

IBM / Technical Newsletter

This Newsletter No. GN28-1311
Date June 30, 1989
Base Publication No. GC28-1828-1
File No. S370-34
Prerequisite Newsletters/
Supplements None

MVS/ESA SPL: Initialization and Tuning

©Copyright IBM Corp. 1988

MVS/System Product - JES3, Program No. 5685-002

MVS/System Product - JES2, Program No. 5685-001

This newsletter contains replacement pages for *MVS/ESA SPL: Initialization and Tuning*. Before inserting any of the attached pages into *MVS/ESA SPL: Initialization and Tuning*, read *carefully* the instructions on this cover. They indicate when and how you should insert pages.

Pages to be Removed

Cover - Edition Notice

iii - vi

3-125 - 3-126

3-141 - 3-142

5-33 - 5-34

5-75 - 5-76

5-125 - 5-128

Attached Pages to be Inserted*

Cover - Edition Notice

iii - vi

3-125 - 3-126

3-141 - 3-142.2

5-33 - 5-34

5-75 - 5-76.2

5-125 - 5-128

*If you are inserting pages from different Newsletters/Supplements and *identical* page numbers are involved, always use the page with the latest date (shown in the slug at the top of the page). The page with the latest date contains the most complete information.

A change to the text or to an illustration is indicated by a vertical line to the left of the change.

Summary of Changes

See the Summary of Changes for a list of the changes included in this newsletter.

Note: *Please file this cover letter at the back of the publication to provide a record of changes.*

IBM Corporation, Information Development, Dept. D58, Building 921-2
P.O. Box 950, Poughkeepsie, New York 12602

...

...

...

...

...

...

...

...

...



MVS/ESA
System Programming Library:
Initialization and Tuning

GC28-1828-1

MVS/System Product:
JES2 Version 3
JES3 Version 3

Second Edition (November, 1988)

This is a major revision of, and obsoletes, GC28-1828-0. See the Summary of Changes regarding new and changed information made to this publication. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

This edition with Technical Newsletter GN28-1311 applies to Version 3 of MVS/System Product (5685-001 or 5685-002) and to all subsequent releases until otherwise indicated in new editions or Technical Newsletters. Changes are made periodically to the information herein; before using this publication in connection with the operation of IBM systems, consult the latest *IBM System/370 Bibliography*, GC20-0001, for the editions that are applicable and current.

References in this publication to IBM products or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product in this publication is not intended to state or imply that only IBM's product may be used. Any functionally equivalent product may be used instead. This statement does not expressly or implicitly waive any intellectual property right IBM may hold in any product mentioned herein.

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

A form for reader's comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, Information Development, Department D58, Building 921-2, PO Box 950, Poughkeepsie, NY 12602. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Summary of Changes

**Summary of Changes
for GC28-1828-1
as Updated June 30, 1989
by Technical Newsletter GN28-1311**

This technical newsletter contains maintenance changes for the following:

- A new parameter in IEASYSxx, MAXCAD=nn, allows you to reserve the number of entries available for SCOPE=COMMON data spaces on all primary address space access lists in the system.
- Support for the following processors:
 - Model 3090 100S
 - Model 3090 250S
 - Model 3090 380S
 - Model 4381 90E

**Summary of Changes
for GC28-1828-1
MVS/System Product Version 3 Release 1.0e**

This book contains information previously presented in *MVS/Extended Architecture SPL: Initialization and Tuning*, GC28-1828-0. The following summarizes the changes to that information.

New Information:

- Library lookaside facility (LLA) now supports production datasets as well as LNKLST data sets.
- SYS1.PARMLIB member CSVLLAxx is updated with new control statements to support LLA.
 - REMOVE allows you to dynamically remove libraries from the list of libraries managed by LLA.
 - MEMBERS allows you to refresh dynamically specified members in LLA-managed production data sets.
 - -LNKLST- is a shorthand method to specify the whole LNKLST concatenation.
 - PARMLIB(dsn) SUFFIX(xx) allows you to specify another parmlib member that contains more LLA control statements to be processed.
 - FREEZE | NOFREEZE establishes whether the DASD directory or the LLA maintained directory is used for the LLA-managed data set.
- The GRSCNFxx parmlib member is updated with 3 new parameters.
 - REJOIN indicates whether the system can automatically rejoin the active ring.
 - TOLINT specifies in seconds the tolerance interval for the expected RSA-message to return.
 - ACCELSYS specifies the threshold for GRS ring acceleration.

