IPS Parameters (IEAIPSxx)

An installation performance specification contains four categories of information:

- Service definition coefficients -- The workload management routines use these values as multipliers, in calculating the service that a transaction has received. (See the topic “Service” earlier in this chapter.)
- Up to 32 workload level numbers -- Workload level numbers are used to specify the relative service rates that jobs of different priorities are to receive at varying levels of system activity. Each number is related to a specific service rate in each performance objective, and the numbers increase in magnitude as contention for system resources increases.
- Up to 64 performance objectives -- Each performance objective is specified in terms of the service rates to be provided at different workload levels. The performance objective number that is specified for each performance group period, identifies the performance objective to be used for that period.

- Up to 255 performance group definitions -- A performance group definition is identified by a performance group number. It may be subdivided into as many as eight different periods. Each period has the following parameters associated with it:
 - period duration
 - units code
 - performance objective number
 - response/throughput bias
 - interval service value

These categories of information are specified by the following parameters:

Performance Group Number

The performance group number is an integer between 1 and 255 that associates a transaction with a performance group definition. All transactions have performance group numbers. 1 is the default for batch jobs and 2 is the default for terminal jobs.

Period Duration

The period duration is a positive integer that indicates the length of the interval during which a specified performance objective is to be associated with a transaction.

Units Code

The units code is specified by the letter "S" or the letter "R". "S" indicates that the indicated period duration is to be measured in terms of accumulated service; "R" indicates that it is to be measured in terms of real-time seconds.

Performance Objective Number

The performance objective number is an integer between 1 and 64, that indicates the performance objective with which the transaction is to be associated for the duration of the indicated period.

Response/Throughput Bias

The response/throughput bias (RTB) is an integer that indicates, for each performance group period, that deviations from the service rate specified in the performance objective can be tolerated in favor of higher system throughput. "1" indicates that for jobs in this performance group period, the resource-use algorithms will produce evaluations that will influence swap decisions; "0" indicates that they will not.

Interval Service Value

The interval service value (ISV) specifies a minimum number of service units that a transaction must receive during each interval of real storage occupancy, before a swap evaluation is made by the workload management routines.

Workload Level

The workload level is an integer between 1 and 128. Up to 32 workload level numbers can be specified (in increasing magnitude representing heavier workload levels of the system). Workload levels are used to differentiate between light and heavy demands for the use of system resources. The same workload level numbers must be specified for each performance objective in an IPS.

Performance Objective

Each performance objective is identified by an integer from 1 to 64 followed by a definition of the service rates that are associated with the workload level numbers specified by the workload level parameter. A service rate must be specified for each workload level number defined in the IPS.

Service Definition Coefficients

The service definition coefficients (CPU, IOC, MSO) are used by the system resources management routines to calculate the service units that are provided to a particular job. Each coefficient is a multiplier on one of the three basic processing units in these calculations. For example, if an installation specifies an IOC coefficient of 3, whenever a job issues an I/O request, the system resources management routines multiply the I/O resource (one processing resource) by the service definition coefficient (3), and that job is considered to have used three service units.

Each of the service definition coefficients may be specified only once in an IPS. If a service definition coefficient is not specified, an IBM-supplied default value will be used.

Service Rates

Service rates are integers that define a target rate that system resources are to be supplied to an associated transaction. When grouped in sets called performance objectives, these define an intended service rate for any system workload condition.

Tuning Parameters (IEAOPTxx)

The parameters in the SYS1.PARMLIB member IEAOPTxx include:

Resource Factor Coefficients

Resource factor coefficients are values in the range 0.0 to 9.9. They specify values that are used as multipliers in the control algorithm to weight the recommendation values produced by the I/O and CPU load balancing algorithms. The two coefficients are IOC and CPU. Each coefficient has a default value of 1.0. The coefficients are used to make the recommendation values produced by the corresponding algorithms either more or less important than the recommendation values supplied by the workload management routines, for which an implied coefficient of 1.0 applies.

Resource Manager Constants

This category contains the enqueue residence value (ERV). The ERV is an integer in the range 0-999999. It is used to calculate the amount of time that an address space will not be swapped because the address space is enqueued upon a system resource required by another address space. The time interval is determined by multiplying the ERV by the (model dependent) time required to execute 10,000 machine instructions. (This procedure is intended to keep the time interval relatively constant on all System/370 models.)

System Activity Measurement Facility (MF/1)

The system activity measurement facility (MF/1) is new in VS2 Release 2. Installations may use it to monitor selected areas of system activity and obtain feedback in the form of trace records and/or formatted reports. Measurements may be gathered independently on:

- CPU activity
- Channel activity and channel-CPU overlap activity
- I/O device activity and contention
 - Unit record devices
 - Graphics devices
 - Direct access storage devices
 - Communication equipment
 - Magnetic tape devices
 - Character reader equipment
- Paging activity
- Workload activity

MF/1 activity is divided into three categories:

- SAMG (system activity measurement gathering) obtains the measurements of system activity requested by the installation.
- SARG (system activity report generation) produces formatted reports of desired system measurements.
- MFC (measurement facility control) controls the collection, tracing and reporting of system activity.

SAMG consists of distinct sets of measurement gathering (MG) routines associated with each class of measurement data. Only those sets associated with the measurement data requested by the system operator are loaded and active during MF/1 reporting intervals.

SARG generates formatted summary reports of information requested by the system operator, collected by SAMG, and routed to it by MFC. Printing of these reports may be done as they are generated, or may be deferred until MF/1 termination.

MFC is the system task that controls MF/1 operation. It causes the loading of the MG routines that will be needed to obtain the desired information, passes control to these routines at the desired intervals, routes collected information to SARG as required, and terminates MF/1 execution at the end of its specified duration.

MF/1 Operation

The operation of MF/1 is controlled by a set of parameters that may be contained in:

- The PARM field of the START command that is issued to start MF/1 processing.

- The PARM field of the EXEC statement in the MF/1 cataloged procedure.
- The MF/1 partitioned data set member (typically a SYS1.PARMLIB member).

The IBM-supplied cataloged procedure for MF/1 names the MF/1 partitioned data set member that contains card image MF/1 control input. Member names are of the form IRBMF1xx, where xx is a decimal digit field. In the absence of any explicit specification, the SYS1.PARMLIB member IRBMF100 is used.

The system operator issues the START command to begin MF/1 monitoring. The parameters controlling the operation of MF/1 are gathered from the above sources in the order stated. Explicit specification of an option in the PARM field of the START command will therefore override any alternative specifications of that option in any other source. If a parameter is not specified in any of the three sources the system default value is used.

The MF/1 options are summarized below. The underlined values are the system default values, as well as the values that appear in the SYS1.PARMLIB member IRBMF100.

CPU/NOCPU

PAGING/NOPAGING

CHAN/NOCHAN

WKLD(PERIOD/GROUP/SYSTEM)/NOWKLD

DEVICE(device list)/NODEVICE

These parameters are used to specify which of the classes of MG routines (CPU activity, paging activity, channel activity, workload activity, I/O device activity) will be loaded and executing during the MF/1 reporting period.

UNITR/NOUNITR

TAPE/NOTAPE

DASD/NODASD

COMM/NOCOMM

GRAPH/NOGRAPH

CHRDR/NOCHRDR

These parameters are the elements of the device list above. They specify the device types (unit record, tape, direct access, communications, graphics, character reader) about which information will be collected.

REPORT(REALTIME/DEFER)/NOREPORT

This parameter specifies whether or not formatted reports of the gathered data will be generated, and when the reports will be generated (i.e., as the data is gathered, or after MF/1 processing terminates).

SYSOUT(class)

This parameter specifies the SYSOUT class to which formatted reports will be directed. Class A is the default.

TRACE/NOTRACE

This parameter specifies whether or not trace records will be written to the SMF data set.

INTERVAL(value/valueM)

This parameter specifies the interval at which data will be gathered for report formatting and/or trace record writing. The range is from 1 to sixty minutes. The default is 15M.

CYCLE(value)

This parameter specifies the frequency at which observations are made of sampled data. The range is from 50 to 999 milliseconds. The default is 250 milliseconds.

STOP(value/valueM/valueH)/NOSTOP

This parameter specifies the desired time interval for MF/1 activity, in minutes or hours (the default is minutes). The operator STOP command will terminate execution at any time, regardless of the value for this parameter. The default value is 15M.

MEMBER(nn)

The value specified by this parameter is added to IRBMF1 to form the name of the partitioned data set member that contains the MF/1 options. The default is 00.

OPI/NOOPI

This parameter specifies whether or not a list of the keyword options to be used will be printed at the operator's console. If the list is printed, the operator will be able to respond with any desired changes to the option list.

MF/1 Reports and Tracing Records

Measurement output is produced in two forms:

- Printed Reports -- The data is formatted and printed for each class of system activity data collected.
- Trace Records -- The data is written as a trace record to the SMF data set. One or more trace records are produced for each class of system activity data collected.

The printed reports and trace records contain the same type of information, but in different forms. The installation can specify desired output as:

- Printed reports only.
- Trace records only.
- Both printed reports and trace records.

The following sections describe the contents of each type of printed report and trace record.

CPU Activity Report

This report provides information about the use of the central processing unit(s). The total CPU wait time that has taken place during an MF/1 reporting interval is generated for each CPU. The CPU wait time is also expressed as a percentage of the reporting interval time.

Paging Activity Report

This report provides detailed information about the demands made on the system paging facilities during the reporting interval. It also distinguishes between the use of real main storage and external page storage. The following categories of information are included:

- Page reclamations -- Page reclamation rates and percentages are produced, for system pageable areas and the user's private address spaces, for both VIO and non-VIO accesses.
- Page-ins -- Page-in rates and percentages are produced for all pages transferred from external page storage to real storage. This information is provided for system pageable areas and the user's private address spaces.
- Page-outs -- Page-out rates and percentages are produced for pages transferred from real storage to external page storage. This information is provided for system pageable areas and the user's private address spaces.
- External page storage counts -- Total page slots are produced, as well as the number of unused slots, unavailable slots, and slots containing either data set pages or virtual address space pages.
- Pageable real storage counts -- The total number of pageable real storage frames is produced, as well as breakdowns into counts and percentages of used and unused page frames.
- Swap counts -- The total number of swaps are produced, as well as average pages per swap-out, and average pages per swap-in.

Workload Activity Reports

These reports provide information about the system workload activity and its relationship to the installation-specified performance objective for each performance group period. This information can be used to evaluate and modify the IPS. The three different workload activity reports that are available provide three different levels of detail. The types of reports desired are specified by input parameters indicated below from the most detailed report to the least detailed report.

- Performance group period report -- (WKLD(PERIOD))
- Performance group report -- (WKLD(GROUP))
- System summary report -- (WKLD(SYSTEM))

Performance Group Period Report: This report provides, for each performance group period within each performance group:

- The total number of service units used by all transactions in that performance group period.
- The average service rate for all transactions within that performance group period.
- The average workload level of all associated transactions.
- The average number of transactions that were concurrently active.
- The total number of transactions that terminated.
- The average elapsed time of a transaction that terminated.

Performance Group Report: This report provides, for each performance group, for one MF/1 reporting interval:

- The total number of service units used by all transactions in that performance group.
- The average number of associated transactions that were concurrently active.

- The total number of associated transactions that terminated during the reporting interval.

System Summary Report: This report provides, for one MF/1 reporting interval:

- The total number of service units used by all transactions.
- The mean workload level of the system.
- The average number of transactions that were concurrently active.
- The total number of transactions that terminated during the reporting interval.

Channel Activity Report

This report provides information about channel loading for all channels in the system. For each active channel, a "channel activity" count provides the total number of successful SIO instructions (excluding the sense SIO instructions) issued to that channel during the reporting interval. A "percent channel active" indication provides the percentage of the total reporting interval during which each channel was busy. A "percent channel busy and CPU waiting" indication provides the percentage of the reporting interval during which each channel was busy while the associated CPU was in a wait state. An installation may use this information, in conjunction with the information contained in the I/O device activity reports, to determine possible bottlenecks associated with channels.

I/O Device Activity Reports

These reports provide information about I/O device loading for all devices in the chosen device class(es). The classes desired are specified in the input parameter string (UNITR, TAPE, DASD, COMM, GRAPH, CHRDR). For each device report, a "device activity count" indicates the number of successful SIO instructions (excluding the sense SIO instructions) issued to the device during the reporting interval. A "percent busy" indication provides the percentage of the reporting interval during which the device was busy. A "mean queue length" value provides an average number of requests that were enqueued and waiting to be serviced on the device during the reporting interval. An installation may use this information, in conjunction with the information contained in the channel activity reports, to determine possible bottlenecks in the I/O subsystem.

Tracing

If TRACE is specified among the MF/1 input parameters, a trace record will be written to the SMF data set, at the end of each reporting interval, for each report requested by the parameters. Each trace record contains information similar to the contents of the corresponding formatted report. Totals, averages, and percentages are not explicitly contained in the trace record, but may be calculated from the explicit data. The SMF records, and their corresponding printed reports are:

- SMF70 -- CPU Activity Report
- SMF71 -- Paging Activity Report
- SMF72 -- Workload Activity Report
- SMF73 -- Channel Activity Report
- SMF74 -- I/O Device Activity Report

Tuning With MF/1

Tuning a system is the process of modifying its hardware or software characteristics to bring performance closer to installation defined objectives. Once an installation's performance objectives are clearly defined, MF/1 can be used to help an installation both specify and meet them. For example, the channel activity report and I/O device activity report can provide information about the efficiency of the system-wide I/O load balance and how it affects system throughput. Analysis of these reports could indicate possible system configuration changes or data set residence modifications.

MF/1 can also be used to determine if the IPS reflects the installation's turnaround/response requirements and, if so, whether the IPS objectives are being met. It can further be used to analyze the specified performance objectives for possible modification.

Tuning Example

Figure 14 shows the division of the installation's jobs into three performance groups. Assume that users associated with performance group 7 -- the time sharing users -- are experiencing poor response time at their terminals. The installation decides to activate MF/1 for a three hour period during a typical day, to obtain six reports on the workload activity for analysis. To initiate this reporting the operator enters:

```
START MF1.MF1,,(NOCPU,NOPAGING,NOCHAN,  
NODEVICE,INTERVAL(30M),STOP(3H))
```

(Assume that there are no EXEC statement parameters in the MF/1 procedure, and that the default SYS1.PARMLIB list IRBMF100 is used.)

With the above command and the standard defaults, the following options will control MF/1 execution in this case:

```
NOCPU  
NOPAGING  
NOCHAN  
WKLD(PERIOD)  
NODEVICE  
REPORT(DEFER)  
SYSOUT(A)  
NOTRACE  
INTERVAL(30M)  
CYCLE(250)  
STOP(3H)  
OPI
```

IPS Values

Performance Group Number	Significance of Performance Group to Installation	FIRST PERIOD					SECOND PERIOD				
		Duration	Units	Performance Objective	RTB	ISV	Duration	Units	Performance Objective	RTB	ISV
7	Terminal Jobs	Interaction	-	2	-	-	-	-	-	-	-
8	High Priority Batch	3000	S	8	-	-	Rest of Job	-	8	-	-
9	Low Priority Batch	Entire Job	-	5	-	-	-	-	-	-	-

Performance Objectives




Figure 14. Tuning Example

Assuming that the operator does not prematurely terminate MF/1 execution with the STOP command, MF/1 will terminate at the end of three hours, and the six reports that were collected at half-hours intervals will then be printed. The installation can now begin its analysis.

Two figures of interest in this case are the average system workload level, and the average active time for all performance group 7 transactions. Suppose that these reports indicate that, during the time that MF/1 monitoring took place, the system was operating approximately at workload level 4 and that the average response time for time sharing users was 30 seconds (considered by the installation to be poor response time).

An analysis of the performance objective specification shows that, at workload level 4, performance objective 2 specifies that performance group 7 transactions receive 40 service units per unit time. The installation might reasonably wish to increase the service rate for its time sharing users in order to improve the response time. This increase would raise the performance objective curve for performance objective 2, as shown in Figure 15.




Figure 15. Changing a Performance Objective

Figure 16 shows that raising the service rate for performance objective 2 raises the system workload level, all other factors being equal. Transactions associated with performance objectives 5 and 8 would, therefore, be affected. Transactions associated with performance objective 5 would be most severely affected because their demand curve slopes most severely to the right of workload level 4.

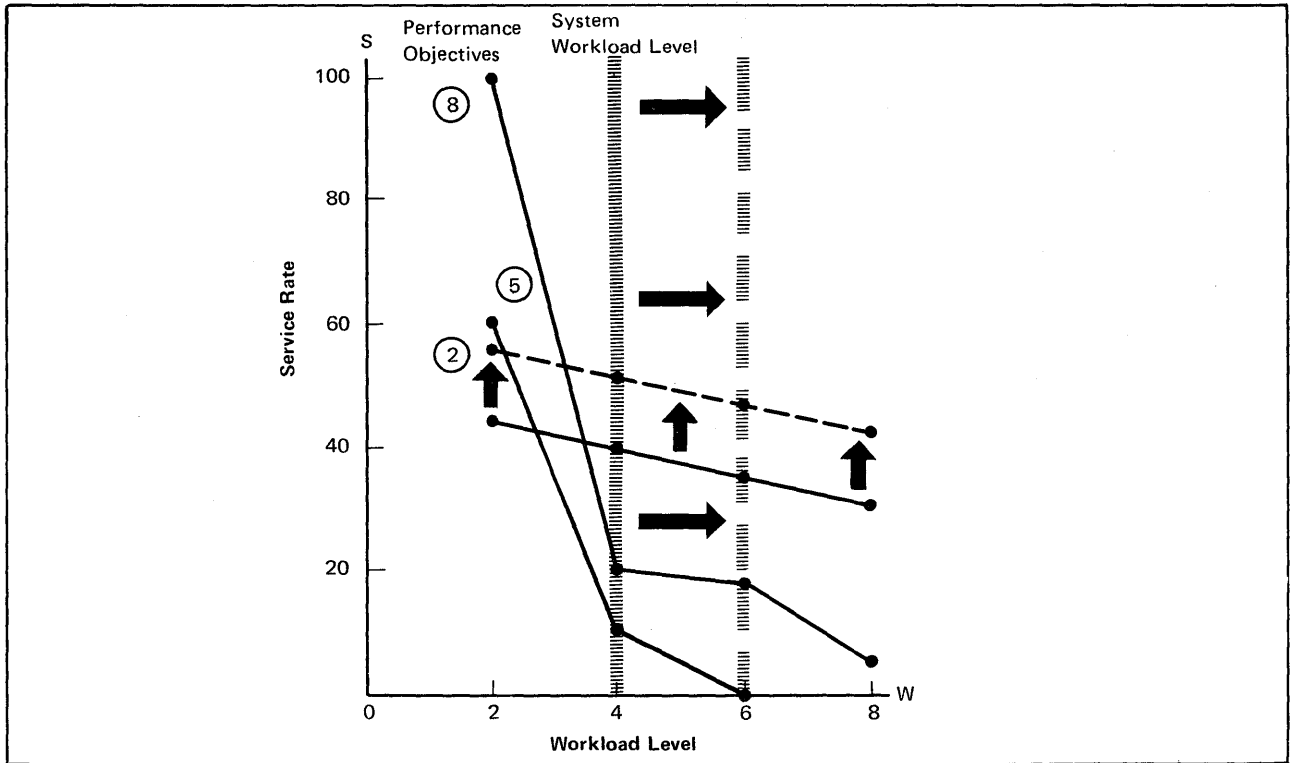


Figure 16. Effects of a Performance Objective Change

Before the installation modifies performance objective 2, it is able to perform further analysis to determine the quantitative effects of such a change. Suppose the MF/1 reports showed an average of 20 jobs associated with performance objective 2, 8 jobs with performance objective 5, and 4 jobs with performance objective 8. Then, some measure of the average system rate of service supplied to all jobs can be obtained by multiplying the number of jobs at each performance objective by the service rate for that performance objective at workload level 4 (i.e., $20 \times 40 + 8 \times 20 + 4 \times 10 = 800 + 160 + 40 = 1000$ total service units per unit time).

In this case, increasing the performance objective 2 curve by 10 service units, results in a total increase in demand of 200 service units (10 times the 20 jobs associated with performance objective 2). Thus, the workload level of the system shifts to the right (increases) to compensate for this increase in demand, while maintaining the system supply of 1000 total service units approximately constant. (See Figure 16.) If the installation considers the service rates provided for all performance objectives at the projected new workload level to be acceptable, it may modify the IPS to contain the new performance objective 2.

The installation may then use MF/1 monitoring to verify the accuracy of the projections; repeating the entire procedure, if desired, until satisfactory results are achieved.

The preceding example is not meant to reflect actual figures or relationships, but rather to indicate the possibilities for tuning that MF/1 can provide. The illustrated method of improving response time is neither the only one, nor necessarily the best -- other alternatives include varying the interval service value parameter, dividing the time sharing performance group into more than one period, and increasing the number of performance groups provided for time sharing users. Any method chosen to improve this or other system performance characteristics will suggest other uses and variations for MF/1.

Reducing I/O Device Allocation Contention Among Job Steps

A major objective of the I/O device allocation design in VS2 Release 2 is to reduce serial processing. A new, unserialized path through the I/O device allocation routines is designed to handle the most common types of device requests, allowing the potential for reduced I/O device allocation contention among job steps. Allocation processing is not serialized for:

- Permanently resident or reserved direct access volumes.
- Teleprocessing lines.
- Job entry subsystem data sets.
- Virtual I/O (VIO) data sets.

An installation can also reduce serialization for allocation requests that do not fall into one of these three categories by defining subsets of device groupings. This will cause serialization to take place for device subsets, rather than for all devices.

Defining a Reserved Mount Characteristic

Because the I/O device allocation routines first attempt to satisfy all direct access requests on permanently resident or reserved volumes, an installation that requires little or no private direct access volume mounting can take advantage of this feature by assigning the reserved mount characteristic to as many direct access volumes as possible.

The reserved mount characteristic can be assigned to direct access volumes in two ways:

- Issuing a MOUNT command (volumes identified by the MOUNT command automatically become reserved).
- Specifying the reserved mount characteristic for the volume in the VATLST00 or the VATLSTxx list of SYS1.PARMLIB.

Defining Device Groupings

When all I/O device allocation requests cannot be satisfied by allocating to permanently-resident or reserved volumes, allocation processing is serialized for device types (e.g., concurrent allocation processing can take place for 2314 devices and 3330 devices). However, this serialization of allocation can be reduced by using the UNITNAME macro instruction at system generation (SYSGEN) time to define subsets of device types (device groupings), and by including each device in not more than one device grouping.

For example, if an installation has two 2314 Direct Access Storage Facilities available, at addresses 130 through 137 and 230 through 237, it could define one device grouping for batch data sets and another for time sharing data sets, making sure that each device address appears in only one device grouping. The UNITNAME macro instruction is specified as indicated below:

```
UNITNAME UNIT=[( 130,4 ), ( 230,4 )],NAME=SYSBATCH  
UNITNAME UNIT=[( 134,4 ), ( 234,4 )],NAME=SYSTSO
```

I/O device allocation processing will, in this case, be serialized for the subset of 2314 devices, specified by the device grouping name SYSBATCH or SYSTSO (e.g., concurrent allocation processing can take place for the 2314s at addresses 130 through 133 and 230 through 234 (SYSBATCH) and for the 2314s at addresses 134 through 137 and 234 through 237 (SYSTSO)).

But if, for example, the 2314 at address 130 is included in both device groupings, allocation processing will be serialized for the entire set of 2314 devices.

System Integrity

A highly desirable property of an operating system is its ability to insure that one program cannot interfere with or modify another program's (system or user) execution unless it is authorized to do so. For reliability and system availability, this is extremely important; for data security, it is essential. VS2 Release 2 provides this capability in the form of system integrity support.

System integrity is defined as the ability of the system to protect itself against unauthorized user access to the extent that security controls cannot be compromised. That is, there is no way for an unauthorized* problem program using any system interface to:

- bypass store or fetch protection, i.e. read or write from or to another user's areas.
- bypass password checking, i.e. access password protected data for which a password has not been supplied.
- obtain control in an authorized state.

In VS2 Release 2 all known integrity exposures have been removed. IBM will accept as valid, any APAR that describes an unauthorized program's use of any system interface (defined or undefined) to bypass store or fetch protection, to bypass password checking, or to obtain control in an authorized state.

*An authorized program in VS2 Release 2 is one that executes in a system key (keys 0-7), in supervisor state, or is authorized via the Authorized Program Facility (APF).

Installation Responsibility

To ensure that system integrity is effective and to avoid compromising any of the integrity controls provided in the system, the installation must assume the responsibility for the following items:

The installation must be responsible for the physical environment of the computing system. Operations personnel and system programmers have, in effect, uncontrolled access to certain portions of the operating system. Problems of this nature are discussed in **The Considerations of Physical Security in a Computer Environment, G520-2700**.

The installation must be responsible for the adoption of certain procedures (e.g. the password protection of appropriate system data sets) that are a necessary complement to the integrity support within the operating system itself. The most significant of these procedures are described briefly in the topic "Installation Procedural Responsibilities."

The installation must ensure that its own modifications and additions to the control program do not introduce any integrity exposures. That is, all installation-written authorized code (e.g. an installation SVC) must perform the same or an equivalent type of validity checking and control that the VS2 Release 2 control program employs to maintain system integrity.

Areas of Concern

Several areas with the potential for integrity exposures have been identified. These areas, and what has been done to eliminate them as potential sources of integrity exposures, are described below. The description is a guideline to aid the installation in preserving system integrity in any modification or addition to the control program. It is also intended to aid in determining whether an impact on existing installation code might occur, where that code is dependent on the use of non-standard interfaces to the control program. There should be no impact on installation-written routines that use standard interfaces (problem program/system interface described in an SRL), because no standard interface has been removed from the system control program in connection with system integrity support. VS2 Release 2 system integrity support is concerned only with restricting the unauthorized problem program; there is no attempt or intention to prevent the installation manager or any authorized system programmer from modifying the system control program.

User-Supplied Addresses for User Storage Areas

A potential integrity exposure exists whenever a routine having a system protection key (key 0-7), accepts a user-supplied address of an area to which a store or fetch is to be done. If the system routine does not adequately validate the user-supplied address to ensure that it is the address of an area accessible to the user for storing and fetching data, an integrity violation can occur when the system key program routine:

- Stores into (overlays) system code or data (e.g. in the nucleus or the system queue area), or into another user's code or data.
- Moves data from a fetch-protected area that is not accessible to the user (e.g. the fetch-protected portion of the common service areas), to an area that is accessible to him.

The elimination of this problem requires that system key routines always verify that the entire area to be stored into, or fetched from, is accessible (for storing or fetching) to the user in question. The primary validation technique is the generally established VS2 Release 2 convention that system key routines obtain the protection key of the user before accessing the user-specified area of storage. For example, VS2 Release 2 data management SVC routines (which generally run in key 5) assume the user's key before modifying a data control block (DCB) or an I/O block (IOB).

User-Supplied Addresses for Protected Control Blocks

A potential integrity exposure exists whenever the control program (system key/privileged mode) accepts the address of a protected system control block from the user. For most system control blocks this situation should not be permitted to exist. However, in certain cases it is necessary to allow the user to provide the address of a system control block that describes his allocation/access to a particular resource (e.g. a data set), in order to identify that resource from a group of similar resources, (e.g. a user may have many data sets allocated). Inadequate validity checking in this situation can create an integrity exposure, since an unauthorized problem program could provide its own (counterfeit) control block in place of the system block and thereby gain the ability to:

- Access a resource in an uncontrolled manner (since the control block in this case would normally define the restrictions, such as read-only for a data set, on the user's allocation to the resource).
- Gain control in a privileged state (because such control blocks may contain the addresses of routines that run in privileged mode or with a system (0-7) key).
- Cause various other problems depending on exactly what data the system may be keeping in the control block involved.

To avoid this type of exposure, the control program must verify, for every address of this type that is accepted from a problem program, that the address is that of:

1. A protected control block created by the control program.
2. The correct type of control program block (e.g. a TCB versus a DEB, or a QSAM DEB versus an ISAM DEB).
3. A control block created for use in connection with the user (job step) that supplied the address.

In VS2 Release 2, verification is generally accomplished by establishing a chain or table of the particular type of control block to be validated. This chain or table is located via a protected and job step-related control block that is known to be valid (i.e. whose address is not allowed to be supplied by the user, and is located via a chain of protected control blocks that begins with a control block known to be valid or fixed at a known location at IPL time, such as the CVT*). Thus a control block can only be entered in the chain/table by an authorized program, which satisfies (1) above; by definition the chain/table establishes the type of control block, which satisfies (2) above; and also by definition each chain/table is located only through a job step-related control block, satisfying point (3).

Resource Serialization

Resource serialization is another potential source of integrity problems. Here an integrity exposure can result when serialization controls on sensitive control program resources are nonexistent, or inadequate to the extent that an unauthorized problem program can directly, or indirectly through a part of the control program (e.g. by issuing an SVC), change the content or status of a system resource (e.g. control block, work area), while another portion of the control program (e.g. another SVC) is using that resource or is in some way dependent on its content or status remaining unchanged for a given period of time.

Adequate control of serial resources becomes even more significant in a multiprocessing environment since some uniprocessor serialization techniques (e.g., disablement), are no longer effective because of the possibility of multiple tasks, even multiple tasks in the same job step, using the same serial resource and running concurrently on different CPUs.

*This does not imply that a system routine must go back to the CVT or similar control block every time it wants to establish a valid chain. Typically, a control block address not too far down on such a chain is available already validated in a register. For example, the first load of an SVC may receive control with a valid TCB address in a register.

To eliminate this as a potential exposure, a locking technique is used to serialize access to the resources in question. However, a locking technique is not effective unless all routines that have the potential for changing the resource, or that depend on its status remaining unchanged for a given period, make use of the (same) locking mechanism. Effective use of a locking technique therefore requires considerable investigation and effort, to determine and keep track of system resources that must be serialized and routines that access such resources.

In VS2 Release 2, a combination of ENQ/DEQ and a new hierarchical locking structure with multiple types of locks, is the primary method used to synchronize the use of serially reusable resources. The locking mechanism is restricted to key 0 programs, which prevents unauthorized problem programs from interfering with the serialization of system resources. For the same reason, ENQ/DEQ is restricted to authorized programs for all major names of the form SYSZxxxx, and for the following explicit major names: SYSDSN, SYSIEECT, SYSIEFSD, SYSIEA01, SYSVTOC, and SYSPSWRD.

Another integrity consideration, related to the above problem, but generally applicable to all validity checking for integrity purposes, is the time-of-check-to-time-of-use problem. This refers to the fact that from the time of a validity check until the completion of the operation associated with that validity check, the variables on which the outcome of the validity check is based must not be allowed to change to the extent that the result of the validity check would be changed. While this is a relatively obvious statement, the requirement it imposes has considerably influenced the direction of the integrity support in VS2 Release 2.

In some cases, this requirement can be met "automatically" through use of hardware serialization mechanisms, such as checking the validity of fetch/store operations by assuming the user's protection key. In other cases, techniques such as saving validated data in protected (storage keys 0-7) core and use of the above-mentioned locking mechanisms are required.

Resource Identification

Resource identification is another area that can be subject to integrity exposures. Exposures can result if the control program does not maintain and use sufficient data to uniquely distinguish one resource from other similar resources. For example, a program must be identified by both name and library to distinguish it from other programs. The consequences of inadequate resource identification are problems such as the ability of an unauthorized problem program to create counterfeit control program code or data, or to cause varying types of integrity problems by intermixing incompatible pieces of control program code and/or data.

The general solution can only be stated as the reverse of the problem; that is, the control program must maintain and use sufficient (protected) data on any control program resource, to distinguish between that resource and other control program or user resources. The following are examples of the controls that VS2 Release 2 employs to comply with the requirement:

- In general, authorized program requests to load other authorized programs are satisfied only from authorized system libraries (see the topic "Control Program Extensions").
- VS2 Release 2 takes explicit steps to ensure that routines loaded from authorized system libraries are used only for their intended purpose. This

includes expanded validity checking to remove any potential for the unauthorized program to specify explicitly which of the authorized library routines are to gain control in any given situation.

- Sensitive system control blocks are validated as being the "correct" blocks to be used in any given control program operation. (See the topic "User-Supplied Addresses of Protected Control Blocks").

SVC Routines Calling SVC Routines

A potential problem area exists whenever a problem program is allowed to use one SVC routine (routine A) to invoke a second SVC routine (routine B) that the problem program could have invoked directly. An integrity exposure occurs if SVC routine B bypasses some or all validity checking based on the fact that it was called by SVC routine A (an authorized program), and in addition, user-supplied data passed to routine B by routine A either is not validity checked by routine A, or is exposed to user modification after it was validated by routine A. This problem will not exist if the user calls SVC routine B directly, because the validity checking will be performed on the basis of the caller being an unauthorized program.

SVC routine A, which is aware that it has been called by an unauthorized program, must ensure that the proper validity checking, etc., is accomplished. However, it is usually not practical for SVC routine A to do the validity checking itself, because of the potential for user modification of the data prior to or during its use by SVC routine B. The general solution should be for SVC routine A to provide an interface to SVC routine B, informing routine B that the operation is being requested with user-supplied data in behalf of an unauthorized problem program (implying that normal validity checking should be performed).

In practice, in VS2 Release 2, most SVC B type routines that could be subject to this problem, use the key of their caller as a basis for determining whether or not to perform validity checking. Therefore most VS2 Release 2 SVC A type routines have simply adopted the convention of assuming the key of their caller before calling the SVC B routine.

Control Program and User Data Accessibility

Important in maintaining system integrity, is the consideration of what system data is sensitive and must be protected from the user, and what data can be exposed to user manipulation. The implications of the exposure of the wrong type of data are obvious.

In general, it is necessary to protect the following types of data:

- Code, and the location of code, that is to receive control in an authorized state.
- Work areas for such code, including areas where it saves the contents of registers.
- Control blocks that represent the allocation or use of system resources.

VS2 Release 2 maintains items such as the above in system storage, or in a separate region in the case of some APF-authorized (see the following section) programs. It may also be necessary to protect, for a limited period, certain data that is normally under the control of the user (e.g. to prevent its modification

during a critical operation). In this case VS2 Release 2 provides fetch protection for such data if:

- The data consists of proprietary information (e.g. passwords).
- The control program cannot determine the nature of the contents of the data area.

Control Program Extensions

The final potential problem area involves the fact that there exists a somewhat hazy distinction between the control program and certain types of problem programs. In most installations, there are problem state/user key (keys 8-15) programs that are actually extensions to the control program in that they are allowed (by means of various special SVCs, etc.) to bypass normal system controls over access to system resources. For example, a special utility program that scans all the data on a pack might be able to avoid the normal system extent checking on a direct access volume. The Authorized Program Facility (APF) was introduced in VS2 Release 1 as a means of identifying these control program extensions as "authorized" programs, and restricting the use of the various "special" SVCs used by such programs to bypass normal system controls. This use of APF is continued in VS2 Release 2 with the restriction of SVCs under APF extended to include certain new SVCs introduced in VS2 Release 2.

If an installation has its own control program extensions and special SVCs that allow the bypass of normal system security or integrity checks (e.g., an SVC that returned control in key 0), and if such SVCs are not currently restricted from use by an unauthorized program, then the APF facility should be used to restrict these SVCs and to authorize the control program extensions that use them. A description of how to use the APF facility can be found in appropriate VS2 publications.

An important VS2 Release 2 addition is that APF has been extended to recognize and accept authorized programs from "authorized" libraries other than SYS1.SVCLIB, SYS1.LINKLIB, and SYS1.LPALIB. Any library may be designated as an authorized library by inclusion of its name in the SYS1.PARMLIB member, IEAAPF00 or IEAPFxx, prior to IPL. Such libraries are the equivalent of the above system libraries in that nonspecific system requests (library not specified) for routines that will execute with system key, privileged mode, or APF authorization, can be satisfied by a module loaded from any authorized library.

Note: The user of APF should also be aware that APF authorization is defined such that any SVC that is restricted by the APF mechanism, can be executed by any system key (0-7) or privileged mode routine. However, this is not true in reverse. An APF authorized program will not automatically be allowed system services that are restricted by system key or privileged mode tests.

Installation Procedural Responsibilities

As noted previously, the installation is responsible for the implementation of procedures that are necessary in certain instances to complement the VS2 Release 2 control program integrity/security support. These responsibilities are discussed below. Note that in some instances, the concern is for basic system integrity, while other items are concerned with a particular type of security.

1. The installation must password protect appropriate system libraries. System integrity obviously cannot be maintained if system code and data is exposed to arbitrary modification by any user on the system. For integrity purposes, it is generally sufficient to protect appropriate libraries from write access; no password is required for read access but a password is required for write access. However, for security purposes, it is necessary to protect certain system data sets (e.g. the PASSWORD data set itself) from read as well as write accesses.
2. Because the checkpoint data set produced by the Checkpoint/Restart facility contains sensitive system data normally protected from the user, maintaining system integrity requires that such data sets be protected from modification (or from being counterfeited) prior to their use by Restart. VS2 Release 2 implements a facility whereby the installation can adopt a set of special procedures/controls over checkpoint data sets that will eliminate their potential for compromising system integrity. The control mechanism involves a combination of:
 - System/operator validation of checkpoint data sets.
 - External labeling procedures for "checkpoint volumes."
 - Offline control of "checkpoint volumes."

Details will be available in later documentation.

3. The installation should be aware of the following password protection requirements, which may necessitate some changes in installation procedure.
 - Because of the characteristics of magnetic tape devices, once a program has access to one file on a tape, it can not be ensured that the program is not able to access other files on that tape, even though it may not possess passwords for those files. For this reasons, VS2 Release 2 has implemented a change in the handling of password-protected tape data sets to attempt to ensure that if a user is allowed access to one data set on a tape volume, he can be presumed also authorized to other data sets on the volume. To this end, whenever a user attempts to create a tape data set whose file sequence number is greater than 1, the protection attributes of the new file will be compared against the protection attributes of the immediately preceding file on the tape (no protection, protected for write, protected for read and write, etc.). If the protection attributes do not match, the new data set will not be created. In addition, if the previous data set is password protected, its password must be supplied before the new data set will be created.

Note that these changes do not eliminate problems caused by an installation library containing already-existing tape files with mixed protection attributes.

- For the reasons given above, the security-conscious installation should ensure that password-protected data sets are created only on tapes which

are unused or completely erased. This is necessary because in certain cases residual information left on the tape after the password data set was created, could result in a user being able to open one of the previous data sets on the tape, thus gaining access to the whole tape.

4. In the following cases, the installation may wish to utilize dynamic allocation and JCL validation exits to achieve additional control over device allocation and user identification.
 - In VS2 Release 2, as previous releases, allocation of 270X lines is shared. The above-mentioned exits can be used to ensure that concurrent allocation of 270X lines does not occur.
 - Fetch/store protect is not available for the 2840 buffer, which may be shared by up to four 2250 graphics devices. The above-mentioned exits can be used to establish a control mechanism that restricts to users of the same security classification, 2250s that are connected to the same 2840.
5. Although APF-authorized programs bypass the usual system control validity checks, in any user interfaces they must provide validity check mechanisms the same as or equivalent to the system validity checks. The effectiveness of APF is dependent on this. However, it may be impractical to provide these checks for some APF-authorized programs, and in these cases their use should be restricted by placing them in APF-authorized password protected libraries. The passwords for these libraries should then be provided only to individuals authorized to use the programs. Installations might consider using this control mechanism for the IEHINIT and IEHDASDR SCP utility routines.
6. If there are two modules with the same name on different authorized libraries and an authorized program attempts to load one of these, it could get the other (either accidentally or deliberately with JOBLIB, STEPLIB, or concatenation). This type of exposure requires the additional restriction that duplicate module names not be permitted across authorized libraries.

Conversion Considerations

This chapter describes significant differences and incompatibilities between OS/VS2 Release 2 and:

- OS/MVT
- OS/VS2 Release 1.

It is divided into the following sections:

- SMF considerations.
- Job control language.
- Operator commands.
- Time sharing considerations.
- Data set conversion.
- Program conversion.
- Multiprocessing considerations

SMF Considerations

System management facilities (SMF) collect and record system and job information. The information can be used for job accounting and management information reports.

Most of the changes to SMF records and exits for VS2 Release 2 have been made to support new system control program design. Because these changes could affect existing data reduction routines, all records and exits should be studied to determine what effect the changes might have.

The following topics describe the changes that have been made to SMF for VS2 Release 2.

Differences from VS2 Release 1 System Management Facilities

The most significant differences from VS2 Release 1 are:

- SMF records contain additional accounting information to record new system environmental characteristics.
- New SMF exits from the control program are provided.
- SMF supports an output limiting facility.
- The SMF data sets, SYS1.MANX and SYS.MANY, must be cataloged.
- The SMFDEFLT list in SYS1.PARMLIB has been changed. Multiple SMF SYS1.PARMLIB members may be maintained. Their names are of the form SMFPRMxx. Some of the parameters in these members are also different. See the topic “IBM-Created Lists” in the chapter “Defining the System” for a more detailed description.
- The same SYS1.PARMLIB member is applicable to background and foreground use.

Differences from MVT System Management Facilities

In addition to the preceding, the most significant differences from MVT are:

- SMF is a standard facility of VS2.
- SMF does not support tape for recording SMF data.

SMF Record Differences for VS2 Release 2

Figure 17 indicates how the SMF records have changed for VS2 Release 2. The following topics explain why some of the more significant changes were made.

New System Resources Manager Data

SMF record types 4, 5, 34, and 35 (step termination, job termination, time-sharing step termination and LOGOFF) provide new information obtained from system resources management routines. Fields in record types 4 and 34 are now used to record step termination data such as VIO page-ins, VIO page-outs, total page-ins, and total page-outs. Fields in record types 5 and 35 are used to record job termination data such as:

- Total number of service units used by a background job (type 5) or a terminal job (type 35).

- The total amount of time that the background job was active (type 5).
- The total amount of time that all transactions were active (type 35).
- The total number of transactions.

This information can be used to approximate TSO terminal response time for a session, and for adjusting IPS values in an attempt to obtain desired response time. (See the chapter “Directing the Use of System Resources” for a description of an installation performance specification (IPS) and service units.)

SMF Recording Replaced by MF/1

Two records, type 1 (wait time) and type 12 (end-of-day), are no longer supported because the system activity measurement facility (MF/1) now records the data in place of SMF. The information from record types 1 and 12 is now in record types 70 (CPU activity) and 71 (paging activity), although record types 70 and 71 are formatted differently from record types 1 and 12. Furthermore, MF/1 records are only written to the SMF data set when MF/1 is operational. Therefore, an installation that wants to use the information that was in record types 1 and 12 should:

- either use exit IEFU83 to obtain the data from the new record types and put it in a form for use by existing data reduction routines,
- or rewrite data reduction routines to recognize the new records and process the new format.

MF/1 processing must also be started so that the data will be gathered and recorded.

Three other new records (types 72, 73, and 74) have been created to record MF/1 data. See the topic “Tracing” in the chapter “Directing the Use of System Resources” for a more detailed description of their contents.