Changed Information:

- The linklist lookaside facility (LLA) is renamed to library lookaside facility (LLA).
- The SYS1.PARMLIB member CSVLLAxx is updated with a changed control statement to support LLA.

LIBRARIES list the names of the LNKLST and production libraries to be managed by LLA.

- The value range was increased for the CWSS and PWSS parameters in the IEAIPSxx member.

Deleted Information:

- Using GTF to track sysevents and the input information for each sysevent has been deleted from *Initialization and Tuning* and moved to *Component Diagnosis and Logic: System Resource Manager, LY28-1592*.

**Summary of Changes
for GC28-1828-0**
MVS/System Product Version 3 Release 1.0

This book contains information previously presented in *MVS/Extended Architecture SPL: Initialization and Tuning*, GC28-1149-5. The following summarizes the changes to that information.

New Information

Information to support MVS™ System Product Version 3 includes

- A new parameter is added to the SVC Parm statement in IEASVCxx. The new parameter supports SVCs issued in access register ASC mode.
- Increase to the value range for the CWSS and PWSS parameters in the IEAIPSxx member.
- A new system library, SYS1.MIGLIB, to be concatenated after SYS1.LINKLIB
- A new SYS1.PARMLIB member COFVLFxx to support virtual lookaside facility (VLF).
- A new SYS1.PARMLIB member CSVLLAxx to support selective refresh of the linklist lookaside (LLA).
- A new SYS1.PARMLIB member EXSPATxx for automatic responses to certain types of excessive spin loops.
- New IPCS support is available in the BLSCECT member of SYS1.PARMLIB.
- Nine new parameters are available to assist in detecting missing interrupt conditions. See the IECIOSxx member.
- A new parameter in IEASYSxx, NSYSLX, allows you to specify the number of slots in the system function table to be reserved for system linkage indexes.

- A new sysevent AVAILPUP (49 hex) to reserve processor storage frames for dumping services.
- A new sysevent CPUTCONV (4A hex) to convert CPU seconds to service units.

Information to support MVS/ESA™ DFP Version 3 Release 1.0 when it becomes available includes

- Support for a new record type in the IEFSSNxx member
- A new SYS1.PARMLIB member, IGDSMSxx, which contains the parameters that initialize the storage management subsystem and define it to the system.
- Descriptions of the new SET SMS and SETSMS commands that change the definitions of the storage management subsystem.

The edition also contains MVS/SP™ Version 3 support that will be effective with MVS/ESA DFP Version 3 Release 1.0.. This information includes a new way of using the VATLSTxx member to define your DASD volumes to the system:

- Use the new VATDEF statement to establish the installation's default use attributes for permanently-resident DASD volumes.
- Use generic volume serial numbers and device types for groups of DASD volumes that are mounted on devices that have the same device type and the same mount and use attributes.

Changed Information:

The following information is changed for MVS/SP Version 3

- The value range for the SRV keyword in the IEAIPSxx member increases to 999999999.
- Clarification and updates are included for the GTFPARM member.

Deleted Information

MVS/SP Version 3 does not support print dump. References to print dump, AMDPRDMP, have been removed.