Record Type	Differences Between VS2 Release 2 and VS2 Release 1	Differences Between VS2 Release 2 and VS1	Differences Between VS2 Release 2 and MVT	Differences Between VS2 Release 2 and MFT
0	*****	*****	Virtual storage statistics added	Virtual storage statistics added
1	_____	_____	_____	_____
2	*****	*****	*****	*****
3	*****	*****	*****	*****
4	SYSIN/SYSOUT EXCP counts eliminated Virtual storage statistics added VIO statistics added Paging statistics added	Virtual storage statistics added VIO statistics added Paging statistics added	SYSIN/SYSOUT EXCP counts eliminated Virtual storage statistics added VIO statistics added Paging statistics added	SYSIN/SYSOUT EXCP counts eliminated Virtual storage statistics added VIO statistics added Paging statistics added
5	System resource manager statistics added	System resources manager statistics added	System resources manager statistics added	System resources manager statistics added
6	Job entry subsystem change	Job entry subsystem change	Job entry subsystem change	Job entry subsystem change
7	*****	*****	*****	*****
8	*****	*****	*****	*****
9	*****	*****	MP statistics eliminated	*****
10	*****	*****	*****	*****
11	*****	*****	MP statistics eliminated	*****
12	_____	_____	_____	_____
13	N/A	_____	N/A	_____
14	VIO statistics added	VIO statistics added	VIO statistics added	VIO statistics added
15	VIO statistics added	VIO statistics added	VIO statistics added	VIO statistics added
17	*****	*****	*****	*****
18	*****	*****	*****	*****
19	*****	*****	*****	*****
20	*****	*****	*****	*****
21	*****	*****	*****	*****
22	New record for MP statistics	New record for MP statistics	New record for MP statistics	New record for MP statistics
26	Job entry subsystem change	New record for job entry subsystem statistics	New record for job entry subsystem statistics	New record for job entry subsystem statistics
30	_____	N/A	_____	N/A
31	This record is now written when the MODIFY TCAM command with the TS option is issued	New record for time sharing statistics	This record is now written when the MODIFY TCAM command with the TS option is issued	New record for time sharing statistics
32	_____	N/A	_____	N/A
33	_____	N/A	_____	N/A

Figure 17. Changes to SMF Records for VS2 Release 2 (Part 1 of 2)

Record Type	Differences Between VS2 Release 2 and VS2 Release 1	Differences Between VS2 Release 2 and VS1	Differences Between VS2 Release 2 and MVT	Differences Between VS2 Release 2 and MFT
34	Virtual storage statistics added VIO statistics added Paging statistics added	New record for time sharing statistics	Virtual storage statistics added VIO statistics added Paging statistics added	New record for time sharing statistics
35	System resources manager statistics added	New record for time sharing statistics	System resources manager statistics added	New record for time sharing statistics
38	N/A	N/A	————	N/A
40	*****	New record for dynamic allocation statistics	*****	New record for dynamic allocation statistics
41	————	N/A	————	N/A
42	————	N/A	————	N/A
43	New record for job entry subsystem statistics	Job entry subsystem change	New record for job entry subsystem statistics	New record for job entry subsystem statistics
44	N/A	————	N/A	N/A
45	New record for job entry subsystem statistics	New record for job entry subsystem statistics	New record for job entry subsystem statistics	New record for job entry subsystem statistics
47	New record for job entry subsystem statistics	Job entry subsystem change	New record for job entry subsystem statistics	New record for job entry subsystem statistics
48	New record for job entry subsystem statistics	Job entry subsystem change	New record for job entry subsystem statistics	New record for job entry subsystem statistics
49	N/A	————	N/A	N/A
62	*****	*****	New record for VSAM statistics	New record for VSAM statistics
63	*****	*****	New record for VSAM statistics	New record for VSAM statistics
64	*****	*****	New record for VSAM statistics	New record for VSAM statistics
67	*****	*****	New record for VSAM statistics	New record for VSAM statistics
68	*****	*****	New record for VSAM statistics	New record for VSAM statistics
69	*****	*****	New record for VSAM statistics	New record for VSAM statistics
70	New record for MF/1	New record for MF/1	New record for MF/1	New record for MF/1
71	New record for MF/1	New record for MF/1	New record for MF/1	New record for MF/1
72	New record for MF/1	New record for MF/1	New record for MF/1	New record for MF/1
73	New record for MF/1	New record for MF/1	New record for MF/1	New record for MF/1
74	New record for MF/1	New record for MF/1	New record for MF/1	New record for MF/1
————	Record not supported			
*****	Same statistics			
N/A	Not applicable			

Figure 17. Changes to SMF Records for VS2 Release 2 (Part 2 of 2)

New Storage Data

Three fields in both record types 4 and 34 now provide new step level accounting information such as:

- Either the size of the V=R address area specified on an ADDRSPC=REAL request, or the total size of the user's private address space for an ADDRSPC=VIRT request.
- The amount of the user's private address space (assigned from the high address) used for the local system queue area (LSQA), the scheduler work area (SWA), and subpools 229 and 230.
- The amount of the user's private address space (assigned from the low address) used for the user's programs.

Job Entry Subsystem Statistics

Some SMF records have been modified and others created to record statistics provided by the job entry subsystem. They are:

- The type 6 (output writer) record contains new information, because the job entry subsystem handles all system output processing.
- The type 26 (job summary) record is produced when a job is ready to be purged from the system. It contains information about the job entry subsystem's processing of the job.
- Record types 43, 45, 47, and 48 record information about the starting or stopping of a subsystem or a subsystem event (e.g., SIGNON).

Multiprocessing Data

The type 22 record is new in VS2 Release 2. It indicates the status of resources unique to a multiprocessing configuration. It serves the same function for CPUs, channels, and storage as record types 8, 9, 10, and 11 do for I/O devices.

Virtual Storage Access Method Data

The virtual storage access method (VSAM) generates new records (types 64, 67, and 68) that contain information similar to existing data management SMF records. VSAM also records statistics when a VSAM data set is opened, allocated, and freed (record types 62, 63, and 69).

Note: Record types 63 and 67 are generated for VSAM recovery processing.

Time Sharing Record Elimination

The following record types are not supported in VS2 Release 2:

- 30-- Start TS
- 32-- Driver
- 33-- Driver Modify
- 41-- Modify TS
- 42-- Stop TS

These records were associated with starting and stopping TSO and with the TSO driver. These functions have been replaced by the MODIFY TCAM command and the system resources manager in VS2 Release 2.

Changes to SMF Exits for VS2 Release 2

Figure 18 indicates the changes that have been made to SMF exits in other systems for VS2 Release 2.

Exit	Difference Between VS2 Release 2 and VS2 Release 1	Difference Between VS2 Release 2 and VS1	Difference Between VS2 Release 2 and MVT	Difference Between VS2 Release 2 and MFT
IEFUJV	This exit is also taken after the internal text is created	This exit is also taken after the internal text is created	This exit is also taken after the internal text is created	This exit is also taken after the internal text is created
IEFUIV	N/A	——	N/A	N/A
IEFUJI	*****	*****	*****	*****
IEFUSI	*****	*****	*****	*****
IEFUSO	New exit for output limiting facility	*****	*****	*****
IEFACTRT	*****	*****	*****	*****
IEFUTL	This exit can now be taken from the foreground There is a new interface to extend the time limit in micro seconds	This exit can now be taken from the foreground There is a new interface to extend the time limit in micro seconds	This exit can now be taken from the foreground There is a new interface to extend the time limit in micro seconds	This exit can now be taken from the foreground There is a new interface to extend the time limit in micro seconds
IEFU83	*****	New exit for selecting records to be written	New exit for selecting records to be written	New exit for selecting records to be written
IEFUJP	New exit receives control when job is ready to be purged from the system	Change in exit parameters	New exit receives control when job is ready to be purged from the system	New exit receives control when job is ready to be purged from the system
—— Exit not supported ***** Unchanged N/A Not applicable				

Figure 18. Changes to SMF Exits for VS2 Release 2

Job Control Language

Job control language statements contain information that is needed to control the processing of jobs.

This section describes JCL facilities, parameters, and subparameters that are either new or have changed from both VS2 Release 1 to VS2 Release 2 and MVT to VS2 Release 2. Figure 19 summarizes the more significant JCL changes.

Changes to VS2 Release 1 JCL

The following topics describe the differences from VS2 Release 1 JCL.

Performance Group Specification

The **PERFORM** parameter is new for VS2 Release 2. It is used to associate a job or job step with any of the several performance groups specified by installation-supplied performance group definitions. Each performance group definition specifies a rate for providing system resources to terminal jobs or batch jobs or job steps under varying conditions (see the chapter: "System Resources Management").

- When **PERFORM=groupnr** is coded on a **JOB** statement, that job receives system resources at the rate defined by the performance group identified by 'groupnr' (a number between 1 and 255).
- When **PERFORM=groupnr** is coded on an **EXEC** statement, the group number value applies only to the job step specified on the **EXEC** statement.
- When **PERFORM.procstep=groupnr** is coded on an **EXEC** statement, the group number value applies to the specified step within the procedure invoked by the **EXEC** statement.

If the **PERFORM** parameter is coded on the **JOB** statement, its value overrides any value for **PERFORM** parameters that may be specified on the **EXEC** statements for that job. If a **PERFORM** parameter is specified for a procedure, its value overrides any value for **PERFORM** parameters that may be specified on **EXEC** statements within the procedure.

If the **PERFORM** parameter is not coded, the default is 1 for background jobs and 2 for foreground users. Therefore if the previous job control language for a job is not changed, the job will be associated with one of only two possible performance group definitions.

Assigning Job Selection Priority

The **PRTY** parameter of the **JOB** statement is ignored in VS2 Release 2 with the **JES2** job entry subsystem. Instead, job selection priority must be specified on a **JES2** control card.

Statement	Parameter	Change	
JOB	ADDRSPC	●	Assigns pageable or nonpageable storage space, and defaults for pageable storage.
	PERFORM	⊙	Associates jobs with performance groups.
	PRTY	⊙	This parameter is ignored when the JES2 job entry subsystem is used. Priority values must be specified on a subsystem control card.
	REGION	⊙	Limits the size of variable GETMAIN requests.
	ROLL	●	This parameter is ignored.
EXEC	ADDRSPC	●	Assigns pageable or nonpageable storage space, and defaults for pageable storage.
	DPRTY	⊙	The default is the value of the APG priority.
	DYNAMNBR	⊙	Limits the number of data sets that a job step may hold for reuse.
	PERFORM	⊙	Associates job steps with performance groups.
	REGION	⊙	Limits the size of variable GETMAIN requests.
	ROLL	●	This parameter is ignored; the job is not failed.
DD	AFF	⊙	This parameter is ignored.
	COPIES	⊙	Specifies multiple copies of an output data set.
	DCB	●	EROPT subparameter specifies the BTAM online terminal test option.
		⊙	FRID subparameter specifies the SYS1.IMAGELIB member to be used to interpret documents read by a 3886 OCR device.
		●	HIARCHY subparameter is ignored.
	DEST	⊙	Specifies the destination for an output data set.
	DYNAM	⊙	Total is added to DYNAMNBR value, if any.
	OUTLIM	○	Limits the number of records to be written to an output data set.
	FREE	⊙	Indicates when a data set is to be freed.
	RETAIN	⊙	This parameter is ignored.
	SEP	⊙	This parameter is ignored.
	SPACE	⊙	Has a new default value for VIO data sets.
	SPLIT	⊙	Converted to an equivalent SPACE parameter.
	SUBALLOC	⊙	Converted to an equivalent SPACE parameter.
UNIT	⊙	The SEP subparameter is ignored.	
<ul style="list-style-type: none"> ● New or changed from MVT ○ New or changed from VS2 Release 1 ⊙ New or changed from MVT and VS2 Release 1 			

Figure 19. VS2 Release 2 JCL Changes

Specifying Data Sets in Paging Space (Virtual I/O Data Sets)

A new Stage I system generation parameter, **VIO**, may be specified on the **UNITNAME** macro instruction to identify unit names that may be used to specify data sets that are to receive virtual I/O (VIO) processing. Any single volume temporary data set defined by a **DD** statement with a unit name for which **VIO=YES** was coded at system generation time, is processed as a virtual I/O data set.

The **DD** statement that defines the temporary data set must meet the following requirements if the temporary data set is to be allocated to the system's paging space:

- The **DSNAME** parameter must not be specified unless an **&** name value is supplied for it. (Data sets with non-system generated names are not allocated to paging space).
- The data set disposition must be new.
- The **UNIT** parameter must specify a unit name that indicates system paging space. (**VIO=YES** must have been specified for that unit name at system generation time.)
- Indexed sequential data set organization must not be specified on the **DSORG** subparameter of the **DCB** parameter.
- No volume serial number may be specified.

If the **SPACE** parameter is not coded for virtual I/O data sets, the default value is 10 primary and 50 secondary blocks with an average block length of 1000.

Requesting Data Set Freeing at Data Set Close Time

The **FREE** parameter of the **DD** statement is new in VS2 Release 2. It is used to specify that a data set is to be freed when that data set is closed.

If **FREE=END** is specified, the data set is freed at the end of the job step. If **FREE=CLOSE** is specified, the data set is freed when the **CLOSE** macro instruction is issued for it.

If the **FREE** parameter is not coded, the default value is **END**. Therefore, if the current job control language for a job is not changed, all data sets will be released at the end of the job step.

If **FREE=CLOSE** is specified for a **SYSOUT** data set, the data set will be sent to the job entry subsystem when the **CLOSE** macro instruction is issued.

Note: **FREE=CLOSE** should not be specified for a data set that is opened and closed more than once during a job step.

Specifying the Number of Resources That May Be Held for Reuse

The **DYNAMNBR** parameter of the **EXEC** statement is new in VS2 Release 2. While it extends and replaces the function of the **DD DYNAM** statement in MVT and VS2 Release 1, it does not specify a value that limits the number of dynamic allocations allowed to a terminal job. Rather, it specifies a value that is used to control the number of not-in-use data sets that may be allocated to a job step. The **DYNAMNBR** parameter value is added to the number of **DD** statements in the job step to determine the number of data sets with the not-in-use attribute that a job step may hold.

The value of the DYNAMNBR parameter must be within the range of 0 to 1,635. If the DYNAMNBR parameter is not coded, the default value is 0.

Requesting Multiple Copies of an Output Data Set

The COPIES parameter of the DD statement is new in VS2 Release 2. It specifies the number of copies of an output data set that are to be produced by the job entry subsystem. The COPIES parameter may be coded only when the SYSOUT parameter is included on the DD statement.

Directing the Destination of an Output Data Set

The DEST parameter of the DD statement is new in VS2 Release 2. It specifies that an output data set is to be routed to a remote station by the job entry subsystem. The DEST parameter may be coded only when the SYSOUT parameter is included on the DD statement.

Format Record Identification for the 3886 OCR Device

A new subparameter (FRID) of the DCB parameter of the DD statement is used to specify the type of document that is supplied as input to a 3886 optical character recognition (OCR) device. The FRID subparameter indicates the member of the SYS1.IMAGELIB data set that is to be used for interpreting the document to be read by the 3886 OCR device.

Output Limit Specification

The OUTLIM parameter on the DD statement is used to specify the limit for the number of logical records to be written to an output data set. It is coded the same way as the MVT OUTLIM parameter.

Increased Data Set Definition Capability

In VS2 Release 2, as many as 1,635 DD statements may be specified for each job step. In both MVT and VS2 Release 1, the limit was 255 DD statements per step.

Increased Volume Sharing Capability

The following additional types of volumes may be shared between jobs and job steps in VS2 Release 2:

- Private direct access volumes mounted in response to nonspecific requests.
- Direct access volumes for which deferred mounting has been requested.

Interpretation of the REGION Parameter

In VS2 Release 2, the address space available to an individual job (user private address space) is restricted only by the size of the individual virtual address space (16 megabytes) minus the space required for system control program use. Therefore, for jobs that specify ADDRSPC=VIRT (requesting V=V storage), the REGION parameter is not used to specify the size of the address space available to the individual job. It is, however, used to specify the maximum value for use by the GETMAIN service routines in servicing variable-length GETMAIN requests. When a variable-length GETMAIN is issued that would cause the address space allocated to a user to exceed the value of the REGION parameter, the user is limited to the REGION parameter value. For jobs that

specify `ADDRSPC=REAL` (requesting `V=R` storage), the value specified is used to limit the size of the address space available to the individual job.

Job Step Dispatching Priority Default Value

In VS2 Release 2, the value of the APG priority will be used as the default dispatching priority for a step when the `DPRTY` parameter is not specified on the `EXEC` statement. In MVT and VS2 Release 1, the default dispatching priority for the step was the job's dispatching priority.

Changes to MVT JCL

In addition to the differences described above, two JCL parameters are new to current MVT users.

Requesting Storage for Program Execution

The `ADDRSPC` parameter is new in VS. It is used to assign pageable or nonpageable storage space to a job or a job step. It may be specified as either `ADDRSPC=VIRT` or `ADDRSPC=REAL`:

- If `ADDRSPC=VIRT` is specified, the job is assigned to pageable (`V=V`) storage.
- If `ADDRSPC=REAL` is specified, the job is assigned to nonpageable (`V=R`) storage.
- If the `ADDRSPC` parameter is not coded, the default value is `VIRT`. Therefore, if the current job control language for a job is not changed, the job will be executed in pageable storage. If the job must be executed in nonpageable storage, `ADDRSPC=REAL` must be specified.

If the `ADDRSPC` parameter is coded on the `JOB` statement, its value overrides any value in the `ADDRSPC` parameters that may be specified in the `EXEC` statements for that job.

Specifying the BTAM Online Terminal Test Option

In VS2, a new value can be specified for the `EROPT` subparameter of the `DCB` parameter of the `DD` statement. If `EROPT=T` is coded, it indicates a request for the basic telecommunication access method (BTAM) online terminal test option.

JES2 JCL Considerations

The JES2 job entry subsystem processes JCL statements in a manner similar to MVT/HASP and VS2 Release 1/HASP. The JCL considerations for JES2 therefore fall into two categories, depending upon whether or not an installation currently has a HASP system. In either case, however, no syntax changes to existing JCL will be required.

JES2 JCL for Non-HASP Users

Three significant points for installations that are unfamiliar with HASP are:

- Extended JCL facilities are provided by JES2 control cards.
- The `PRTY` parameter on the `JOB` statement cannot be used to specify job priority and will be ignored if coded.

- The cataloged procedure to be used for job conversion can no longer be associated with an input stream but rather with a job class.

JES2 Control Cards: The JES2 control cards are identified by a /* in columns 1 and 2 and a JES2 control verb beginning in column 3. They are used to:

- Supply information for the JES2 separator page and card.
- Identify a job selection priority (for queuing the job).
- Supply information for calculating a job selection priority if such a priority is not explicitly stated.
- Identify volumes required by the job, so that the operator will cause the job to be held until the required volumes become available.
- Route system output data sets to the local printer or punch, if the job was submitted from a remote work station.
- Send a message to the central computer operator.
- Specify the DDNAME of the cataloged procedure library to be used for conversion of the job's JCL.
- Specify extended output processing setup options to be associated with data set groups.

Assigning Job Selection Priority: A JES2 control card is used to assign a selection priority to a job. If this control card is not included, the job entry subsystem assigns a priority determined either by the estimated execution time and estimated number of lines of code (values contained on another JES2 control card) or by the JES2 generation default values.

If the PRTY parameter is coded on a JOB statement, it will be ignored.

JES2 JCL for HASP Users

The two most significant points are:

- In addition to the job copies and job routing facilities, it is now possible to specify data set copies and data set routing.
- Two new control cards have been created.

/*JOBPARM Control Card: The /*JOBPARM control card was new in HASP for VS2 Release 1 and is coded similarly for JES2. It specifies job-related parameters that are used by the JES2 job entry subsystem to:

- Establish the default job selection priority.
- Control the job's processing.
- Build the JES2 separator page and card.

These parameters extend the information that previously could be specified only on a JOB statement accounting information parameter by adding information such as the name of the cataloged procedure library to be used for JCL conversion.

/*OUTPUT Control Card: The /*OUTPUT control card was new in HASP for VS2 Release 1, and is coded similarly for JES2. It specifies characteristics and/or options for groups of output data sets.

It is used in conjunction with the form number subparameter in the SYSOUT parameter of the DD statement.

Any system output data set whose DD statement form number subparameter matches the FORMS code specified on as /*OUTPUT control card, will be processed according to the specifications of the /*OUTPUT control card.

Operator Command Language

Operator commands are statements to the control program issued via a console device or in an input stream. They are used to request information from the control program, alter normal operations, initiate new operations, or terminate existing operations.

In MVT all data coming into and out of the system was handled by system routines and, consequently, controlled by system operator commands. With the introduction of a job entry subsystem, the handling of the reading, writing, spooling, and queueing of data is a function of the job entry subsystem in control. All data flow prior to and immediately after execution is handled by a job entry subsystem and controlled by its unique operator commands. Hence in OS/VS2 Release 2, some of the historical system commands and the functions they performed will be replaced by operator commands of the job entry subsystem.

This section does not describe all of the operator commands needed to run jobs in VS2 Release 2. Rather, it describes the manner in which the operator command language has been changed for this release. Also included is a brief functional description of those commands related to JES2, the job entry subsystem that is shipped with the initial version of Release 2.

Other areas of support include multiprocessing, VTAM, time sharing, and resource management. All have caused changes to the command language. Support for miscellaneous functions is described at the end of the section.

Changes to Operator Commands

Figure 20 contains operator commands and operands that have changed for VS2 Release 2. Functions that are no longer supported by operator commands are usually controlled through the job entry subsystem. The specific job entry subsystem commands are not listed, but their general functional description is included under "JES Support". Note that a command may have changes supporting several separate functions.

Command	Operand		Change
CANCEL	USERID	⊙	Used to cancel jobs in execution (only).
	u = userid	⊙	Cancels time sharing terminal session.
	IN OUT ALL BRDR	⊙	No longer supported. The job entry subsystem provides functionally equivalent support.
CHNGDUMP	SET	⊙	Overrides the SYS1.PARMLIB defaults for SYSABEND or SYSUDUMP storage dumps.
	DEL	⊙	Eliminates the SET operand overrides.
CONTROL	M	⊙	Replaced by the T operand.
	T	⊙	Supports the TRACK command, which has replaced MONITOR A.
DISPLAY	JOBS	⊙	Active batch-type jobs, and started tasks and jobs.
	TS	⊙	Active time-sharing jobs. Functionally equivalent to DISPLAY USER.
	INIT	⊙	Active initiators.
	A		All the above.
	jobname	⊙	No longer supported. The job entry subsystem provides functionally equivalent support.
	N	⊙	
	Q	⊙	
	SQA	⊙	
USER	⊙	Replaced by DISPLAY TS.	
R [,LIST]	⊙	The optional LIST operand shows the first 67 characters of outstanding reply messages.	
MATRIX	○	Used with multiprocessing to selectively display the status of such system resources as CPUs, channels, devices, and storage. Generally compatible with M65 MP.	
NET	⊙	Monitors the status of VTAM.	
HALT	EOD	•	Initiates closedown of entire operating system.
	TP	•	Initiates closedown of TCAM.
	NET	⊙	Initiates closedown of VTAM.
HOLD	Q	⊙	No longer supported. The job entry subsystem provides functionally equivalent support.
	Q=LIST	⊙	
	jobname	⊙	
	TP=stationname		Used to intercept a TCAM station.
MODIFY	procname		Used for jobs started from the system console or the system input stream.
	BRDRQ	⊙	No longer supported. The job entry subsystem provides functionally equivalent support.
	CLASS	⊙	
	JOBCLASS	⊙	
	OUTCLASS	⊙	
	PAUSE	⊙	

Figure 20. Changes to Operator Commands for VS2 Release 2 (Part 1 of 3)

Command	Operand		Change
MODIFY (Continued)	DRIVER	⊙	No longer supported because time sharing has been integrated into the system.
	HOLD	⊙	
	REGSIZE	⊙	
	SMF	⊙	
	SUBMIT	⊙	
	USERS	⊙	
	NET	⊙	Optionally changes some of the operating characteristics of VTAM.
MONITOR	A	⊙	Functionally replaced by the TRACK command.
MOUNT	LSQA=	⊙	This keyword can no longer be included. LSQA is now dynamically expandable during execution. Including this keyword will cause the command to be rejected.
MSGRT	MN=A	⊙	This operand is no longer supported. MSGRT establishes default routing for TRACK, STOPTR, DISPLAY and CONTROL commands.
QUIESCE (Q)		⊙	The short form of the command (Q) is not supported. It is not necessarily issued before every reconfiguration, as required by M65 MP.
RELEASE	Q	⊙	No longer supported. The job entry subsystem provides functionally equivalent support.
	Q= LIST	⊙	
	jobname	⊙	
	TP=stationname		Used to release an intercepted TCAM station.
RESET	jobname, PERFORM=	⊙	Changes the performance group for the jobstep currently running.
	CLASS	⊙	No longer supported. The job entry subsystem provides functionally equivalent support.
	OUT	⊙	
	PRTY	⊙	
SEND	SAVE	⊙	The message is not sent until the next time the userid is logged on -- even if it is currently logged on.
	OPER	⊙	Specifies that the message text is to be sent to a specific console.
	BRDCAST	⊙	Specifies that the message is to be sent to all active consoles.
	CN	⊙	Specifies that the message is to be sent to a specific operator's console. Messages sent via the OPER or CN keywords will have respectively OPER or CN appended to the message.
SET	IPS= nn	⊙	This optional parameter can specify a member of SYS1.PARMLIB which contains the system resources manager parameters under which the system is to run.
	Q PROC AUTO	⊙	No longer supported. The job entry subsystem provides functionally equivalent support.
START	LSQA=	⊙	This keyword is no longer supported since LSQA is dynamically expandable during execution. If this keyword is included, the command will be rejected.
	NET	⊙	Initiates startup of VTAM.

Figure 20. Changes to Operator Commands for VS2 Release 2 (Part 2 of 3)

Command	Operand		Change
STOPTR		●	Stops the periodic displays activated through the TRACK command.
TRACK	JOB TS INIT A	●	Requests periodic displays of the number and names of jobs and users.
		●	As opposed to MONITOR A, each graphics console is allowed a TRACK command. The TRACK command is rejected from a nonCRT console or TSO terminal.
VARY	CPU	●	Changes the status of a CPU to ONLINE or OFFLINE. Generally compatible with M65MP.
	CHANNEL	●	Generally compatible with M65MP. The channel range is now 0 to F. A channel with an online unallocated tape drive can be taken offline by VARY.
	PATH	○	Causes a path (logical connection) to be logically made available or removed from the system.
	STORAGE	●	Changes the status of an area of real storage to ONLINE or OFFLINE. M65MP Box number designations are not allowed. Address ranges may be entered in multiples of K. (K = 1024)
	NET	●	Used to add or remove VTAM terminals.
WRITELOG	START	●	Reinitiates the log function.
<ul style="list-style-type: none"> ● New or changed from MVT ○ New or changed from VS2 Release 1 ● New or changed from MVT and VS2 Release 1 			

Figure 20. Changes to Operator Commands for VS2 Release 2 (Part 3 of 3)

Multiprocessing Support

The MVT VARY commands (VARY STORAGE, VARY PATH, VARY CUPU, VARY CHANNEL) are supported in VS2 Release 2. There is no MP support in VS2 Release 1.

The SET command is modified to identify the CPU which has the acceptable clock value, or to cause the time zone constant to be restored to its IPL value.

The QUIESCE command suspends system activity, i.e., all CPUs are placed in manual state.

The DISPLAY MATRIX command is generally compatible with M65MP. The syntax has been changed for compatibility with TSO command language, and a new operand displays the highest real address. Two CPUs (0 and 1) and 16 channels (0-F) are supported.

VTAM Support

The NET operand on the DISPLAY, HALT, START, and VARY commands identify them with VTAM.

Time Sharing (TSO) Support

Since time sharing is now integrated into the system, MODIFY command operands such as USERS, SUBMIT, REGSIZE, DRIVER, SMF, and HOLD are not supported. Time sharing is started and stopped by the operator issuing a MODIFY TCAM with a TS=START or STOP operand. There is no TSO driver to modify. Its functions are supported by the system resources manager.

The background reader (BRDR) support by the START, STOP, and CANCEL commands is dropped since there is no longer a background reader. Reader functions are supported by the job entry subsystem.

System Resources Management Support

START TS and the driver parameters of MODIFY are no longer supported. The functions of the driver are performed by the system resources manager.

The IPS parameter of the SET command specifies the member of SYS1.PARMLIB which contains the parameters which are used by the system resources manager.

Job Entry Subsystem Support

In VS2 Release 2, the job entry subsystem (JES2 or JES3) functionally replaces the SYS1.SYSJOBQE data set and the OS readers and writers. Support is therefore removed for operands associated with these functions. Commands affected by these changes are CANCEL, DISPLAY, HOLD, MODIFY, RELEASE, RESET, and SET for queues, and SET, START, and STOP for readers and writers.

The job entry subsystem available with the initial version of VS2 Release 2 is JES2, which is generally compatible with HASP II.

Some of the operator commands entered will be JES commands. For the most part, the JES2 operator commands are identical to those of the historical system commands, but with a special character identifier preceding the verb.

A brief description of JES2 command functions follows.

Job Related Functions: The job entry subsystem has the ability to:

- Release all jobs or specific jobs from the hold queue.
- Display, for each active job entered through JES2, job related information (e.g., that the job is either executing, being purged, or on a device).
- Display the number of jobs queued for special forms processing for either remote or local, or both.
- Display jobs queued for JES2 services, and which services they are waiting for.
- To place jobs currently in the system and entered through JES2, on the hold queue.
- To cancel specific jobs (JES2 will issue SVC 34 to cancel jobs in execution in the system). Jobs cancelled will have their current activity deleted and will be scheduled for purge processing.
- Display for specific jobs entered through JES2, their current status (e.g., the job is in system execution, or queued for print processing).

- To schedule a job for purge processing after its current activity has completed.
- To set the new base number for automatic job number assignments (i.e., all jobs entered through JES2 will be assigned unique job identifiers starting from the operator-set base).

Device List Functions: Device list commands perform the following functions:

- Backspace a printer device.
- Restart a current activity on an output device, (i.e., the current activity on the designated device will be terminated and, in the case of printer/punch devices, the activity will be returned to the appropriate queue in order of priority and made eligible for selection. In the case of remote job entry lines, JES2 will, upon completion of the current line I/O, cancel all activities on the line).
- Forward space a printer device.
- Interrupt current activity on a printer device, (i.e., the current activity on the designated printer(s) will, if active, be checkpointed and terminated, with the activity requeued on the appropriate JES2 queue).
- Repeat a current activity on a printer/punch device, (i.e., after the current activity has been completed, the activity will be requeued on the appropriate queue to schedule a repetition of the activity).
- START/STOP a device.
- Place all jobs to be read into JES2 from a specified reader, on the hold queue.
- Set a printer or punch to handle different characteristics (e.g. forms, FCB loads, and UCS trains/chains).

System Functions: System commands which control the ability of JES2 to process jobs through the system may be broken down into two groups:

- Initiator Commands -- those commands which control the actual selection and submission of jobs from JES2 to the system for processing.
- JES Commands -- those commands which control the ability of JES2 to process jobs for any function.

Initiator Commands: Initiator commands perform the following function:

- Display for all JES2 initiators, their status (e.g., the job classes, their processing, whether they are active, inactive (waiting for new jobs), or draining.
- Stop an initiator.
- Start an initiator.
- Set the job classes of jobs that the initiator will handle.
- With respect to jobs read into JES2 and on an input queue, change that job's queuing priority and job class.

JES START/STOP Commands: With JES2 START/STOP commands, the facility exists:

- For an operator to stop JES2, which will stop all JES2 job processing and prevent any new function from beginning.
- For an operator to start JES2.

Miscellaneous Display Function: Display functions include:

- The ability to display the device addresses and volume serial numbers of all online direct access storage devices.
- The ability to display the status (JES2) of all devices (readers, printers, punches) on a remote line.
- The ability to display the status (JES2) of all JES2 controlled, non-direct access devices attached to the local system as well as the corresponding hardware address of the device.

Remote Job Entry Function: With remote job support the operator has the ability to:

- Display a message at a remote terminal, whether logged-on, signed-on, or inactive.
- Route a job's output to a remote terminal, whether that terminal is logged-on, signed-on, or inactive.

Miscellaneous Support

General support for VS2 Release 2 is provided with several new functions or commands.

Displays: The DISPLAY, TRACK and STOPTR commands are used to request or change the display of information concerning currently active tasks. Only jobs running under initiators and started tasks are displayed.

Each graphics console is allowed a TRACK command, which functionally replaces the MONITOR A command. TRACK will be rejected from a console with no CRT. The TRACK and STOPTR commands are used for timer or interval driven periodic displays, as distinguished from the event driven displays of the MONITOR and STOPMN commands.

DISPLAY USER, N, Q, SQA, and jobname, are not supported. These functions are supported through the job entry subsystem.

The M operand of the CONTROL command is replaced by the T operand -- an equivalent function. CONTROL T affects the TRACK command instead of the MONITOR A command.

Dumps: Temporary changes to the defaults defined in SYS1.PARMLIB may be specified or eliminated with the CHNGDUMP command -- new in this release.

Messages: The SEND command is modified to allow the message to be entered into the broadcast data set for a specified user, or to be sent to an operator console.

LSQA: The START and MOUNT commands no longer support the LSQA keyword. The local system queue area (LSQA) is dynamically expandable rather than fixed in size.

System Log: The system log function is started automatically during IPL, and can now be restarted after being closed.

Time Sharing Considerations

Changes to the time sharing (TSO) facilities for VS2 Release 2 support system resources management, VSAM, allocation, the job entry subsystem, the EDIT function, accounting, and the integration of time sharing into the system. Changes to the time sharing commands of MVT and VS2 Release 1 are described in Figure 21.

System Resources Management Support

The system resources management routines allow an installation to group together into performance groups, user transactions with similar performance requirements. (See the chapter "Directing the Use of System Resources.")

An operand of the LOGON command allows the user to specify a performance group number.

The ACCOUNT facility specifies the performance groups that the time-sharing user is authorized to request at LOGON time.

VSAM Support

Both VSAM and non-VSAM data sets can exist in VS2 Release 2. A combination of the TSO utilities available in MVT and VS2 Release 1, and access method services are used to manipulate data sets in VS2 Release 2.

The following access methods services commands are fully supported for interactive use at terminals: DEFINE, ALTER, CNVTCAT, DELETE, LISTCAT, REPRO, PRINT, EXPORT, IMPORT, and VERIFY. These commands are discussed in the topic "Data Set Conversion" in this chapter.

The DELETE and LISTCAT utilities are not supported. Their functions are supported for the VSAM catalog, with commands of the same name, by access method services.

LISTALC and LISTDS support both VSAM and non-VSAM data sets.

The ALTER command provides the functions of PROTECT and RENAME. The latter two apply only to non-VSAM data sets.

Allocation Support

The ALLOCATE, FREE, and ATTRIB commands support dynamic allocation through device type and disposition specification, concatenation and deconcatenation of data sets, multivolume and multiunit requests, and setting DCB parameters relating to tape volumes.

ALLOCATE and PROFILE are added as subcommands of the EDIT command.

Command	Subcommand	Parameter		Change
ACCOUNT	ADD	DEST(stationid)/ <u>LOCAL</u>	⊙	Specifies the default destination for SYSOUT data sets defined by this user. The default is the central output facility.
		MOUNT/ <u>NOMOUNT</u>	⊙	Specifies that dynamic allocation requests can cause volume mounting.
	PERFORM(group)/ <u>NOPERFORM</u>	⊙	Specifies the performance group (if any) that the TSO user is authorized to request at LOGON.	
	CHANGE	The new parameters are the same as those listed under the ADD subcommand.	⊙	The functions are basically the same as described under the ADD subcommand, except that there are no default values.
	SYNC		⊙	Builds a new BROADCAST data set and synchronizes it with the UADS.
ALLOCATE		DATASET	⊙	Extended to allow concatenation of cataloged data sets.
		DEST (stationid) HOLD/ <u>NOHOLD</u>	⊙	These parameters support redirected SYSOUT. DEST allows the user to specify a workstation to which all SYSOUT data sets are directed. HOLD/NOHOLD enables the user to place the SYSOUT data sets on a hold queue for later processing.
		DUMMY AVBLOCK/TRACKS/ CYLINDERS RELEASE ROUND KEEP/DELETE/ CATALOG/ UNCATALOG	• • • ⊙ ⊙	These parameters and their functions are directly analogous to their JCL counterparts.
		UNIT (unit-type) UCOUNT (unit-count)/ <u>PARALLEL</u>	⊙ ⊙	These parameters provide device support. UNIT specifies the kind and maximum number of devices allocated to the file or data set. The (unit-type) may be an installation - defined group name, a generic device - type, or a specific device address. The unit count specifies the maximum number of devices to be simultaneously mounted for the file or data set.
		VOLUME MAXVOL POSITION PRIVATE LABEL VSEQ	⊙	Through these parameters, multivolume data sets are supported. The user can define the volume serial numbers, limit, relative position of the data set, and volume and label information. These parameters have corresponding JCL statements in MVT and VS2 Release 1.
		SYSOUT (class)	⊙	Class can be specified for SYSOUT.
				OPTCD RECFM BUFOFF DSORG LIMCT DEN TRTCH DIAGNS

Figure 21. Changes to Time Sharing Commands for VS2 Release 2 (Part 1 of 3)

Command	Subcommand	Parameter		Change
CANCEL/STATUS		jobname (jobid)	⊙	The jobid is optionally needed to uniquely identify a particular job, if more than one has been submitted with the same jobname.
EDIT	ALLOCATE	Same as command	⊙	ALLOCATE is added as an EDIT subcommand.
	SEND	Same as command	⊙	SEND is added as an EDIT subcommand.
	RUN	LIB (dslist)	⊙	Concatenated load libraries are supported for the RUN subcommand.
	DELETE (DEL)		⊙	The DELETE subcommand abbreviation is changed to DEL.
	END		⊙	The EDIT input data set is freed upon command termination.
	PROFILE	Same as command	⊙	PROFILE changes are available under EDIT.
EXEC		PROMPT	⊙	This parameter allows prompting during the execution of CLIST operations.
FREE		DEST (stationid)	⊙	These parameters support redirected SYSOUT. The option is provided to change the job entry subsystem user to which SYSOUT data sets are directed, and to override the current HOLD specification.
		HOLD/NOHOLD	⊙	
		KEEP/DELETE/ CATALOG/ UNCATALOG/ SYSOUT (class)	⊙	
LOGOFF		HOLD/ <u>DISCONNECT</u>	⊙	A telecommunications line may be held (i.e., not dropped) at LOGOFF time.
LOGON		<u>MOUNT/NOMOUNT</u>	⊙	Specifies that dynamic allocation requests can cause volume mounting.
		PERFORM/ <u>NOPERFORM</u>	⊙	Specifies the performance group (if any) that is authorized for this user.
		RECONNECT	⊙	Allows the user to log back onto the system and pick up the session -- within an installation - defined period of time.
OPERATOR	CANCEL		⊙	The jobname, unit address, and identifier operands are no longer supported. Their functions are handled by the job entry subsystem.
	DISPLAY	A jobname N = LIST Q = LIST USER	⊙	The parameters are no longer supported. Displaying job and queue data is a function of the job entry subsystem.
		R [,LIST]	⊙	Replaced by DISPLAY TS.
		TS [,LIST]	⊙	DISPLAY TS LIST provides information formerly obtained with DISPLAY A or DISPLAY USERS.
	MODIFY		⊙	No longer supported.
	MONITOR	DSNAME	⊙	The DSNAME and SPACE parameters are no longer supported for these subcommands.
STOPMN	SPACE	⊙		

Figure 21. Changes to Time Sharing Commands for VS2 Release 2 (Part 2 of 3)

Command	Subcommand	Parameter		Change
OPERATOR (continued)	SEND	SAVE	•	The message is placed in the broadcast data set, and will not be printed at the receiving terminal until the <u>next</u> LOGON -- even if the user is currently logged onto the system.
		DELETE/LIST/ALL BRDCST	⊙ ⊙	ALL is now the default. Specifies that the message is to be sent to all active consoles.
		CN	•	Specifies that the message is to be sent to a specific operator console. Messages sent via the OPER or CN keywords will have respectively OPER or CN appended to the message.
OUTPUT		DEST (stationid)	⊙	Specifies the remote station user to which SYSOUT data sets -- for jobs and classes specified in other parameters -- are to be directed.
		HOLD/ <u>NOHOLD</u>	⊙	Output data sets are held on a queue until the user wants to print them to his terminal (HOLD), or are to be printed immediately by the system writer.
		KEEP/ <u>NOKEEP</u>	⊙	Specifies whether to keep the SYSOUT data set after printing it.
		NEWCLASS DELETE	⊙	These parameters perform functions previously performed by variations of the NOPRINT parameter.
		NOPRINT	⊙	No longer supported. Functionally replaced by NEWCLASS and DELETE.
PROFILE		LIST	⊙	Prints out the user's profile.
		MODE/ <u>NOMODE</u>	⊙	Mode messages are printed under the EDIT command.
SEND		USER (*)	⊙	The message is sent to the stationid of the terminal issuing the request.
		SAVE	⊙	The message is placed in the broadcast data set, not to be printed until the next LOGON -- even if the recipient is currently logged onto the system.
		WAIT/ <u>NOWAIT</u>		WAIT -- The transmitting terminal is locked until the receiving terminal can accept the message. NOWAIT -- If the receiving terminal is busy, the message is not sent.
		CN	⊙	Specifies that the message is to be sent to a specific operator console. Messages sent via the OPER or CN keywords will have respectively OPER or CN appended to the message.
SUBMIT			⊙	The account number under which the user logs on is included on the JOB card generated by the SUBMIT command. The job is printed with the JOB SUBMITTED message.
TIME			⊙	The local time-of-day and the cumulative service units are printed.
<ul style="list-style-type: none"> • New or Changed from MVT. ⊙ New or Changed from VS2 Release 1. ⊙ New or Changed from MVT and VS2 Release 1. 				

Figure 21. Changes to Time Sharing Commands for VS2 Release 2 (Part 3 of 3)

Job Entry Subsystem Support

The FIB (foreground initiated background) commands SUBMIT, CANCEL, STATUS, and OUTPUT, interface with the job entry subsystem to submit or access a background job.

The restriction that was previously specified at SYSGEN limiting the number of jobs submitted by a user, is removed in VS2 Release 2.

ALLOCATE, FREE, OUTPUT, and SUBMIT allow SYSOUT data sets to be routed to a remote station.

Some of the operands of the DISPLAY subcommand are no longer supported, since displaying job and queue data are functions of the job entry subsystem.

Edit Support

The EDIT command is used to create, modify, display, and store data. FBS (fixed-block standard) record format is used for new data sets created by the SAVE command, when fixed records are being edited.

ALLOCATE, SEND, and PROFILE are now subcommands of EDIT.

Initial entry into EDIT places the current line reference at the top of the data set. Initial entry through the null line ENTRY continues to position this reference at the end of the data set.

Accounting Support

A utility (see the topic "Reformatting the UADS" in this chapter) is provided to convert the SYS1.UADS data set to the VS2 Release 2 format. LOGON supports both the old and new format. However, since all updating of accounting fields is done in the new format, accounting fields of an old format UADS must be reinitialized.

Time Sharing Integration Support

Under MVT and VS2 Release 1, time sharing (TSO) was an option which was started by the operator. A TSO task was created which could be modified with a subcommand of the OPERATOR command. This MODIFY subcommand is no longer supported, since there is no time sharing task.

Time sharing is started and stopped by the operator issuing a MODIFY TCAM TS=START or STOP. All the installation performance optimization (e.g., driver Q parameters), performance objectives (e.g., background percentage), and system control functions previously specified by START or MODIFY TS commands, are performed by the system resources management routines.

The SYS1.UADS and SYS1.BROADCAST data sets must be mounted at system initialization since they are allocated at that time.

The SPACE and DSNAME parameters of the MONITOR subcommand are not supported, since there is no time sharing task.

Data Set Conversion

Changes to data sets in VS2 Release 2 require the conversion of existing MVT or VS2 Release 1 data sets.

- The system catalog (SYSCTLG) must be converted to a key-sequenced VSAM data set.
- The broadcast data set and the user attribute data set must be converted to the new VS2 Release 2 format.

Catalog Conversion

The VS2 Release 2 system is capable of handling both OS and VSAM data sets, however Release 2 will support only VSAM catalogs. This will necessitate the creation of VSAM catalogs corresponding to existing OS catalogs, and creates a need for the capability to update both types of catalogs to reflect new or updated non-VSAM data sets for installations which run both OS and VS2 Release 2. Utilities are therefore provided to update the VS2 Release 2 catalog to reflect information in an OS catalog, and to update an OS catalog to reflect the changes in the VS2 Release 2 catalogs for non-VSAM data sets.

The VSAM catalog is updated on a VS2 Release 2 system; the OS catalog on the system that has the catalogs to be updated.

The structure of VSAM catalogs is different from that of OS catalogs. For VSAM data sets, the VSAM catalog functions somewhat like an OS VTOC in that the VSAM catalog entry contains the cylinder and track address of the data set, as well as volume and device type information. For non-VSAM data sets, the catalog contains the same information as its OS counterpart. As in OS, a non-VSAM data set can be found either through the catalog or through the VTOC of the volume containing the data set. Figures 22 and 23 show simplified relationships between OS and VSAM catalogs.