Parameter	Use of the Parameter
APF	Names the parmlib member (IEAAPFxx) that contains authorized data set names.
APG	Specifies the priority value of the automatic priority group for use by the system resource manager, if one is not specified in the selected IPS.
CLOCK	Completes the name of the parmlib member (CLOCKxx) that prompts the operator to initialize the TOD clock during NIP and specifies the difference between the local time and Greenwich Mean Time (GMT).
CLPA	Tells NIP to load the link pack area with the modules contained in the LPALST concatenation. It also purges VIO data set pages that were used in the previously initialized system. Thus, CLPA implies CVIO.
CMB	Specifies the I/O device classes for which measurement data is to be collected, in addition to the DASD and tape device classes.
CMD	Completes the name of the parmlib member (COMMNDxx) that contains commands to be issued internally during master scheduler initialization.
CON	Completes the name of the parmlib member (CONSOLxx) that centralizes control of the console configuration for your installation. CONSOLxx also contains the initialization values for communication tasks, the characteristics for the hardcopy log, and default routing codes for messages that do not have routing information.
CSA	Specifies the sizes of the virtual common service area and extended common service area in multiples of 1K bytes.
CVIO	Deletes previously used VIO data set pages from the paging space. This parameter is automatically included when CLPA is specified.
DUMP	Specifies whether SYS1.DUMP data sets for SVC dump are to be on direct access device(s) or tape devices. This parameter can also indicate that <i>no</i> SYS1.DUMP data sets are to be made available for SVC dumps.
DUPLEX	Specifies a duplex page data set name or overrides the existing duplex data set name. This parameter is ignored on quick starts and warm starts (non-CLPA IPLs).
FIX	Completes the name of one or more parmlib members (IEAFIXxx) that contain names of modules from SYS1.SVCLIB, the LNKLST concatenation, and the LPALST concatenation that are to be placed in a fixed LPA that lasts for the duration of the IPL.
GRS	Specifies whether the system is to start a global resource serialization complex, join an existing complex, or not be part of a complex.
GRSCNF	Completes the name of the parmlib member (GRSCNFxx) that contains the information needed to initialize a system that is to be part of a global resource serialization complex.
GRSRNL	Completes the name of one or more parmlib members (GRSRNLxx) that contain resource name lists (RNLs).
ICS	Completes the name of the parmlib member (IEAICSxx) that contains the installation control specification. The installation control specification is used by the system resources manager (SRM) to assign performance groups.
IOS	Completes the name of the parmlib member (IECIOSxx) that contains missing interrupt handler (MIH) statements used to modify MIH time intervals and HOTIO statements used to modify recovery actions specified in the hot I/O detection table (HIDT).
IPS	Completes the name of the parmlib member (IEAIPSxx) from which the system resources manager (SRM) will obtain the installation performance specification.
LNK	Completes the name of one or more parmlib members (LNKLSTxx) that contain names of data sets that are to be concatenated to SYS1.LINKLIB.
LNKAUTH	Specifies whether all data sets in the LNKLST concatenation are to be treated as APF authorized or whether only those that are named in the APF table are to be treated as APF authorized.
LOGCLS	Specifies the JES output class for the log data sets.
LOGLMT	Specifies the maximum number of WTLs (messages) for a log data set. When the limit is reached, the data set is scheduled for sysout processing.
LPA	Completes the name of one or more parmlib members (LPALSTxx) that contain names of data sets that are to be concatenated to SYS1.LPALIB for the purpose of building the pageable LPA (PLPA and extended PLPA).
MAXCAD	Specifies the maximum number of SCOPE = COMMON data spaces to be allowed during an IPL.
MAXUSER	Specifies a value that the system uses (along with the RSVSTRT and RSVNONR parameter values) to build the address space vector table (ASVT). The system uses the ASVT to locate the various address space control blocks.