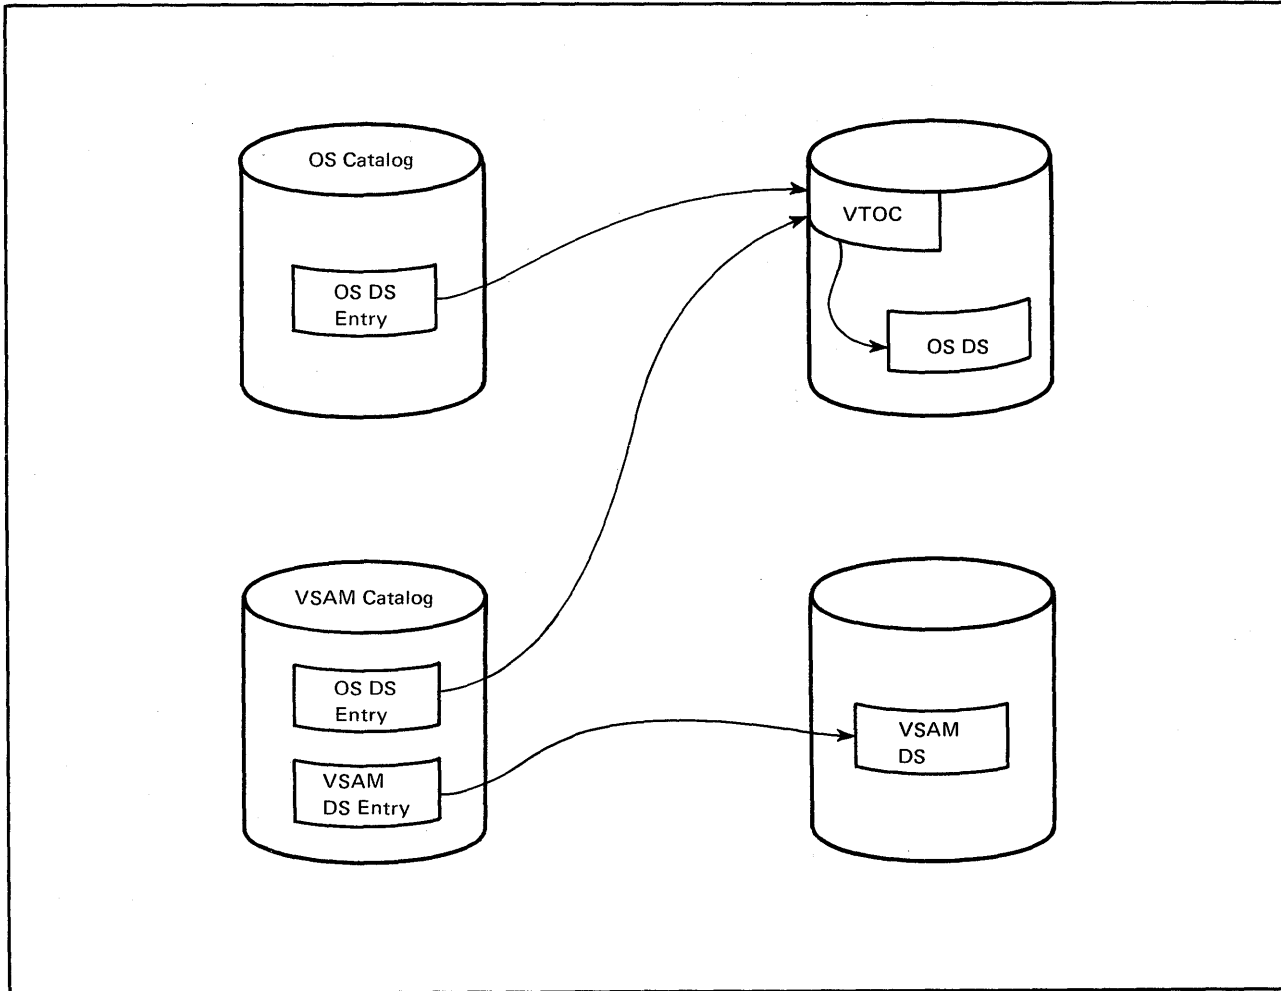


Figure 22. Simplified Catalog Structure - OS vs. VSAM

OS Catalog	VSAM Catalog
Contains only OS data set entries	Contains entries for either OS or VSAM data sets
One catalog during initialization	One catalog during initialization
System catalog must be on the IPL volume	System master catalog need not be on the IPL volume
CVOL	User or private catalog
System-controlled search strategy	User-controlled search strategy
All catalogs are named SYSCTLG	Catalogs may be named by the user
Cannot be password protected	Can be password protected

Figure 23. Some OS/VSAM catalog relationships

The order of search for LOCATE may be altered in VS2 Release 2 by using VSAM DD statements defining a specific catalog to be used as the primary catalog for a single job or job step. The order of search for the specified data set is then:

- First, all JOBCAT/STEPDAT catalogs in the order they are specified.
- Second, the catalog indicated by a catalog connector (analogous to CVOL pointer) in the data set name,

or

the master catalog, if there is no connector in the data set name.

VS2 Release 2, through the access method services, provides the means to create VSAM catalogs and to update the VSAM catalogs with entries from OS catalogs. It also provides through the IEHUCAT utility, the means for updating OS catalogs with the non-VSAM data set entry changes to the VS2 Release 2 catalogs. Some changes to the catalog functions of system utilities and to the CAMLST macro have resulted.

Catalog Utility and Macro Changes

The following changes have been incorporated into the IEHLIST, IEHPROGM, and IEHMOVE utilities, and the CAMLST macro:

- IEHLIST -- LISTCTLG has been eliminated. It has been replaced by the access method services function, LISTCAT.
- IEHMOVE -- All catalog-related functions have been deleted or replaced by access method services.

Deleted Functions

MOVE DSGROUP
COPY DSGROUP
MOVE CATALOG
COPY CATALOG

- IEHPROGM -- All catalog-related functions have been removed and replaced by access method services, except for CATLG and UNCATLG which are still available for non-VSAM data sets.

Function	Access Method Services Replacement
BLDX	none
DLTX	none
BLDA	DEFINE
DLTA	DELETE
BLDG	DEFINE
CONNECT	DEFINE or IMPORT
RELEASE	DELETE or EXPORT

- CAMLIST -- The CAMLST macro is supported as follows:

Operation	Supported
LOCATE by name	yes
LOCATE by TTR	no
CATLG	yes - non-VSAM only
UNCATLG	yes - non-VSAM only
RECATLG	yes - non-VSAM only
BLDX	no
DLTX	no
BLDA	no - use DEFINE
DLTA	no - use DELETE
BLDG	no - use DEFINE
LNKX	no - use IMPORT or DEFINE
DRPX	no - use EXPORT or DELETE

Creating or Updating the VSAM Catalog

Creating and/or updating VSAM catalogs is accomplished on the VS2 Release 2 system through the access method services (Figure 24).

Function	Access Method Services Verb
Define a VSAM data set or catalog, and catalog non-VSAM data sets.	DEFINE
Define and delete aliases for catalog names and non-VSAM data set names.	DEFINE, DELETE
List VSAM and non-VSAM catalog entries or records of a data set.	LISTCAT, PRINT
Alter entries on aliases in a specified VSAM catalog	ALTER
Print data set records.	PRINT
Move or copy a VSAM catalog.	REPRO
Convert to a VS2 Release 2 catalog.	CNVTCAT
Convert a sequential or an indexed sequential data set to the VSAM format.	REPRO
Convert a VSAM data set to sequential format.	REPRO
Make a data set portable from one operating system to another.	EXPORT, IMPORT
Copy a VSAM data set for reorganization.	REPRO
Create a back-up data set.	EXPORT
Support generation data groups (GDGs).	DEFINE, DELETE

Figure 24. Functions of the Access Method Services

The process of creating the VS2 Release 2 catalogs and updating them with entries from the OS catalogs might typically go like this:

- At SYSGEN, the DATASET macro is used to create the new VSAM master catalog.
- The new system's catalogs are created using DEFINE. For example, one master catalog might be defined to correspond to the OS system catalog, and one user or private catalog defined for each OS CVOL.

- CNVTCAT is run using an OS system catalog as input and the VS2 Release 2 master as the target catalog. All CVOL pointers are related to their corresponding VSAM user catalogs. The result is that the VSAM master catalog is a functional replacement for the input catalog.
- CNVTCAT is run for all OS CVOLS with the VS2 Release 2 catalog(s) as target.
- ALTER can be used to password protect the catalog(s).

Updating an OS Catalog From a VSAM Catalog

The IEHUCAT utility is used to update the OS catalogs to reflect all actions taken for the non-VSAM data sets in a VS2 Release 2 catalog. IEHUCAT is run on the system that has the catalogs to be updated.

VS2 Release 2 catalog management will optionally write to the SMF file, records reflecting additions, deletions, and alterations to the catalogs. IEHUCAT uses this information from the SMF file to update the OS catalogs.

Reformatting the UADS

In order to use any of the new keywords on the ACCOUNT command subcommands, the Release 2 format of the user attribute data set (UADS) must be used. The Account facility utility program named UADSREFM is used to create a new UADS or to reformat an existing UADS. This program reformats the UADS by reading a member from the existing UADS, copying its structure, incorporating the required new information, and writing the reformatted member to the new UADS. This reformatting process also eliminates any inefficient spacings that may have resulted from previous manipulations of a member's structure.

The UADSREFM utility program is invoked in an Account Facility input stream that includes a UADSREFM control card in the control data set.

If the VS2 Release 2 UADS format is not used, an error message will be issued when any of the new ACCOUNT subcommand keywords are used.

Formatting the Broadcast Data Set

The SYNC subcommand of the ACCOUNT command is used to create a broadcast data set in the Release 2 format. The SYNC subcommand also synchronizes the user identifications in the UADS with the user identifications in an existing broadcast data set. It should therefore be used immediately after the creation of a new UADS or the reformatting of an existing UADS.

The SYNC subcommand may be used from the foreground or included in the control data set of an Account Facility input stream.

If the VS2 Release 2 format of the broadcast data set is not used, an error message will be issued the first time an attempt is made to access it. Processing will continue, but the broadcast data set will be inaccessible.

Program Conversion

OS/VS2 Release 2 is generally upward compatible with MFT, MVT, VS1 and VS2 Release 1. Object programs and load modules that operate in MFT, MVT, or VS1 and follow IBM programming conventions described in *OS/VS Supervisor Services and Macros GC27-6979*, *OS/VS Data Management Services Guide GC26-3768*, and *OS/VS Data Management Macros GC26-3793*, will operate without change in this release. Some exceptions are:

- Programs sensitive to the PSW format (if migration is from an OS system).
- Programs that modify, are modifications of, or depend upon implementation of, the MFT, MVT, VS1, or VS2 Release 1 control program.
- Programs sensitive to the OS catalog structure.
- Programs that reference areas not normally available to a problem program through system functions.
- In a multiprocessing environment with shared real storage, multi-tasking programs which reference the same area should be reviewed for synchronization dependencies.
- Programs dependent upon MVT or VS2 Release 1 integrity exposures.
- EXCP, NOTE or POINT to a SYSIN or SYSOUT data set.
- User-written appendages which might now run enabled.

The capability to execute with virtual addresses equal to real (ADDRSPC=REAL) for programs that must operate entirely in real storage is provided for:

- programs that use EXCP or XDAP with user written appendages.

Note: Programs that use EXCP should observe the conventions outlined in the channel characteristics manual for the host system.

- programs that dynamically modify channel programs.
- programs that are time dependent (such as MICR).

TCAM Conversion

OS TCAM message control programs must be reassembled to run in the OS/VS2 environment. This reassembly allows the MCPs to benefit from the virtual storage capability of OS/VS2. Under OS/VS2, TCAM runs as a subsystem in a virtual address space. Certain TCAM elements, such as the buffer pool, I/O appendages, control blocks, and tables, are fixed in real storage for the duration of the TCAM task.

In most cases, TCAM application programs need only to be linkage edited.

Multiprocessing Considerations

OS/VS2 Release 2 supports two CPUs configured to share storage units -- a mode called tightly coupled multiprocessing. The two CPUs must be of the same model type; either two model 158 CPUs or two model 168 CPUs.

All available real storage is addressable by both CPUs. The only copy of the control program that exists in storage treats each CPU as a resource; nominally either CPU can be dispatched to perform any ready task at any time. Because two CPUs increase the demand for service from the single set of supervisor routines, a system of locks has been introduced to permit functions to be executed in parallel, one by each CPU. Where each CPU requests the same service at the same time, the locks ensure that the service is rendered in sequence.

In most cases the control program manages the multiprocessing configuration in such a way that it is logically indistinguishable from a uniprocessor. However certain unique aspects, in addition to those described in the **Introduction to OS/VS2 Release 2**, should be noted, including:

- Hardware configuration considerations.
- Operator control of configuration.
- Programmer considerations.

Hardware Configuration Considerations

Some hardware considerations for VS2 Release 2 are:

- Except for the presence of two CPUs, the minimum configuration is the same as for a uniprocessor.
- The real storage provided is typically double that provided for a single CPU of the same model.
- I/O devices need be accessible through only one channel path by one CPU. Unit record equipment is typically attached this way. For a device attached in this manner, no attachment of another control unit may be made to the channel position of the same number on the other CPU.
- Tape and direct access devices are normally accessible by both CPUs. In a typical case, control units are equipped with programmable 2-channel switches, and each path is attached at the same channel of each CPU. This type of I/O configuration is called symmetrical.
- 2914 manually controlled switching units can be used to improve the availability of asymmetrically attached I/O.

Operator Control of Configuration

Multiprocessing requires some additional education of operators. They must be aware of the following facets of the system:

- By means of the VARY command, one CPU can be dynamically added to or deleted from the configuration. The VARY command can also be used to associate a channel with a tape or direct access device.
- Addressable storage is configured from a collection of units. The real address range to be represented by a unit can be set by means of hardware dials.

Addressable real storage can be dynamically decremented or incremented, and need not be contiguously addressable.

- Two uniprocessors can be dynamically configured from a multiprocessor.
- A primary use of the VARY and QUIESCE commands is to logically delete elements (storage units, CPU, channels, control units, and devices) from the configuration for maintenance purposes.
- If a CPU encounters a check condition, the system automatically attempts to continue operation without using the offending element. The operator is informed of the action taken, and under certain circumstances controls the follow-on procedure.

Programmer Considerations

A problem program which executes successfully on a uniprocessor under OS/VS2 Release 2 will, in almost all cases, run successfully on a multiprocessing configuration under the same control program. Exceptions that do arise usually involve multitasking where, under multiprocessing, two tasks can be executed simultaneously -- one by each CPU. The locking structure is described in the **Introduction to OS/VS2 Release 2, GC28-0661**. Situations that must be considered are:

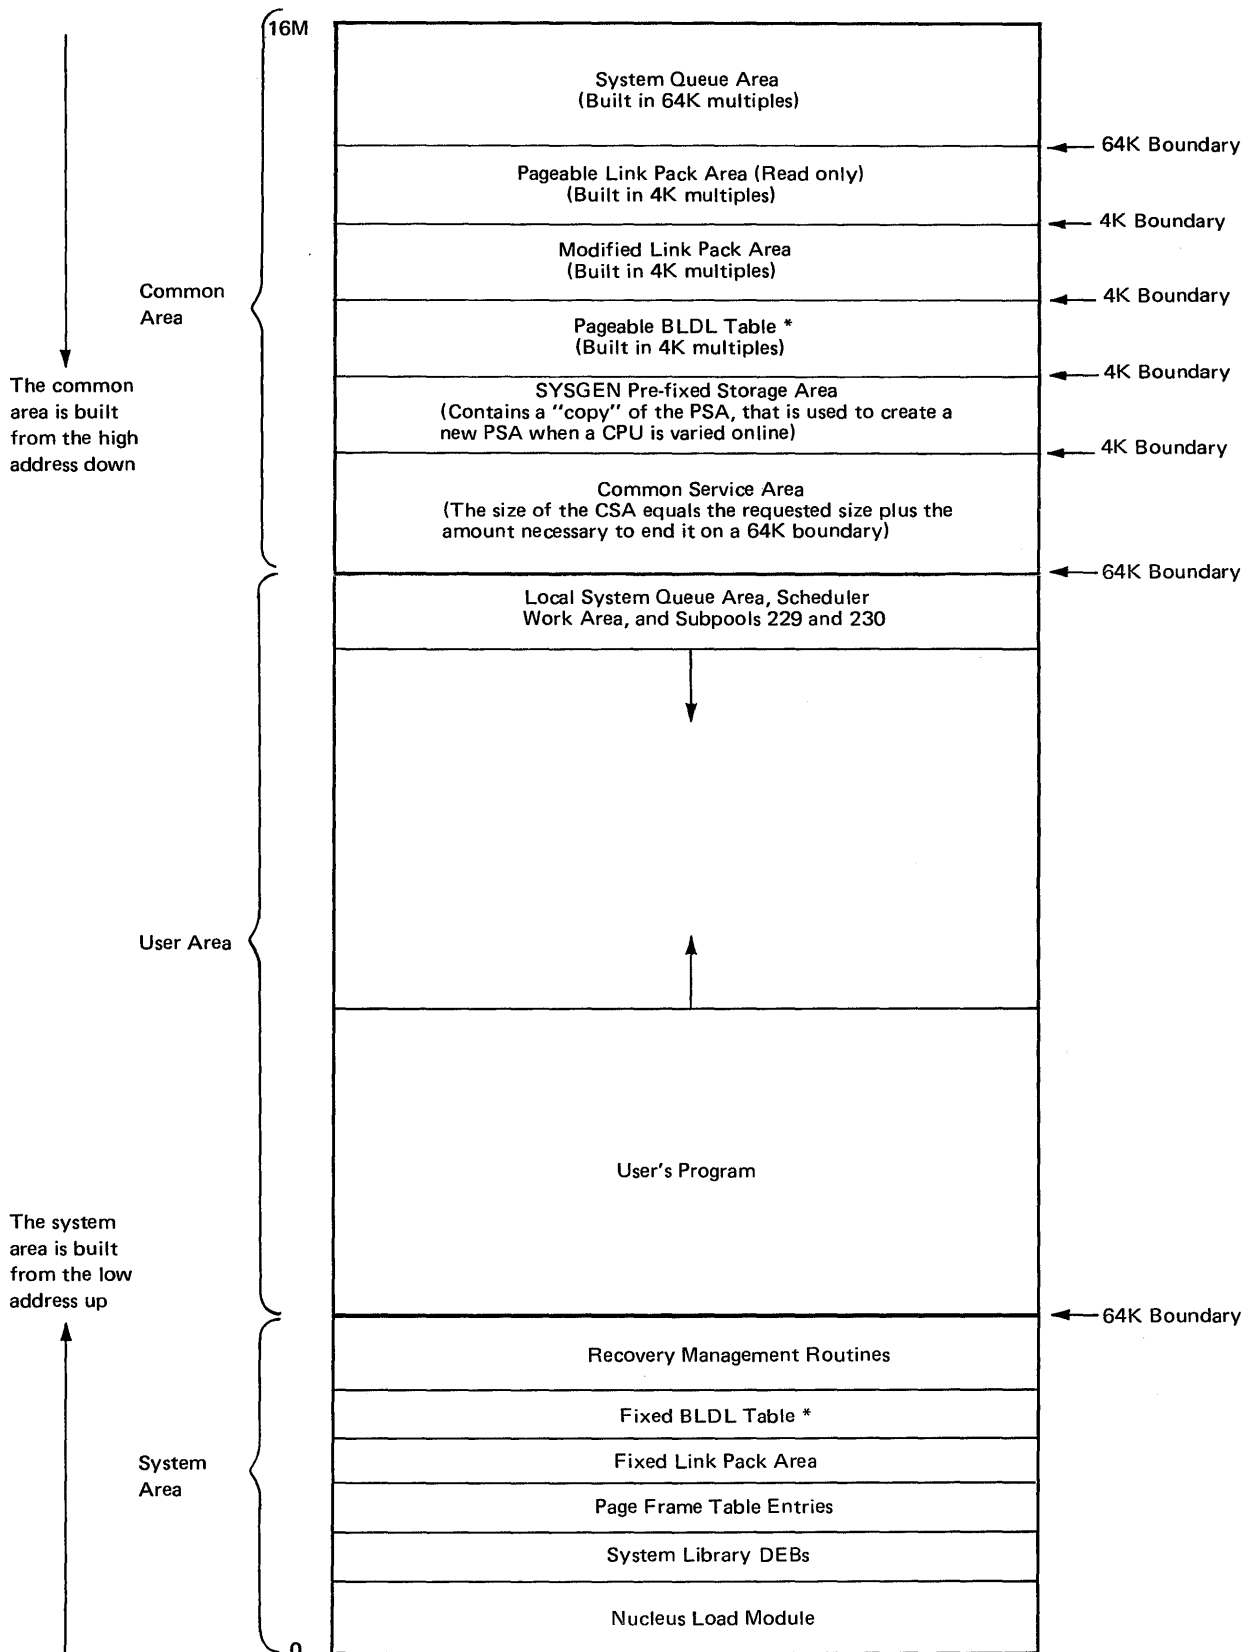
- Implied wait dependencies, where a low priority task that POSTs a higher priority task, will not be resumed on a uniprocessor until the higher priority task issues a WAIT. With multiprocessing, the lower priority task can be resumed immediately by the second CPU. Both tasks might then be executed concurrently.
- Similarly, a task that issues ATTACH may find that the attached task immediately commences execution on the second CPU.
- If two tasks share the same data, and if logical consistency is necessary, explicit methods of ensuring serial access (e.g., WAIT/POST or ENQ/DEQ) must be used since the priority differences between the tasks will no longer ensure sequential access. Furthermore, data which is fetched from or stored into such shared fields, even by a single instruction such as MVC, may not be internally consistent if these programmed sequencing controls are not used.
- Instruction streams (i.e., individual instructions) must not be dynamically modified by other instructions being executed by a CPU.
- Where only one CPU has a feature required for execution of a problem program, CPU affinity scheduling for that program can be specified at SYSGEN time. For example, an emulator might be installed on only one CPU in a Model 168 multiprocessing configuration. CPU affinity scheduling can be used to ensure that only the appropriate CPU is used to execute emulator-dependent job steps. CPU affinity scheduling is not required when both CPUs have installed the same emulator feature.

Appendix A: JES2 Differences from MVT/HASP and VS2 Release 1/HASP

MVT/HASP	VS2 Release 1/HASP	JES2
Tape input (via HASP START command)	Tape input (via HASP START command)	Internal reader (via a cataloged procedure)
Execution task monitor	Dynamic dispatcher	System resources management routines
Binary synchronous code adapter and synchronous transmitter receiver communications facilities	Binary synchronous code adapter and synchronous transmitter receiver communications facilities	Binary synchronous code adapter communications facilities only
Pseudo devices for SYSIN/SYSOUT via QSAM, BSAM, or EXCP	Pseudo devices for SYSIN/SYSOUT via QSAM, BSAM, or EXCP	Subsystem interfaces for SYSIN/SYSOUT via QSAM and BSAM only
Installation-written output writers only via the system output writer	Installation-written output writers only via the system output writer	Subsystem interface allows installation-written routines to obtain SYSOUT data sets from JES2
HASP console services <u>or</u> multiple console support	Multiple console support with HASP interfaces	Multiple console support with JES2 interfaces
Purge card	Purge record written to SMF data set	Purge record written to SMF data set
System output queued by four types	System output queued by forms, printer set-up, and 36 output classes	System output queued by forms, printer set-up, and 36 output classes
No equivalent function	Forms, printer set-up, and routing information may be specified on a HASP OUTPUT control card	Forms, printer set-up, and routing information may be specified on a JES2 OUTPUT control card
Restricted checkpoint/restart support for jobs using SYSIN/SYSOUT	Restricted checkpoint/restart support for jobs using SYSIN/SYSOUT	Full checkpoint/restart support
No SMF SYSIN record counts No SMF output limiting support No SMF output record	No SMF SYSIN record counts No SMF output limiting support SMF output record is written to the SMF data set	Full SMF support
Numeric forms designation	Alphanumeric forms designation	Alphanumeric forms designation
TSO SYSOUT data sets are written by the system output writer	TSO SYSOUT data sets are written by the system output writer	Time sharing SYSOUT data sets are processed by JES2 and are retrievable via the OUTPUT command
Background jobs scheduled from a terminal must use system facilities rather than HASP	TSO users may submit jobs to HASP for scheduling, inquire of their status, and cancel them	Time sharing users may submit jobs to JES2 for scheduling, inquire of their status, and cancel them

Figure 25. JES2 Differences From MVT/HASP and VS2 Release 1/HASP

Appendix B: Virtual Storage Map



* The pageable and fixed BLDL tables are mutually exclusive.

Figure 26. Virtual Storage Map

Glossary

This glossary provides definitions of OS/VS2 terms. For definitions of terms not included, see IBM Data Processing Glossary, GC20-1699.

ABP: Actual block processor.

ACR: Alternate CPU recovery.

actual block processor: One of several programs that translates I/O requests into the proper format for the I/O supervisor.

address space: The virtual storage assigned to a batch or terminal job, a system task, or a task initiated by the START command. Each address space consists of the same range of addresses.

address space identifier: A unique, system-assigned identifier for an address space.

address translation: The process of changing the address of an item of data or an instruction from its virtual address to its real storage address. See also dynamic address translation.

alternate CPU recovery: a facility that attempts system recovery when a CPU fails by transferring work to another CPU.

APAR: Authorized program analysis report. A request for correction of a problem caused by a defect in a current unaltered release of a program. A PTF or corrected code is issued to the customer and the correction is incorporated into subsequent releases of the program.

APF: Authorized program facility.

APG: Automatic priority group.

ASID: Address space identifier.

ASM: Auxiliary storage manager.

ASP: See JES3.

asymmetric I/O: In a tightly coupled multiprocessing environment, an I/O device configuration in which any I/O device is physically attached to only one CPU.

authorized library: A library that can contain authorized programs.

authorized program: A system program or user program that is allowed to use restricted functions.

authorized program facility: A facility that permits the identification of programs that are authorized to use restricted functions.

automatic priority group: A group of jobs, each having its own address space at a single installation-specified priority level, that are dispatched according to a system resources manager algorithm that attempts to provide optimum use of CPU and I/O resources.

auxiliary storage: Data storage other than main storage.

auxiliary storage management: A facility that controls the allocation and release of pages on external page storage, and schedules I/O operations on the page data set.

basic control (BC) mode: A mode in which the features of a System/360 computing system and additional System/370 features, such as new machine instructions, are operational on a System/370 computing system. See also extended control (EC) mode.

BC mode: Basic control mode.

block processor: See actual block processor.

change bit: A bit associated with a page in real storage. The change bit is turned "on" by hardware whenever the associated page in real storage is modified.

channel program translation: In a channel program, replacement by software of virtual addresses with real addresses.

common area: The area of virtual storage that is addressable by all address spaces.

common service area: A part of the common area that contains data areas addressable by all address spaces, but protected during its use by the key of the requester. Abbreviated CSA.

communication task: A function that handles all communication between programs and operator consoles.

control program keys: Protection keys (0-7) that are reserved for control program use.

control registers: A set of registers used for operating system control of relocation, priority interruption, program event recording, error recovery, and masking operations.

converter: The routine that transforms spooled JCL statements into an internal format, and processes commands entered through the input stream.

CPU affinity: In a multiprocessing configuration, a function that allows a unit of work to be directed to a particular CPU for execution.

cross memory services lock: A global suspend lock that is used for services that apply to more than one private address space.

CSA: Common service area.

DAT: Dynamic address translation.

demand paging: Transfer of a page from external page storage to real storage at the time it is needed for execution.

disabled page fault: A page fault that occurs when interruptions are disallowed by the CPU.

DSS: Dynamic support system.

dynamic address translation: (1) The change of a virtual storage address to a real storage address during the execution of an instruction. (2) A hardware feature that performs the translation. Abbreviated DAT.

dynamic area: The portion of virtual storage that is divided into address spaces that are assigned to job steps and system tasks.

dynamic support system (DSS): An interactive debugging facility that allows authorized maintenance personnel to monitor and analyze events and alter data.

EC mode: Extended control mode.

enabled page fault: A page fault that occurs when interruptions are allowed by the CPU.

ESTAE/ESTAI exit routines: Task recovery exit routines from which percolation to a higher-level exit may occur.

extended control (EC) mode: A mode in which all the features of a System/370 computing system, including dynamic address translation, are operational. See also basic control (BC) mode.

external page storage: The portion of auxiliary storage that is used to contain pages.

external page table (XPT): An extension of a page table that identifies the location on external page storage of each page in that page table.

external slot: See slot.

fixed: Not capable of being paged out. A fixed page in a private address area can be swapped out.

fixed BLDL table: A BLDL table that the user has specified to be fixed in the lower portion of real storage.

fixed link pack area: An extension of the link pack area that occupies fixed pages in the lower portion of real storage.

fixed page: a page in real storage that is not to be paged out.

fragmentation: The inability to assign real storage locations to virtual addresses because the available spaces are smaller than the page size.

frame: Same as page frame.

FRR: Functional recovery routine.

functional recovery routine: A recovery routine that is used by locked programs, SRBs, and supervisor control routines.

global: Pertaining to more than one private address space.

global lock: A lock that protects serially reusable resources related to more than one private address space.

global services: Services that apply to more than one private address space. Also used when referring to multiple JES3 processors.

graphics access method (GAM): A facility that supports the 2250 Display Unit and the 2260 Display Station through the use of graphic programming services (GPS) and the graphic subroutine package (GSP).

hardware error recovery management system: A facility that attempts recovery from hardware malfunctions. It consists of the machine check handler (MCH) and the channel check handler (CCH).

HASP: See JES2.

installation performance specification (IPS): A set of installation-supplied control information used by the workload management routines. An IPS includes performance group definitions, performance objectives, workload level numbers and coefficients used to define the service rate.

internal writer: A facility in the job entry subsystem (JES2 or JES3) that allows user written output writers to write data on devices not directly supported by the job entry subsystem.

internal reader: A facility that transfers jobs to the job entry subsystem (JES2 or JES3).

interaction: The acceptance by a system of a line of input from a terminal, processing of the line, and a return of data, if any, to the terminal.

interprocessor communication: In a multiprocessing system, a function that provides communications between CPUs that are under control of the same control program.

interval service value (ISV): In the system resources manager, a category of information contained in a period definition, which specifies the minimum amount of service that an associated job will receive during any interval that the address space is in real storage.

invalid page: A page that cannot be directly addressed by the dynamic address translation feature of the central processing unit.

IPC: Interprocessor communication.

IPS: Installation performance specification.

ISV: Interval service value.

JES2: A functional extension of the HASP II program that receives jobs into the system and processes all output data produced by the job.

JES3: A functional extension of ASP Version 3, that performs job entry, scheduling, and output processing functions. It also provides loosely coupled multiprocessing support for up to four System/370 processors.

LCH: Logical channel queue.

LGN: Logical group number.

link pack area (LPA): An area of virtual storage containing selected reenterable and serially reusable routines that are loaded at IPL time and can be used concurrently by all tasks in the system.

link pack update area: An area in virtual storage containing modules that are additions to or replacements for link pack area modules.

local lock: A suspend lock that protects the resources assigned to a particular private address space.

local service: A supervisory service that applies to only one private address space. Also used when referring to a single JES3 processor.

local system queue area (LSQA): One or more segments associated with each address space that contain system control blocks.

lock: A means of serialization used by supervisory control program routines. See global lock, local lock, spin lock, suspend lock.

logical channel queue: A queue of operations to be performed on a channel.

logical group: A collection of pages that are related to each other. An address space may be composed of multiple groups: one for the LPA, one for the SWA, and one for the private address space.

logical group number: An identifier of a logical group. It is the base value used by the real storage manager and the auxiliary storage manager to compute the identifier of a logical page. All the pages of a VIO data set may be represented by a single LGN.

LPA: Link pack area.

LSQA: Local system queue area.

master address space: The virtual storage used by the master scheduler task.

MCP: Message control program.

Message control program: A program that is used to control the sending or receiving of messages to or from remote terminals.

MF/1: System activity measurement facility.

MIC: Missing interruption checker.

missing interruption checker: A program that periodically tests active I/O operations in order to detect missing I/O interruptions and inform the operator of abnormal conditions before they can affect system operation.

MP: Multiprocessing

multiple address space: A feature that provides each user with a private address space.

multiprocessing (MP) system: A computing system employing two or more interconnected processing units to execute programs simultaneously.

node: An addressable point in a teleprocessing network.

page: (1) A fixed-length block of instructions, data, or both, that can be transferred between real storage and external page storage. (2) To transfer instructions, data, or both, between real storage and external page storage.

page data set: A data set in external page storage, in which pages are stored.

page fault: A program interruption that occurs when a page that is marked not in real storage is referred to by an active page. Synonymous with page translation exception.

page fixing: Marking a page nonpageable so that it remains in real storage.

page frame: A block of real storage that can contain a page. Synonymous with frame.

page-in: The process of transferring a page from external page storage to real storage.

page-out: The process of transferring a page from real storage to external page storage.

page reclamation: The process of making addressable the contents of a page in real storage that has been marked invalid. Page reclamation can occur after a page fault.

page stealing: To take away an assigned page frame from a user to make it available for another purpose.

page table (PGT): A table that indicates whether a page is in real storage and correlates virtual addresses with real storage addresses.

page translation exception: A program interruption that occurs when a virtual address cannot be translated by the hardware because the invalid bit in the page table entry for that address is set. Synonymous with page fault. See also segment translation exception, translation specification exception.

paging: The process of transferring pages between real storage and external page storage.

paging device: A direct access storage device on which pages (and possibly other data) are stored.

paging rate: The average number of page-ins and page-outs per unit of time.

parallel: Pertaining to the simultaneity of two or more processes.

PER: Program event recording.

percolation: In error recovery, the passing of control to a higher-level recovery routine from another recovery routine along a preestablished path.

performance group: An installation-defined category that includes users with similar requirements and associates them with a performance specification.

performance objective: A category of information contained in an installation performance specification. Each performance objective specifies the service rate(s) that an associated job is to receive, for a number of different workload levels.

period definition: In the system resources manager, a category of information contained in a performance group definition, that indicates which performance objective is to be followed, either during a particular real-time period or until a particular amount of service has been accumulated by an associated job.

PGT: Page table.

private address space: An address space assigned to a particular user.

program event recording (PER): A hardware feature used to assist in debugging programs by detecting and recording program events.

program management: Routines that perform supervisory services related to the moving of programs from the system library or user libraries to virtual storage.

quick cell facility: A high-performance storage allocation technique using a fixed block size.

RAS: Reliability, availability, serviceability.

real address: The address of a location in real storage.

real storage: The storage of a System/370 computing

system from which the central processing unit can directly obtain instructions and data, and to which it can directly return results.

real storage management: Routines that control the allocation of pages in real storage. Abbreviated RSM.

reconfiguration: The process of placing a CPU, main storage, or channel offline for maintenance; and adding or removing components while the system is in operation.

recovery routine: A routine that is entered when an error occurs during the performance of an associated operation. It isolates the error, assesses the extent of the error, indicates subsequent action, and attempts to correct the error and resume operation.

recovery termination manager: A program that handles all normal and abnormal termination of tasks by passing control to a recovery routine associated with the terminated function.

reference bit: A bit associated with a page in real storage; the bit is turned "on" by hardware whenever the associated page in real storage is referred to. There is a reference bit in each of two storage keys associated with each page frame.

refreshable: A refreshable load module can be replaced by a new copy during its execution without changing either the sequence or results of processing.

relocate hardware: See dynamic address translation.

relocation interrupt: See page fault.

remote terminal access method (RTAM): A facility that controls operations between the job entry subsystem (JES2 or JES3) and remote terminals.

response/throughput bias (RTB): In the system resources manager, a category of information contained in a period definition, which indicates how the workload manager is to weigh trade-off between satisfying some system throughput objective, and satisfying the IPS-specified service rate.

RMTGEN: Generation of remote work stations for remote job entry.

RSM: Real storage manager.

RTAM: Remote terminal access method.

RTB: Response/throughput bias.

scheduler work area (SWA): One or more segments in virtual storage that contain most of the job management control blocks, such as the JCT, JFCB, SCT, and SIOT. There is one SWA for each initiator.

secondary storage: Auxiliary storage.

segment: A contiguous 64K area of virtual storage.

segment table (SGT): A table used in dynamic address translation to control user access to virtual storage segments. Each entry indicates the length, location, and availability, of a corresponding page table.

segment translation exception: A program interruption that occurs when a virtual address cannot be translated by the hardware because the invalid bit in the segment table

entry for that address is set. Synonymous with program interruption code 16.

service: See service rate.

service rate: In the system resources manager, a measure of the rate that system resources (service) are provided to individual jobs. It is used by the installation to specify performance objectives, and used by the workload management routines to track the progress of individual jobs.

service unit: Same as service.

slot: A continuous area on a paging device in which a page can be stored.

software recording facility: A facility used by functional recovery routines (FRRs) to write system error records.

spin lock: A lock that prevents a CPU from doing work until the lock is cleared. Contrast with suspend lock.

SQA: System queue area.

SRB: Service request block.

SRM: System resources manager.

suspend lock: A lock that prevents requesters from doing work on a CPU, but allows the CPU to continue doing other work. Contrast with spin lock.

SWA: Scheduler work area.

swapping: A paging technique that writes the active pages of a job to external page storage, and reads pages of another job from external page storage into real storage.

symmetric processors: Processors with identical configurations.

symmetric storage configurations: Machine configurations with identical storage units.

system activity measurement facility (MF/1): A facility that collects information about system resource usage such as paging activity and the use of the CPU, channels, and I/O devices, to produce trace records and reports.

system key: A key that protects system data from damage or modification by unauthorized users.

system queue area (SQA): An area of virtual storage reserved for system-related control blocks. It contains fixed pages and is assigned protection key zero. Abbreviated SQA.

system resources manager: A group of programs that controls the use of system resources in order to satisfy the installation's performance objectives.

thrashing: A condition in which the system can do little useful work because of excessive paging.

timer supervision: Routines that provide the date and time of day, measure time intervals, schedule activities, and set the interval timer.

TPIOS: A facility that supports programmed telecommunications control units (TCUs) and generates channel programs for the channel scheduler.

transaction: In the system resources manager, a performance measurement made on an address space basis. It is a job or a job step for either batch or remote batch entry. It is a terminal interaction for time sharing (TSO) jobs.

translation tables: Page tables and segment tables.

uniprocessing: Sequential execution of instruction by a central processing unit or independent use of a CPU in a multiprocessing system.

VIO: Virtual I/O.

virtual address: An address that refers to virtual storage and must, therefore, be translated into a real storage address when used.

virtual equals real address area (V=R): An area of virtual storage that has the same range of addresses as real storage, and is used for a program or part of a program that is not to be paged or swapped during execution.

virtual I/O (VIO): A facility that pages data into and out of external page storage, although to the problem program, the data appears to be read from or written to direct access storage devices.

virtual memory: See virtual storage.

virtual storage: Addressable space that appears to the user as real storage, from which instructions and data are mapped into real storage locations. The size of virtual storage is limited by the addressing scheme of the computing system and by the amount of auxiliary storage available, rather than by the actual number of real storage locations.

virtual storage access method (VSAM): An access method for direct or sequential processing of fixed and variable length records on direct access devices.

virtual storage management (VSM): Routines that allocate address spaces and virtual storage areas within address spaces, and keep a record of free and allocated storage within each address space.

VSAM: Virtual storage access method.

VSM: Virtual storage management.

V=R area: See virtual equals real address area.

workload manager: A part of the system resources manager that allows an installation to determine the performance that any group of users will receive, monitors the workload, and schedules resources accordingly.

XPT: External page table.

Indexes to OS/VS publications are consolidated in the OS/VS Master Index, GC28-0602, and the OS/VS Master Index of Logic, GY28-0603. For additional information about any subject listed below, refer to other publications listed for the same subject in the Master Index.

/*JOBPARM control card 97
/*OUTPUT control card 97-98

A operand

DISPLAY command 100
MONITOR command 101
TRACK command 102
ABP (actual block processor) 125
access method services 115
 highlights 12
 initializing system data sets 26
 replacement for IEHLIST, IEHMOVE, IEHPROGM functions 113
 used to create or update a VSAM catalog 115-116
 verbs 115
access methods
 BTAM 21
 GAM 21,126
 ISAM 21
 QTAM 13
 RTAM 126
 TCAM 21,13
 VSAM 10,129
 VTAM 10,21
ACCOUNT command 107
ACR (alternate CPU recovery)
 definition 125
actual block processor (ABP)
 definition 125
ADD subcommand 107
address range 11
address space
 definition 125
 master scheduler 35,36
 multiple virtual 11
 swapping by the system resources manager 50
 time sharing user 35
address space identifier (ASID)
 definition 125
address translation
 definition 125
ADDRSPC job control language parameter 96,93
ADDRSPC=REAL (see virtual equals real address area)
AFF job control language parameter 93
affinity, CPU 119
 definition 125
AFFINITY system generation macro instruction 20,25
ALGLIB system generation macro instruction 25
ALGOL system generation macro instruction 25
ALL operand 100
ALLOCATE command 107
ALLOCATE subcommand 108
 EDIT command 110
 job entry subsystem support 110
ALTER verb 115
alternate CPU recovery (ACR)
 definition 125
ALTSYS system initialization parameter 36
AMBLIST service aid 13
APAR
 definition 125
 for integrity exposures 75
 for type I programs 15
APF (authorized program facility) 81
 authorized program library list 40
 definition 125
 SYSGEN specification with SVCTABLE macro 21
APF system initialization parameter 30,29

APG (automatic priority group)
 definition 125
 job step dispatching priority default value 96
 system initialization parameter specification 30
 system resources manager algorithm 51
APG system initialization parameter 30,29
 difference from VS2 Release 1 34
appendage, installation-written 117,39
ASB (automatic SYSIN batching) reader 13
ASID (address space identifier)
 definition 125
ASM (auxiliary storage manager)
 definition 125
ASP 10
 definition 125
ASSEMBLER system generation macro instruction 25
asymmetric I/O 118
 definition 125
ATTRIBUTE command 107
authorized libraries 81
 definition 125
 duplicate module name restriction 83
 routines used only for their intended purpose 79
 specification in SYS1.PARMLIB list IEAAPFxx 40
authorized appendage list (IEAAP00) 39-40
authorized program facility (APF) 81
 authorized program library list 40
 definition 125
 SYSGEN specification with SVCTABLE macro 21
authorized programs 75,125
 definition 125
 requests to load other authorized programs 79
authorized program library list (IEAAPFxx) 40
authorized state 75
authorized system programmer 77
AUTO operand 101
automatic priority group (APG)
 definition 125
 job step dispatching priority default value 96
 system initialization parameter specification 30
 system resources manager algorithm 51
automatic SYSIN batching (ASB) reader 13
auxiliary storage
 definition 125
 paging activity report of use 65-66
auxiliary storage management
 definition 125
AUXLIST system initialization parameter 35

BACKSPACE addressed to SYSIN/SYSOUT DCB 13,14
basic control (BC) mode
 definition 125
basic telecommunications access method (BTAM) 21
 online terminal test option 96
 system generation specification 21
BLDL system initialization parameter 30,29
BLDLF system initialization parameter 30,29
BRDCAST operand 101
BRDR operand 100
BRDRQ operand 100
broadcast data set (SYS1.BROADCAST) 46
 formatting 116
 mounted for system initialization 110
 new format description 47
BTAM (basic telecommunications access method)
 online terminal test option 96
 system generation specification 21

CAMLST macro instruction 113-114
CANCEL command
 ALL 100
 BRDR 103,100
 IN 100
 OUT 100