Figure 3-7 (Part 1 of 2). Overview of IEASYSxx Parameters

Parameter	Use of the Parameter
MLPA	Completes the name of one or more parmlib members (IEALPAxx) that names modules from SYS1.SVCLIB, the LNKLST concatenation, and the LPALST concatenation that are to be placed in a modified LPA that lasts for the duration of the IPL.
MSTRJCL	Completes the name of MSTRJCLxx, a module in SYS1.LINKLIB, that contains the JCL used to start the master scheduler address space.
NONVIO	Designates one or more local page data sets that are not to be used for VIO paging.
NSYSLX	Specifies the number of slots in the system function table to be reserved for system linkage indexes (LXs).
OPI	Indicates whether the operator is to be allowed to override particular parameters, or all parameters, contained in IEASYSxx.
OPT	Completes the name of a parmlib member (IEAOPTxx) that contains parameters to be used by various algorithms of the system resources manager.
PAGE	Gives the names of new page data sets to be used as additions to or replacements for existing page data sets. The first-named data set is used for the PLPA and extended PLPA pages. The second-named data set is used for MLPA and CSA. The third and all subsequently named data sets are used as local page data sets. Replacement is possible only if the parameter is placed in IEASYSxx, and the operator selects this member by entering SYSP=xx. The PAGE parameter, when specified by the operator, can only add temporarily (until the next cold or quick start) to a parmlib page data set list.
PAGTOTL	Specifies the total number of page and swap data sets that can be allocated for the life of the IPL.
PAK	Completes the name of one or more parmlib members (IEAPAKxx) that contain groups of names of modules in the LPALST concatenation that are executed together or in sequence.
PURGE	Demounts all mass storage system volumes.
RDE	Specifies that the reliability data extractor feature is included. For information on RDE, see <i>SYS1.LOGREC Error Recording</i> .
REAL	Specifies the maximum amount of real storage, in 1K blocks, that can be allocated for concurrent ADDRSPC=REAL jobs.
RER	Specifies that the reduced error recovery procedures for magnetic tapes are in effect if they are included on the OPTCD parameter of a data definition (DD) statement or on the DCB macro instruction. If the DD statement or the DCB macro does not specify the reduced error recovery procedures, all requests for them are ignored.
RSU	Specifies the number of storage units to be made available for storage reconfiguration.
RSVNONR	Specifies the number of ASVT entries to be reserved for replacing those entries marked non-reusable for the duration of an IPL.
RSVSTRT	Specifies the number of address space vector table (ASVT) entries to be reserved for address spaces created in response to a START command.
SCH	Specifies a parmlib member (SCHEDxx) from which the master scheduler will obtain its parameters. This member centralizes control over the eligible device table, the size of the master trace table, the program properties table, and the completion codes that are eligible for automatic restart.
SMF	Specifies a parmlib member (SMFPRMxx) from which SMF will obtain its parameters.
SQA	Gives the number of 64K blocks of virtual system queue area to be created at IPL (in addition to the system's minimum virtual SQA and extended SQA).
SSN	Completes the name of the parmlib member, IEFSSNxx, that contains the information to be used in identifying subsystems that are to be initialized.
SVC	Completes the name of the parmlib member (IEASVCxx) that contains the information an installation supplies to define its own SVCs. NIP processing places these SVC in the SVCTABLE.
SWAP	Specifies the names of new swap data sets to be used as additions to or replacements for existing swap data sets.
SYSNAME	Specifies the name of the system being initialized. Global resource serialization uses this name to identify the different systems in a complex.
SYSP	Specifies one or more alternate system parameter lists (IEASYSxx) that are to be read by NIP in addition to IEASYS00. SYSP may be specified only by the operator.
VAL	Names one or more parmlib members (VATLSTxx) that contain "mount" and "use" attributes of direct access devices.
VRREGN	Gives the default real-storage region size for an ADDRSPC=REAL job step that does not have a REGION parameter in its JCL.

Figure 3-7 (Part 2 of 2). Overview of IEASYSxx Parameters

part of the LNKLST concatenation. If a library is in the LNKLST concatenation, but is not APF-authorized, referencing this library via a JOBLIB or STEPLIB DD statement causes the library to be considered unauthorized for the duration of the job or step.

LNKAUTH = $\left(\begin{array}{l} \text{LNKLST} \\ \text{APFTAB} \end{array} \right)$

Meaning and Use: This parameter specifies whether all data sets in the LNKLST concatenation are to be treated as APF-authorized when accessed as part of the concatenation or whether only those data sets named in the APF table are to be treated as APF-authorized.

Value Range: Not applicable

Default Value: LNKLST, meaning that all of the data sets in the LNKLST concatenation are to be treated as APF-authorized when accessed as part of the concatenation. If the default for the LNKAUTH system parameter is taken (LNKAUTH=LNKLST) or is specified in IEASYSxx or by the operator, libraries in the LNKLST concatenation are authorized when accessed as part of the LNKLST concatenation. If a library is in the LNKLST concatenation, but is not also APF-authorized, referencing this library via a JOBLIB or STEPLIB DD statement causes the library to be considered unauthorized for the duration of the job or step.

Associated Parmlib Member: None

LOGCLS = x

Meaning and Use: This parameter specifies the JES output class for the log data sets. A log data set is queued to this class when its WTL limit has been reached. (The limit is specified by the LOGLMT initialization parameter.)

Example: LOGCLS=L

In this example, the current log data set is queued to output class L when the limit on the number of WTLs has been reached.

If the specified LOGCLS value is invalid, or an I/O error occurs while the IEASYSxx member is being read, master scheduler initialization prompts the operator for a replacement LOGCLS value. If prompting is forbidden (the OPI operand was specified), the default value A is assigned.

For the other log parameter, see LOGLMT.

Value Range: A single alphabetic or numeric character: A-Z or 0-9.