USERID 100
 U 100
 CANCEL/STATUS command 108
 job entry subsystem support 110
 CANCEL subcommand 108
 catalog (see VSAM catalog)
 catalog connector 113
 catalog utility and macro changes 113-114
 CENPROCS system generation macro instruction 27,25
 CHAN MF/1 parameter 64
 CHANGE subcommand 107
 channel activity measurement 63,67
 channel activity report 67
 tuning example 68
 channel-CPU overlap activity measurement 63
 CHANNEL operand 102
 channel program translation 125
 CHANNEL system generation macro instruction 20,25
 channel-to-channel adapter 10
 CHECKER system generation macro instruction 25
 checkpoint/restart
 checkpoint data sets 82
 JES2 support 121
 CHNGDUMP command
 DEL 100
 SET 100
 CHRDR MF/1 parameter 64
 CKPTREST system generation macro instruction 21,25
 CLASS operand
 MODIFY command 100
 RESET command 101
 CLPA system initialization parameter 30,29
 CMD system initialization parameter 30,29
 CN operand 101
 CNVTCAT verb 115
 COBOL F
 compiler 15
 library 15
 COBLIB system generation macro instruction 25
 COBOL system generation macro instruction 25
 COMM MF/1 parameter 64
 command library (SYS1.COMDLIB) 46
 adding installation-written members 27
 commands
 operator 99-105
 time sharing (TSO) 106-110
 common area 123
 definition 125
 common service area (CSA) 123
 definition 125
 system initialization parameter specification 30
 communication task
 definition 125
 initialization 28
 concatenation
 JOBLIB and STEPLIB integrity considerations 83
 libraries containing APF-authorized programs 40
 with SYS1.LINKLIB 41
 configuration
 minimum for multiprocessing 118
 minimum for starter system 19
 CONSOLE system generation macro instruction 21,25
 control algorithm 51
 CONTROL command
 M 105,100
 T 105,100
 control function 51
 control program extensions 81
 conversational remote job entry (CRJE) 13
 conversion considerations 85-119
 converter
 definition 125
 COPIES job control language parameter 95,93
 counterfeit control block 77
 CPQE system initialization parameter 35
 CPU activity measurement 63,65
 CPU activity report 65
 CPU affinity 119
 definition 125
 CPU execution 52
 CPU load balancing algorithm 50
 weighted by resource factor coefficients 62,51
 CPU MF/1 parameter 64
 CPU operand 102
 CPU resource factor coefficient 62
 CPU service definition coefficient 62,52-53
 CRJE (conversational remote job entry) 13
 cross memory services lock
 definition 125
 CSA (common service area) 123
 definition 125
 system initialization parameter specification 30
 CSA system initialization parameter 30,29
 CTRLPROG system generation macro instruction 21,25
 CVIO system initialization parameter 31,29
 CVOL 112
 unsupported MVT function 13
 CYCLE MF/1 parameter 65

 DASD MF/1 parameter 64
 DAT (dynamic address translation) 125
 data reduction routines 86,87
 data sets
 catalog conversion 111-116
 concatenation 40,41
 freeing 94
 system 45-48
 SYS1.BROADCAST conversion 116
 SYS1.UADS conversion 116
 temporary 10
 DATAMGT system generation macro instruction 21,25
 DATASET system generation macro instruction 21,25
 for adding installation-written members to system
 libraries 26
 for initializing new system data sets 26
 DCB job control language parameter 93
 DCMLIB system generation macro instruction 25
 DD statement
 difference from MVT 93
 difference from VS2 Release 1 93
 maximum specification 95
 dedicated work files 13
 DEFINE verb 115
 defining the system 17-48
 DEL operand 100
 DELETE subcommand 108
 DELETE verb 115
 dependencies, compatibility 117
 DEQ macro instruction
 ensuring serial access for multiprocessing 119
 major name restrictions for authorized programs 79
 DEST job control language parameter 95,93
 device allocation algorithm 51
 device groupings 74
 device list functions of JES2 commands 104
 DEVICE MF/1 parameter 64
 direct system output (DSO) 13
 disabled page fault
 definition 125
 DISPLAY command
 A 100
 INIT 100
 jobname 105,100
 JOBS 100
 MATRIX 102,100
 N 105,100
 NET 102,100
 Q 105,100
 R 100
 SQA 105,100
 TS 100
 USER 105,100
 DISPLAY subcommand 108
 job entry subsystem support 110
 DPRTY job control language parameter 96,93
 DRIVER operand 101
 driver, TSO 14
 DSO (direct system output) 13
 DSS (dynamic support system)
 definition 126

NUCMAP system initialization parameter 32
 system data set SYS1.DSSVM 45,47,48
 dump data set (SYS1.DUMP) 46
 device specification by a system initialization
 parameter 31
 maximum number specified 34
 dump default lists
 IEAABDxx 40
 IEADMPxx 40
 overriding defaults with the CHNGDUMP
 command 105,100
 DUMP system initialization parameter 31,29
 difference from VS2 Release 1 34
 DYNAM job control language parameter 93,94
 dynamic address translation (DAT)
 definition 125
 dynamic allocation validation exit 83
 dynamic dispatcher 121
 dynamic support system (DSS)
 definition 126
 NUCMAP system initialization parameter 32
 system data set SYS1.DSSVM 45,47,48
 DYNAMNBR job control language parameter 94,93

EC (extended control) mode
 definition 126
 EDIT command 108
 EDIT system generation macro instruction 21,25
 EDITOR system generation macro instruction 27,25
 emulator 119
 specified by AFFINITY macro instruction 20
 EMULATOR system generation macro instruction 25
 enabled page fault
 definition 126
 END subcommand 108
 ENQ macro instruction
 ensuring serial access for multiprocessing 119
 major name restrictions for authorized programs 79
 ENQ/DEQ algorithm 51
 enqueue residence value (ERV) 62,51
 EOD operand 100
 EROPT JCL subparameter of DCB parameter 96,93
 error recording data set (SYS1.LOGREC) 45
 change from MVT and VS2 Release 1 47
 ERV (enqueue residence value) 62,51
 EXCP addressed to SYSIN/SYSOUT DCB 13,14
 EXEC command 108
 EXEC statement 93
 execution task monitor 121
 EXPORT verb 115
 extended control (EC) mode 126
 external page storage
 definition 126
 MF/1 counts 66,65
 external page table (XPT)
 definition 126

fetch protection 77,83
 FIB (foreground initiated background) commands 110
 file protection attributes 82
 FIX system initialization parameter 31,29
 fixed
 definition 126
 fixed BLDL table 123
 definition 126
 system initialization parameter specification 30
 fixed link pack area (LPA) 123
 definition 126
 specifying routines with the FIX parameter 31
 fixed LPA list (IEAFIXxx) 40
 fixed page
 definition 126
 foreground initiated background (FIB) commands 110
 FORTLIB system generation macro instruction 25
 FORTRAN G
 compiler 15
 library 15

FORTRAN H
 compiler 15
 library 15
 FORTRAN system generation macro instruction 25
 fragmentation
 definition 126
 virtual storage 11
 FREE command 108
 job entry subsystem support 110
 FREE job control language parameter 94,93
 freeing data sets 94
 FRID job control language subparameter of DCB
 parameter 95,93
 Full ANS COBOL V2
 compiler 15
 library 15
 functions not supported
 MVT 13-14
 VS2 Release 1 14

GAM (graphics access method)
 definition 126
 system generation specification 21
 GDG (generation data group) 115
 generalized trace facility (GTF)
 replacement for trace and TSO trace 14
 TRACE system initialization parameter 35
 GENERATE system generation macro instruction 21,25
 generation data group (GDG) 115
 GENTSO system generation macro instruction 25
 GETMAIN, restriction for variable-length 95
 GJOBCTL system generation macro instruction 25
 GJP (graphic job processor) 13
 global
 definition 126
 global lock
 definition 126
 global services
 definition 126
 glossary 125-129
 GRAPH MF/1 parameter 64
 graphic job processor (GJP) 13
 graphics access method (GAM)
 definition 126
 system generation specification 21
 GRAPHICS system generation macro instruction 27,25
 GTF (generalized trace facility)
 replacement for trace and TSO trace 14
 TRACE system initialization parameter 35

HALT command
 EOD 100
 NET 102,100
 TP 100
 hard copy log 31
 HARDCPY system initialization parameter 31,29
 hardware error recovery management system
 definition 126
 HASP
 definition 126
 summary of differences from JES2 121
 HELP system generation macro instruction 25
 HIARCHY job control language subparameter of DCB
 parameter 93
 hierarchy support
 unsupported MVT function 13,36
 highlights of VS2 Release 2 10-11
 HOLD command
 jobname 100
 Q 100
 TP 100
 HOLD operand 101
 HRAM system initialization parameter 36
 HSVC system initialization parameter 36

IBM-created SYS1.PARMLIB lists 39-42
 IEAABDxx (dump default list)
 description 40
 overriding defaults with the CHNGDUMP command 105,100
 IEAAPFxx (authorized program library list) 40
 IEAAPF00 (authorized appendage list) 39
 IEABLDxx (resident BLDL list) 41-42
 IEADMPxx (dump default list)
 description 40
 overriding defaults with the CHNGDUMP command 105,100
 IEAFIXxx (fixed LPA list) 40
 IEAIPSxx (see installation performance specification)
 IEALOD00 (LPA directory load list) 41
 IEALPAXx (modified link pack area list) 43
 IEAOPTxx (system resources management tuning parameter list) 62,43-44
 IEAPAK00 (LPA packing list) 41
 IEASYsxx (system parameter list)
 description 42
 recommendations for values 37
 IEBUPDAT utility program 13
 IEBUPDTE utility program 13
 IEFACTRT SMF exit 91
 IEFUIV SMF exit 91
 IEFUJI SMF exit 91
 IEFUJP SMF exit 91
 IEFUJV SMF exit 91
 IEFUSI SMF exit 91
 IEFUSO SMF exit 91
 IEFUTL SMF exit 91
 IEFU83 SMF exit 91
 IEHIOSUP utility program 13
 IEHLIST utility program 113,13,14
 IEHMOVE utility program 113,13,14
 IEHPROGM utility program 113,26,13,14
 IEHUCAT utility program 116,113
 IMAGELIB system generation macro instruction 25
 IMBDMAP program 13
 IMCOSJQD service aid 14
 IMPORT verb 115
 IN operand 100
 indexed sequential access method (ISAM) 21
 INIT operand
 DISPLAY command 100
 TRACK command 102
 initial program loader (IPL) 28
 initialization, master scheduler 28
 initialization, new system data sets 26
 initialization, system 28-44
 parameters 28-37
 installation-created SYS1.PARMLIB lists 43-44
 installation performance specification (IPS) 52-62
 definition 126
 IEAIPSxx list 60-62,40
 parameters 60-62
 tuning example 68-72
 using SMF data to adjust 87
 installation verification procedure (IVP) 18,19
 integrity, system 75-83
 areas of concern 77-81
 definition 75
 installation responsibility 76
 procedural responsibilities 82-83
 interaction 53
 definition 126
 internal reader
 definition 126
 JES2 facility 121
 internal writer
 definition 126
 interprocessor communication (IPC)
 definition 126
 INTERVAL MF/1 parameter 64
 interval service value (ISV) 61
 definition 126
 IOC resource factor coefficient 62
 IOC service definition coefficient 62,52-53
 IOCONTRL system generation macro instruction 25
 IODEVICE system generation macro instruction 21,25
 I/O device activity measurement 63
 I/O device activity report 67
 tuning example 68
 I/O device allocation 73
 I/O load balancing algorithm 50
 weighted by resource factor coefficients 62,51
 I/O measure 52
 IPC (interprocessor communication)
 definition 126
 IPL (initial program loader) 28
 IPS (see installation performance specification)
 IPS operand 101
 IPS system initialization parameter 31,29
 IQAORDER list 38
 IRBMF1xx (MF/1 parameter list) 41
 ISAM (indexed sequential access method) 21
 ISV (interval service value) 61
 definition 126
 IVP (installation verification procedure) 18,19

 JCL (job control language) 92-98
 (see also parameters, JCL)
 changes to MVT 96,93
 changes to VS2 Release 1 92-96
 for JES2 96-98
 summary of changes 93
 JCL validation exit 83
 JES2
 definition 126
 differences from MVT/HASP and VS2 Release 1/HASP 121
 generation 22-23
 highlights 10
 JCL 96-98
 operator commands 103-105
 SMF record support 90
 time sharing commands 110
 JES2 control cards 96-98
 JES2 generation 22-23
 options 22-23
 steps 22
 relationship to SYSGEN 22
 JES3
 definition 126
 highlights 10
 job class, JES2 23
 job control language (JCL) 92-98
 changes to MVT 96,93
 changes to VS2 Release 1 92-96
 for JES2 96-98
 summary of changes 93
 job entry subsystem
 data set 45,47
 generation 18-19,21-23
 highlights 10
 (see also JES2; JES3)
 job purge SMF exit 91
 job queue data set (SYS1.SYSJOBQE) 47
 unsupported MVT function 14
 unsupported VS2 Release 1 function 14
 job selection priority 97
 JOB statement 93
 job step dispatching priority default value 96
 JOBCAT catalog 113
 JOBCLASS operand 100
 jobname operand
 DISPLAY command 100
 HOLD command 100
 RELEASE command 101
 JOBPARM control card 97
 JOBS operand
 DISPLAY command 100
 TRACK command 102

 key-sequenced data set 111

 link library (SYS1.LINKLIB) 45
 adding installation-written members 27
 link library list (LNKLSTxx) 41

link pack area (LPA)
 defined 126
 fixed LPA 123
 modified LPA 123
 pageable LPA 123
 link pack area library (SYS1.LPALIB) 45
 adding installation-written members 27
 LINKLIB system generation macro instruction 27,25
 LISTCAT verb 115
 LNK system initialization parameter 31,29
 LNKSTxx (link library list) 41
 load balancing 50
 LOADER system generation macro instruction 27,25
 local lock
 definition 126
 local service
 definition 126
 local system queue area (LSQA) 123
 definition 126
 local time constant (PARMTZ) 43
 locking technique 79,118
 log data sets 47
 specification of output class with LOGCLS
 parameter 31
 specification of WTL limit with LOGLMT
 parameter 31
 LOGCLS system initialization parameter 31,29
 LOGLMT system initialization parameter 31,29
 LOGOFF command 108
 LOGON command 108
 system resources manager support 106
 loosely coupled multiprocessing (MP) 10
 LPA (link pack area)
 definition 126
 fixed LPA 123
 modified LPA 123
 pageable LPA 123
 LPA directory load list (IEALOD00) 41
 LPA packing list (IEAPAK00) 41
 LSQA (local system queue area) 123
 definition 126
 LSQA operand
 MOUNT command 105,101
 START command 105,101
 LSQACEL system initialization parameter 35

 M operand 100
 MACLIB system generation macro instruction 27,25
 macro instructions, system generation
 AFFINITY 20,25
 ALGLIB 25
 ALGOL 25
 ASSEMBLER 25
 CENPROCS 27,25
 CHANNEL 20,25
 CHECKER 25
 CKPTREST 21,25
 COBLIB 25
 COBOL 25
 CONSOLE 21,25
 CTRLPROG 21,25
 DATAMGT 21,25
 DATASET 21,25
 DCMLIB 25
 EDIT 21,25
 EDITOR 27,25
 EMULATOR 25
 FORTLIB 25
 FORTRAN 25
 GENERATE 21,25
 GENTSO 25
 GJOBCTL 25
 GRAPHICS 27,25
 HELP 25
 IMAGELIB 25
 IOCONTRL 25
 IODEVICE 21,25
 LINKLIB 27,25
 LOADER 27,25
 MACLIB 27,25
 OUTPUT 25
 PAGE 27,25
 PARMLIB 27,25
 PL1 25
 PLILIB 25
 PTOP 25
 RESMODS 27,25
 RPG 25
 SECONSLE 27,25
 SCHEDULR 21,25
 SORTLIB 25
 SORTMERG 25
 SUPRVSOR 25
 SVCTABLE 21,25
 SYSUTILS 25
 TELCMLIB 25
 TSO 21,25
 TSOPTION 25
 UCS 27,25
 UNITNAME 21,25
 macro library (SYS1.MACLIB) 45
 adding installation-written members 27
 main storage (see real storage)
 main storage hierarchy support 13
 main storage occupancy algorithm 50
 master address space 35,36
 definition 127
 master scheduler initialization 28
 MATRIX operand of DISPLAY command 102,100
 MAXUSER system initialization parameter 32,29
 MCP (message control program)
 definition 127
 TCAM 117
 measurement facilities
 MF/1 63-72
 SMF 86-91
 measurement facility control (MFC) 63
 MEMBER MF/1 parameter 65
 MEMBERS system generation parameter 26
 message control program (MCP)
 definition 127
 TCAM 117
 MFC (measurement facility control) 63
 MF/1 (system activity measurement facility) 63-72
 definition 128
 highlights 11
 operation 63-64
 parameters 64-65
 reports 65-67
 SYS1.PARMLIB member 41
 tracing 67
 tuning example 68-72
 MF/1 parameter list (IRBMF1xx) 41
 MF/1 reports 65-67
 channel activity 67
 CPU activity 65
 I/O device activity 67
 paging activity 65-66
 workload activity 66-67
 MIC (missing interruption checker)
 definition 127
 initialization 28
 MIN system initialization parameter 35
 minimum configuration
 for multiprocessing 118
 for starter system 19
 missing interruption checker (MIC)
 definition 127
 initialization 28
 MLPA system initialization parameter 32,29
 MN operand 101
 MOD system initialization parameter 36
 modified link pack area 123
 system initialization parameter specification 31
 modified link pack area list (IEALPAXx) 43
 MODIFY command
 BRDRQ 100
 CLASS 100
 DRIVER 103,101
 HOLD 103,101
 JOBCLASS 100

NET 102,101
 OUTCLASS 100
 PAUSE 100
 procname 100
 REGSIZE 103,101
 SMF 103,101
 SUBMIT 103,101
 TCAM 103
 USERS 103,101
 MODIFY subcommand 108
 MONITOR command 105,101
 MONITOR subcommand 108
 MOUNT command 101,105
 MP (see multiprocessing)
 MPA system initialization parameter 35
 MPS system initialization parameter 36
 MSGRT command 101
 MSO service definition coefficient 62,52-53
 multiple address spaces 11
 definition 127
 specifying maximum number with MAXUSER
 parameter 32
 multiple locks 118
 multiple virtual address spaces 11
 definition 127
 specifying maximum number with MAXUSER
 parameter 32
 multiprocessing (MP), tightly coupled
 considerations for VS2 Release 2 118-119
 definition 127
 highlights 10
 operator command support 102
 SMF record support 90
 system integrity considerations 78
 multiprocessing (MP), loosely coupled 10
 multiprogramming 11
 multitasking
 considerations for a multiprocessing
 environment 78,119

 N operand 100
 NET operand 102
 DISPLAY command 100
 HALT command 102
 MODIFY command 101
 START command 101
 VARY command 102
 NIP (nucleus initialization program) 28
 NOCHAN MF/1 parameter 64
 NOCHRDR MF/1 parameter 64
 NOCOMM MF/1 parameter 64
 NOCPU MF/1 parameter 64
 NODASD MF/1 parameter 64
 NODEVICE MF/1 parameter 64
 NOGRAPH MF/1 parameter 64
 non-standard control program interface 77
 not-in-use attribute 94
 not-in-use data sets 94
 NOTE addressed to SYSIN/SYSOUT DCB 13,14
 NOOPI MF/1 parameter 65
 NOPAGING MF/1 parameter 64
 NOREPORT MF/1 parameter 64
 NOTAPE MF/1 parameter 64
 NOTRACE MF/1 parameter 64
 NOUNTR MF/1 parameter 64
 NOWKLD MF/1 parameter 64
 nucleus 123
 nucleus initialization program (NIP) 28
 nucleus library (SYS1.NUCLEUS) 45
 adding installation-written members 27
 nucleus-only generation 18,13
 NUCMAP system initialization parameter 32,29

 OPER operand 101
 operator
 commands 99-105
 control of multiprocessing configurations 118-119
 reducing interaction during system initialization 37
 OPERATOR command 108

operator commands 99-105
 job entry subsystem support 103-105
 multiprocessing support 102
 summary of changes 100-102
 system resources management support 103
 time sharing support 103
 VTAM support 102
 operator intervention restriction 32,29
 OPI MF/1 parameter 65
 OPI system initialization parameter 32,29
 OPT system initialization parameter 32,29
 ordered seek queueing 14
 OUT operand
 CANCEL command 100
 RESET command 101
 OUTCLASS operand 100
 OUTLIM job control language parameter 95,93
 OUTPUT command 109
 job entry subsystem support 110
 OUTPUT control card 97-98
 output limit specification 95
 OUTPUT system generation macro instruction 25

 page
 definition 127
 page data sets 45
 adding paging space to the system 48
 change from VS2 Release 1 47
 definition 127
 highlights 12
 naming (PAGE parameter) 32
 page fault
 definition 127
 page fixing
 definition 127
 page frame
 definition 127
 MF/1 counts 66
 page frame table entries 123
 page-in
 counts in SMF records 86
 definition 127
 rates and percentages in MF/1 reports 66
 page-out
 counts in SMF records 86
 definition 127
 rates and percentages in MF/1 reports 66
 page reclamation
 definition 127
 rates and percentages in MF/1 reports 66
 page replacement algorithm 51
 page stealing 51
 definition 127
 PAGE system generation macro instruction 27,25
 PAGE system initialization parameter 32,29
 adding paging space to the system 48
 difference from VS2 Release 1 34
 page table (PGT)
 definition 127
 page translation exception
 definition 127
 pageable BLDL table 123
 system initialization parameter specification 30
 pageable link pack area 123
 creation using the CLPA parameter 30
 SYS1.LPALIB 45
 transient area replacement 14
 pageable real storage counts 66
 paging
 definition 127
 replacement for rollout/rollin 13
 replacement for scatter load 13
 paging activity measurement 63,65-66
 paging activity report 65-66
 paging device
 definition 127
 PAGING MF/1 parameter 64
 paging rate
 MF/1 counts 66
 definition 127

paging space (page data sets) 45
 adding paging space to the system 48
 change from VS2 release 1 47
 definition 127
 highlights 12
 naming (PAGE parameter) 32
 PAL system initialization parameter 35
 parameters, JCL
 ADDRSPC 96,93
 AFF 93
 COPIES 95,93
 DCB 93
 DEST 95,93
 DPRTY 96,93
 DYNAM 93,94
 DYNAMNBR 94,93
 FREE 94,93
 HIARCHY 93
 OUTLIM 95,93
 PERFORM 93,92
 PRTY 96-97,92,93
 REGION 95-96,93
 RETAIN 93
 ROLL 93
 SEP 93
 SPACE 93
 SPLIT 93
 SUBALLOC 93
 SYSABEND 40
 SYSUDUMP 40
 UNIT 93
 parameters, MF/1 64-65
 example 68-70
 parameters, system initialization
 ALTSYS 36
 APF 30,29
 APG 30,29
 AUXLIST 35
 BLDL 30,29
 BLDLF 30,29
 CLPA 30,29
 CMD 30,29
 CPQE 35
 CSA 30,29
 CVIO 31,29
 DUMP 31,29
 FIX 31,29
 HARDCPY 31,29
 HRAM 36
 HSVC 36
 IPS 31,29
 LNK 31,29
 LOGCLS 31,29
 LOGLMT 31,29
 LSQACEL 35
 MAXUSER 32,29
 MIN 35
 MLPA 32,29
 MOD 36
 MPA 35
 MPS 36
 NUCMAP 32,29
 OPI 32,29
 OPT 32,29
 PAGE 32,29
 PAL 35
 QBF 36
 RAM 36
 REAL 32,29
 RERP 36
 RSVC 36
 SMF 33,29
 SQA 33,29
 SQACEL 35
 SQS 36
 SYSP 33,29
 TMSL 35,36
 TRACE 35
 TSOAX 35
 VAL 34,29
 VRREGN 34,29
 WTOBFRS 34,29
 WTORPLY 34,29
 PARMLIB system generation macro instruction 27,25
 PARM TZ (local time constant) 43
 password protection
 PASSWORD data set 46
 system data sets 82,76
 system libraries 82
 tape data sets 82-83
 VSAM catalog 112
 PATH operand 102
 PAUSE operand 100
 PDS system generation parameter 26
 percent channel active value 67
 percent channel busy and CPU waiting value 67
 PERFORM job control language parameter 93,92
 performance groups 53
 definition 127
 performance group number 61,53
 JCL specification 92
 time sharing LOGON command 106
 performance group periods 53-61
 definition 127
 tuning example 68-72
 performance group period report 66
 performance group report 66-67
 performance objective number 62,52-60
 performance objectives 53-60
 definition 127
 effect of changing 71
 motive for specifying 53
 tuning example 68-72
 performance specifications (see installation performance specification)
 period duration parameter 61
 PGT (page table)
 definition 127
 physical addresses for devices 19
 physical environment, computing system 76
 PL1 system generation macro instruction 25
 PL/1 F
 compiler 15
 library 15
 syntax checker 15
 PL1LIB system generation macro instruction 25
 POINT addressed to SYSIN/SYSOUT DCB 13,14
 pre-fixed storage area (PSA) 123
 PRESRES list 44,38
 PRINT verb 115
 priority I/O queueing 14
 priority, job selection 97
 priority, job step dispatching 96
 private address space
 definition 127
 size recorded by SMF 90
 PROC operand 101
 procedure library (SYS1.PROCLIB) 46
 adding installation-written members 27
 processor-only generation 18
 procname operand 100
 PROFILE subcommand 108
 EDIT command 110
 program conversion 117
 programming conventions 117
 protected system control block 77-78
 PRTY job control language parameter 96-97,92,93
 PRTY operand 101
 PSA (pre-fixed storage area) 123
 pseudo device 121
 PTOP system generation macro instruction 25

 Q operand
 HOLD command 100
 RELEASE command 101
 SET command 101
 QBF system initialization parameter 36
 queued telecommunications access method (QTAM) 13
 QUIESCE command 102,101
 QTAM (queued telecommunications access method) 13

R operand 100
RAM system initialization parameter 36
reader
 internal 126,121
 OS 13,14
real address
 definition 127
 range set for multiprocessing 118-119
real storage
 ADDRSPC=REAL 117
 definition 127-128
 installation's influence of importance 53
 multiprocessing considerations 118-119
 MVT 11
 paging activity report of use 65-66
 shared 10
real storage management (RSM)
 definition 128
real storage occupancy
 basic processing resource 52
 resource-use algorithm 50
 service definition coefficient MSO 62
 service received during interval of 62
REAL system initialization parameter 32,29
 difference from VS2 Release 1 35
reconfiguration 118-119
 definition 128
recovery routine
 definition 128
recovery management routines 123
recovery termination manager
 definition 128
 initialization 28
refreshable
 definition 128
 modules in the pageable link pack area 45
REGION job control language parameter 95-96,93
REGSIZE operand 101
RELEASE command
 jobname 101
 Q 101
 TP 101
relocate hardware (dynamic address translation)
 definition 125
remote job entry (RJE) 13
remote terminal access method (RTAM)
 definition 128
remote terminal programs 23
REPORT MF/1 parameter 64
REPRO verb 115
RERP system initialization parameter 36
reserved mount characteristic 74
RESET command
 CLASS 101
 jobname 101
 OUT 101
 PRTY 101
resident BLDL list (IEABLDxx) 41-42
RESMODS system generation macro instruction 27,25
resource factor coefficients 62
 use of 51
resource identification 79
resource manager constants 62
resource-use algorithms 50-51
resource-use routines 50-51
response/throughput bias (RTB) 61
 definition 128
response time 53
 approximating from SMF records 87
 MF/1 tuning example 68-72
RETAIN job control language parameter 93
RJE (remote job entry) 13
RMTGEN 23
 definition 128
 options 23
ROLL job control language parameter 93
rollout/rollin 13
RPG system generation macro instruction 25
RSM (real storage management)
 definition 128
RSVC system initialization parameter 36
RTAM (remote terminal access method)
 definition 128
RTB (response/throughput bias) 61
 definition 128
RUN subcommand 108

SAMG (system activity measurement gathering) 63
SARG (system activity report generation) 63
satellite graphic job processor (SGJP) 13
SAVE operand 101
scatter load 13
scheduler work area (SWA) 123
 definition 128
 SYS1.SYSJOBQE replacement 14,47
SCHEDULR system generation macro instruction 21,25
secondary storage (auxiliary storage)
 definition 125
 paging activity report of use 65-66
SECONDSLE system generation macro instruction 27,25
security, system 82-83
segment
 definition 128
segment table (SGT)
 definition 128
SEND operator command 105
 BRDCAST 101
 CN 101
 OPER 101
 SAVE 101
SEND subcommand
 EDIT command 110,108
 OPERATOR command 109
SEND time sharing command 109
SEP job control language parameter 93
serialization techniques 78-79
 disablement 78
 ENQ/DEQ 79,119
 locking 79
 WAIT/POST 119
service definition coefficients 62,60
 use of 52-53
service rate 52-60
 definition 128
 effect of change for a performance objective 71
 IPS specification 62
 tuning example 68-72
 workload activity reports 66-67
service units 52-60
 definition 128
 MF/1 workload activity reports 66-67
 SMF record counts 86
 tuning example 68-72
SER0 routine 13
SER1 routine 13
SET command
 used to identify the CPU with an acceptable clock value 102
 used to restore time zone constant to IPL value 102
 used to specify an IPS 101
SET operand 100
SGJP (satellite graphic job processor) 13
SGT (segment table)
 definition 128
SIO instruction counts 67
slot
 definition 128
 MF/1 counts 66
SMF (system management facilities) 86-91
 differences from MVT 86
 differences from VS2 Release 1 86
 exits 91
 records 86-91
 system initialization parameter specification 33
SMF data sets (SYS1.MANX/SYS1.MANY) 46
 change from MVT 48
 change from VS2 Release 1 47
SMF operand 101
SMF parameter list (SMFPRMxx) 42
SMF system initialization parameter 33,29
SMFDEFLT parameter list 42,38

SMFPRMxx (SMF parameter list) 42
software recording facility
 definition 128
SORT/MERGE 15
SORTLIB system generation macro instruction 25
SORTMERG system generation macro instruction 25
SPACE job control language parameter 93
SPACE system generation macro parameter 26
spin lock
 definition 128
SPLIT job control language parameter 93
SQSA (system queue area) 123
 definition 128
 system initialization parameter specification 33
SQA operand 100
SQA system initialization parameter 33,29
SQACEL system initialization parameter 35
SQS system initialization parameter 36
stage I 18
stage II 18
standard control program interface 77
START command
 BRDR 103
 LSQA 105,101
 NET 102,101
 TS 103
START operand 102
starter system 19
 minimum configuration 19
 physical addresses for devices 19-20
STPCAT catalog 113
STIDP instruction 36
STOP command
 BRDR 103
 STOP MF/1 parameter 65
 STOPMN subcommand 108
 STOPTR command 105,102
 STORAGE operand 102
 storage map, virtual 123
 store protection 83
SUBALLOC job control language parameter 93
SUBMIT command 109
 job entry subsystem support 110
SUBMIT operand 101
suspend lock
 definition 128
SUPRVSOR system generation macro instruction 25
SVC library (SYS1.SVCLIB) 45
 adding installation-written members 27
SVC routines calling SVC routines 80
SVC routines, installation-written
 integrity considerations 81
 SYSGEN specification with SVCTABLE macro 21
SVCTABLE system generation macro instruction 21,25
SWA (scheduler work area) 123
 definition 128
 replacement for SYS1.SYSJOBQE 14,47
swap counts 66
swapping
 definition 128
 ERV effect on 62
 ISV effect on 61
 use by the system resources manager 52
symmetric processors
 definition 128
symmetric storage configuration
 definition 128
symmetrical I/O 118
SYNC subcommand 116,107
SYSABEND job control language parameter 40
SYSCTLG data set 47
 conversion to VSAM format 111-116
SYSDSN major name 79
SYSGEN (see system generation)
SYSIEA01 major name 79
SYSIEECT major name 79
SYSIEFSD major name 79
SYSOUT MF/1 parameter 64
SYSP system initialization parameter 33,29
SYSPSWRD major name 79
system activity measurement facility (MF/1) 63-72
 definition 128
 highlights 11
 operation 63-64
 parameters 64-65
 reports 65-67
 SYS1.PARMLIB member 41
 tracing 67
 tuning example 68-72
system activity measurement gathering (SAMG) 63
system activity report generation (SARG) 63
system area 123
system catalog (SYSCTLG) 47
 conversion to VSAM format 111-116
 highlights 10
system data sets 45-48
 changes to MVT 48
 changes to VS2 Release 1 47-48
 optional 46-47
 password protection 82
 required 45-46
system definition 17-48
system environment recording routines (SER0 and SER1) 13
system generation
 changes 24-27
 JES2 generation 22-23
 job entry subsystem generation 18,22-23
 macro instructions (see macro instructions, system generation)
 procedures 18-27
 RMTGEN 23
 starter system 19-20
 unsupported VS2 Release 1 macro instructions 27
system initialization 28-44
 parameters 28-37
 (see also parameters, system initialization)
 system parameter library (SYS1.PARMLIB) 37-44
system integrity 75-83
 areas of concern 77-81
 definition 75
 installation responsibility 76
 procedural responsibilities 82-83
system key 77
 definition 128
system parameter library (SYS1.PARMLIB) 37-44
 IBM-created lists 39
 installation-created lists 43
 required system data set 45
 summary of changes 38
system management facilities (SMF) 86-91
 differences from MVT 86
 differences from VS2 Release 1 86
 exits 91
 records 86-91
 system initialization parameter specification 33
system parameter list (IEASYSxx) 42
system parameter specification 35
system protection key 77
 definition 128
system queue area (SQA) 123
 definition 128
 system initialization parameter specification 33
system resources management 50-62,68-72
 definition 128
 highlights 11
 JCL support 92
 operator command support 103
 parameters 60-62
 SMF records 86-87
 time sharing command support 106
 tuning example 68-72
system resources management tuning parameter list (IEAOPTxx) 62,43-44
system restart 31
system security 82-83
system summary report 67
SYSUDUMP job control language parameter 40
SYSUTILS system generation macro instruction 25
SYSVTOC major name 79

SYS1.ACCT data set 47
 unsupported MVT function 14
 SYS1.BROADCAST data set 46
 formatting 116
 mounted for system initialization 110
 new format description 47
 SYS1.CMDLIB data set 46
 adding installation-written members 27
 SYS1.DCMLIB data set 46
 SYS1.DSSVM data set 45
 change from MVT 48
 change from VS2 Release 1 47
 specifying creation of a new nucleus map with the
 NUCMAP parameter 32
 SYS1.DUMP data set 46
 device specification by a system initialization
 parameter 31
 maximum number specified 34
 SYS1.HELP data set 46
 SYS1.IMAGELIB data set 46
 adding installation-written members 27
 use for 3886 OCR devices 95
 SYS1.LINKLIB data set 45
 adding installation-written members 27
 SYS1.LOGREC error recording data set 45
 change from MVT and VS2 Release 1 47
 SYS1.LPALIB data set 45
 adding installation-written members 27
 SYS1.MACLIB data set 45
 adding installation-written members 27
 SYS1.MANX data set 46
 change from MVT 48
 change from VS2 Release 1 47
 SYS1.MANY data set 46
 change from MVT 48
 change from VS2 Release 1 47
 SYS1.NUCLEUS data set 45
 adding installation-written members 27
 SYS1.PAGE data set 47
 SYS1.PARMLIB (system parameter library) 37-44
 IBM-created lists 39
 installation-created lists 43
 required system data set 45
 summary of changes 38
 SYS1.PARMLIB lists
 changes to 38
 IEAABDxx 40
 IEAAPFxx 40
 IEAAP00 39
 IEABLDxx 41-42
 IEADMPxx 40
 IEAFIXxx 40
 IEAIPSxx 40
 IEALOD00 41
 IEALPAXx 43
 IEAOPTxx 43-44
 IEAPAK00 41
 IEASYSxx 42
 IQAORDER 38
 IRBMF1xx 41
 LNKLSTxx 41
 PARMTZ 43
 PRESRES 44,38
 SMFDEFLT 42,38
 SMFPRMxx 42
 VATLSTxx 44
 SYS1.PROCLIB data set 46
 adding installation-written members 27
 SYS1.SAMPLIB data set 46
 SYS1.STGINDEX data set 46
 SYS1.SVCLIB data set 45
 adding installation-written members 27
 SYS1.SYSJOBQE data set 47
 unsupported MVT function 14
 unsupported VS2 Release 1 function 14
 SYS1.SYSVLOGX (see log data sets)
 SYS1.SYSVLOGY (see log data sets)
 SYS1.TELCMLIB data set 46-47
 adding installation-written members 27
 SYS1.UADS (user attribute data set) 46
 formatting 116
 mounted for system initialization 110
 new format description 48
 SYS1.VTAMLIB data set 47
 adding installation-written members 27
 SYS1.VTAMLST data set 47
 adding installation-written members 27
 SYS1.VTAMOBJ data set 47

 T operand 100
 tape data set password protection 82-83
 TAPE MF/1 parameter 64
 TCAM (telecommunications access method)
 conversion considerations 117
 system generation specification 21
 TELCMLIB system generation macro instruction 25
 telecommunications access method (TCAM)
 conversion considerations 117
 system generation specification 21
 telecommunications library (SYS1.TELCMLIB) 46-47
 adding installation-written members 27
 teleprocessing applications program 53
 TESTRAN program 14
 throughput 61
 tightly coupled multiprocessing (MP)
 considerations for VS2 Release 2 118-119
 definition 127
 highlights 10
 operator command support 102
 SMF record support 90
 TIME command 109
 time dependency 117
 time-of-check-to-time-of-use problem 79
 time sharing
 accounting support 110
 allocation support 106
 command changes 107-109
 edit support 110
 job entry subsystem support 110
 SMF record elimination 90
 system resources manager support 106
 time sharing integration support 110
 VSAM support 106
 time slicing 14
 TMSL system initialization parameter 35,36
 TOD clock messages 28
 TP operand
 HOLD command 100
 RELEASE command 100
 trace 14
 TRACE MF/1 parameter 64
 TRACE system initialization parameter 35
 trace, system option
 during system initialization 35
 unsupported MVT and VS2 Release 1 function 14
 trace, TSO 14
 tracing, MF/1 67
 SMF records 86
 TRACK command 105
 A 102
 INIT 102
 JOBS 102
 TS 102
 transaction 53
 definition 129
 transient areas 14
 translation tables 129
 definition 129
 TS operand
 DISPLAY command 100
 TRACK command 102
 TSO (see time sharing)
 TSO driver 14
 TSO system generation macro instruction 21,25
 TSO trace 14
 TSOAUX system initialization parameter 35
 TSOPTION system generation macro instruction 25

turnaround time 53,68
 tuning parameter list (IEAOPTxx) 62,43-44
 tuning with MF/1 68-72
 type I programs 15

u operand 100
 UADS (see SYS1.UADS)
 UADSREFM utility program 116
 UCS system generation macro instruction 27,25
 unallocating data sets (see freeing data sets)
 uniprocessing 129
 UNIT job control language parameter 93
 UNITNAME system generation macro instruction 21,25
 defining device groupings 74
 specifying VIO processing 94
 UNTR MF/1 parameter 64
 units code parameter 61
 unserialized path through I/O device allocation 73
 unsupported MVT functions
 job control language summary 93
 operator command summary 100-102
 SMF record summary 88-89
 summary 13-14
 system data sets 48
 system initialization parameters 35-36
 time sharing command summary 107-109
 unsupported VS2 Release 1 functions
 job control language summary 93
 operator command summary 100-102
 SMF record summary 88-89
 summary 14
 system data sets 47-48
 system generation macro instructions 27
 system initialization parameters 35
 time sharing command summary 107-109
 user address space (see private address space)
 user area 123
 user attribute data set (SYS1.UADS) 46
 formatting 116
 mounted for system initialization 110
 new format description 48
 user key 81
 USER operand 100
 user protection key 81
 USERID operand 100
 user-supplied addresses 77
 USERS operand 101
 utility programs
 IEHIOSUP 13
 IEHLIST 113,13,14
 IEHMOVE 113,13,14
 IEHPROGM 113,26,13,14
 IEHUCAT 116,113

V=R (virtual equals real) address area
 ADDRSPC JCL parameter specification 96
 candidate programs 117
 definition 129
 REAL system initialization specification 32
 VRREGN system initialization specification 34
 VAL system initialization parameter 34,29
 validity checking 76
 fetch/store operations 79
 SVC routines calling SVC routines 80
 unauthorized program requests for authorized library
 routines 80
 user-supplied addresses 77
 VARY command 118
 CHANNEL 102
 CPU 102
 NET 102
 PATH 102
 STORAGE 102
 VATLSTxx (volume attribute list) 44
 VIO (virtual I/O)
 catalog for data sets (STGINDEX) 46
 definition 129
 highlights 10
 JCL requirements 94
 MF/1 report data 66

 SMF record support 86
 SYSGEN specification with the UNITNAME
 macro 21
 system initialization parameter specification 31
 virtual address
 definition 129
 virtual equals real address area (V=R)
 ADDRSPC JCL parameter specification 96
 candidate programs 117
 definition 129
 REAL system initialization specification 32
 VRREGN system initialization specification 34
 virtual I/O (VIO)
 catalog for data sets (STGINDEX) 46
 definition 129
 highlights 10
 JCL requirements 94
 MF/1 report data 66
 SMF record support 86
 SYSGEN specification with the UNITNAME
 macro 21
 system initialization parameter specification 31
 virtual memory (virtual storage)
 definition 129
 highlights 11
 virtual storage
 definition 129
 highlights 11
 virtual storage access method (VSAM)
 catalog 111-116
 definition 129
 highlights 10
 SMF record support 90
 time sharing support 106
 virtual storage management (VSM)
 definition 129
 virtual storage map 123
 virtual telecommunications access method (VTAM)
 command support 100-102
 highlights 10
 system data sets 47
 volume attribute list (VATLSTxx) 44
 volume sharing 95
 volume table of contents (VTOC) 111
 VRREGN system initialization parameter 34,29
 VSAM (virtual storage access method)
 catalog (see VSAM catalog)
 definition 129
 highlights 10
 SMF record support 90
 time sharing support 106
 VSAM catalog 111-116
 creating 115-116
 master catalog 45
 relationships between OS and VSAM 112
 search strategy 113
 structure 111-112
 updating 115-116
 use of access method services 115-116,113
 VSM (virtual storage management) 129
 VTAM (virtual telecommunications access method)
 command support 100-102
 highlights 11
 system data sets 47
 VTAM data sets
 SYS1.VTAMLIB 47
 SYS1.VTAMLST 47
 SYS1.VTAMOBJ 47
 VTOC (volume table of contents) 111

WKLD MF/1 parameter 64
 workload activity measurement 63
 workload activity reports 66
 performance group period report 66
 performance group report 66-67
 system summary report 67
 workload level 54-61
 function of initiated jobs 58
 MF/1 data in reports 66-67

tuning example 68-72
used to express service relationships between groups of
users 54
workload management 50
definition 129
WRITELOG command 102
writer
internal 126
OS 13,14
WTOBFRS system initialization parameter 34,29
WTORPLY system initialization parameter 34,29

XDAP addressed to SYSIN/SYSOUT DCB 13,14
XPT (external page table)
definition 126

2-channel switches 118
3886 OCR device 95

Your views about this publication may help improve its usefulness; this form will be sent to the author's department for appropriate action. Using this form to request system assistance or additional publications will delay response, however. For more direct handling of such requests, please contact your IBM representative or the IBM Branch Office serving your locality.

Possible topics for comment are:

Clarity Accuracy Completeness Organization Index Figures Examples Legibility

Cut or Fold Along Line

What is your occupation? _____
Number of latest Technical Newsletter (if any) concerning this publication: _____
Please indicate in the space below if you wish a reply.

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. Elsewhere, an IBM office or representative will be happy to forward your comments.

Cut or Fold Along Line

Your comments, please . . .


This manual is part of a library that serves as a reference source for system analysts, programmers, and operators of IBM systems. Your comments on the other side of this form will be carefully reviewed by the persons responsible for writing and publishing this material. All comments and suggestions become the property of IBM.

Fold

Fold

First Class
Permit 81
Poughkeepsie
New York

Business Reply Mail
No postage stamp necessary if mailed in the U.S.A.



Postage will be paid by:

International Business Machines Corporation
Department D58, Building 706-2
PO Box 390
Poughkeepsie, New York 12602

Fold

Fold

OS/VS2 Planning Guide for VS2 Printed in U.S.A. GC28-0667-0



International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, New York 10604
(U.S.A. only)

IBM World Trade Corporation
821 United Nations Plaza, New York, New York 10017
(International)