Default Value: A, which represents output class A.

Associated Parmlib Member: None

LOGLMT = nnnnnn

Meaning and Use: This parameter specifies the maximum number of WTLs (messages) allowed for each log data set. The value is used by log processing to determine when a log data set should be scheduled for sysout processing by JES. When the value is reached, log processing issues a simulated WRITELOG command to close and free the current log data set, and to allocate and open a new log data set.

Example: LOGLMT=004852

In this example, when 4,852 WTLs have been issued to a log data set, the data set is scheduled for sysout processing on the output class specified by the

LOGCLS parameter. Log processing then allocates and opens a new log data set.

If the specified value is invalid or an I/O error occurs while the IEASYSxx member is being read, master scheduler initialization prompts the operator for a replacement LOGLMT value. If prompting is forbidden (the OPI operand was specified), the default value of 500 is assigned.

For the other log parameter, see LOGCLS.

Value Range: 000000-999999

Default Value: 500

Associated Parmlib Member: None

LPA = $\left\{ \begin{array}{l} \text{aa} \\ (\text{aa,bb,...[,L]}) \end{array} \right\}$

Meaning and Use: This parameter specifies one or more LPALSTxx parmliib members. The two alphameric characters, represented by aa (or bb, etc.), are appended to LPALST to form the name of the LPALSTxx member(s). If the L option is specified, the system displays (at the operator's console) the names of the data sets successfully concatenated to SYS1.LPALIB.

The LPALSTxx member(s) list data sets that are to be concatenated to SYS1.LPALIB. (For information on the use, contents, and syntax of LPALSTxx, see the description of member LPALSTxx.)

Value Range: Any two alphameric characters.

Default Value: None.

Associated Parmlib Member: LPALSTxx

MAXCAD = nnn

Meaning and Use: Reserves the number of entries available for SCOPE=COMMON data spaces on all primary address space access lists (PASN-ALs) in the system.

Reference

For a description of data spaces, see *Extended Addressability*

A SCOPE=COMMON data space can be used by all programs in the system. It provides a commonly addressable area similar to the common storage area (CSA). SCOPE=COMMON data spaces are used by MVS and can also be used by subsystems and applications that have a need for common storage.

Each SCOPE=COMMON data space uses one entry on all PASN-ALs in the system. Because the maximum number of entries in a PASN-AL is 254, each SCOPE=COMMON data space your program creates and adds to the PASN-AL, reduces the number of SCOPE=SINGLE and SCOPE=ALL data spaces that a program can address through its PASN-AL. Therefore, it is recommended that installations allow each subsystem or application only one SCOPE=COMMON data space.

When selecting a value for MAXCAD (that is selecting a maximum number of SCOPE=COMMON data spaces), allow for the SCOPE=COMMON data spaces that MVS uses, then allow for the SCOPE=COMMON data spaces that subsystems or applications use.

If you code an invalid number (less than 10 or greater than 125) the system uses the default number of 25 and issues an informational message.

Value Range: 10-125

Default Value: 25

Associated Parmlib Member: None

MAXUSER = nnnnnn

Meaning and Use: This parameter specifies a value that, under most conditions, the system uses to limit the number of jobs and started tasks that can execute concurrently during a given IPL. (The number includes time sharing jobs, batch jobs, started system tasks, the master scheduler, and JES2 or JES3.) However, when the system is heavily used, it can use the value specified on the RSVSTRT system parameter to allow more concurrent jobs and started tasks than the number specified by MAXUSER.

Note: The system uses the values specified for the MAXUSER, RSVSTRT, and RSVNONR system parameters to determine the number of entries in the address space vector table (ASVT), which is used to locate the various address space control blocks.

Assume, for example, that MAXUSER specifies 500 and RSVSTRT specifies 5. If there is an attempt to start an address space (using the START command), and none of the 500 ASVT entries defined by the MAXUSER parameter is available (meaning heavy system use), but an entry defined by the RSVSTRT parameter is available, the system uses that entry. Thus, when the system is heavily used, there can be more concurrent jobs and started tasks in the system than the number defined by MAXUSER.

If the system's recovery has to reconstruct the ASVT because the ASVT was written over in error, some entries reserved for replacement (via RSVNONR) might be added to the normal queue of available entries. Therefore, the absolute limit to the number of concurrent jobs and started tasks is the sum of

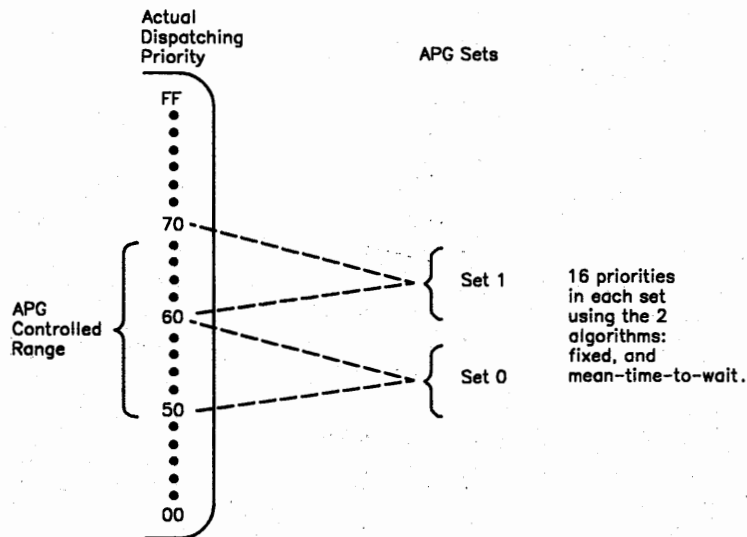


Figure 5-6. Extended Dispatching Control

The time slice function allows a job to have two different priorities within this range (DP-base priority and TSDP-time slice priority.) The amount of time the job will be at each priority is set by the installation via the TSGRP, TSPTRN, and TUNIT parameters. The TSGRP parameter defines a time slice group. Each job is associated with a time slice group by specifying this parameter within the related performance group period. The TSPTRN defines a set number of intervals (1-64) and determines which address spaces will be raised to their time slice priority during those intervals. In this way the user may control the relationship of resource utilization between jobs. For example, if two time slice groups are defined TSGRP=1 and TSGRP=2, and it is required that one group should be favored 75% of the time while the other is favored for 25% of the time, the TSPTRN parameter would be: TSPTRN=(1,1,1,2). The actual time duration represented by each interval within the TSPTRN parameter is set by the TUNIT parameter. Figure 5-7 illustrates the SRM timing parameters that are discussed in this section.

The following table relates the SRM seconds found in Figure 5-7, to wall clock time as follows:

Processor Model	SRM sec/real time sec.
3090 Model 100S	6.7114
3090 Model 120E	8.6957
3090 Model 120S	9.2593
3090 Model 150E	11.2360
3090 Model 150S	14.7059
3090 Model 170S	18.1818
3090 Model 180E	19.6078
3090 Model 180S	25.0000
3090 Model 200E	19.6078
3090 Model 200S	25.0000
3090 Model 250S	14.7059
3090 Model 280E	19.6078
3090 Model 280S	25.0000
3090 Model 300E	19.6078
3090 Model 300S	25.0000
3090 Model 380S	25.0000
3090 Model 400E	19.6078
3090 Model 400S	25.0000
3090 Model 500E	19.6078
3090 Model 500S	25.0000
3090 Model 600E	19.6078
3090 Model 600S	25.0000
4381 Model 90E	5.9524
4381 Model 91E	7.3529
4381 Model 92E	7.3529

Finally, PVLDP sets the priority of jobs in the performance group zero. Any job which is marked privileged in the program properties table will be placed in performance group zero.

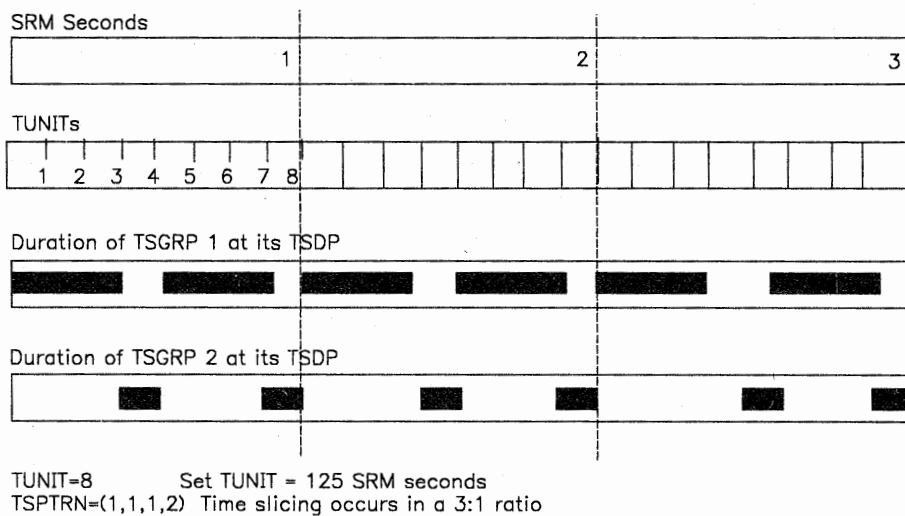


Figure 5-7. Example of SRM Timing Parameters in the IPS

Selecting Service Definition Coefficients (Worksheet Number 1)

The purpose of the service definition coefficients (SDCs) is to allow the individual service components to be weighted. For example, there will probably be a greater demand for the resource in least supply. Using coefficients, service provided by such a resource can be given added importance. This is not to imply that the service values are to be used for accounting purposes, since, as will be seen later, the service consumed by a job is not necessarily repeatable.

The coefficients should be high enough to yield a range of service rates sufficiently high for the workload management function to be effective, but they must not be so high that service rates become excessively large. This also results in ineffective workload management control.

Changes to the coefficients may require changes to other parameters in the IPS that are dependent on service value specifications (for example, interval service value and duration).

An increase in a service definition coefficient will numerically raise the system service capacity and the service rate of users, though not, of course, affecting the system's physical capacity for work.

For example, consider the I/O component of service for the following:

- With $IOC = 1.0$, a service rate of 100 represents 100 I/O requests/second.
- With $IOC = 2.0$, a service rate of 100 represents 50 I/O requests/second.

Three alternatives exist for setting the SDCs:

1. Use the values in the default IPS (IEAIPS00). These should generate reasonable service values for the workload management function. The default coefficients are:
 $CPU = 10.0$
 $IOC = 5.0$
 $MSO = 3.0$
 $SRB = 10.0$
2. Use previous installation SDCs. The installation may wish to include the coefficient for SRB service for several reasons:
 - To obtain information through RMF on SRB usage
 - To monitor SRB usage
3. Define new coefficients. To do this, a more detailed understanding of the individual service components is required. These are discussed separately in the following paragraphs.

CPU and SRB Service: SRM calculates CPU and SRB service based on the task and SRB execution time, the CPU mode, and the number of processors online, which should make the IPS independent of the processor complex.

The following table describes the service consumed per second of execution time by CPU model. The values listed are internal SRM constants. The "total system absorption rate" reported by RMF will not equal the values listed here because these do not include certain types of system processing.

Processor Model	Service Units Per Second of Task or SRB Execution Time	Seconds of Task or SRB Execution Time Per Service Unit
3090 Model 100S	281.88	0.003584
3090 Model 120E	365.22	0.002738
3090 Model 120S	388.89	0.002571
3090 Model 150E	471.91	0.002119
3090 Model 150S	617.65	0.001619
3090 Model 170S	763.64	0.001310
3090 Model 180E	823.53	0.001214
3090 Model 180S	1050.00	0.000952
3090 Model 200E	765.88	0.001306
3090 Model 200S	997.50	0.001033
3090 Model 250S 1 way partitioned	586.76	0.001704
3090 Model 280E partitioned	823.53	0.001214
3090 Model 280S partitioned	1050.00	0.000952
3090 Model 280E single image	765.88	0.001306
3090 Model 280S single image	997.50	0.001003
3090 Model 300E	724.71	0.001380
3090 Model 300S	966.00	0.001035
3090 Model 380S single image	966.00	0.001035
3090 Model 380S 1 way partitioned	1050.00	0.000952
3090 Model 380S 2 way partitioned	997.50	0.001003
3090 Model 400E partitioned	765.88	0.001306
3090 Model 400S partitioned 2 way	997.50	0.001003
3090 Model 400E single image	700.00	0.001429
3090 Model 400S single image	924.00	0.001082
3090 Model 500E partitioned 2 way	765.88	0.001306
3090 Model 500S partitioned 2 way	997.50	0.001003
3090 Model 500E partitioned 3 way	724.71	0.001380
3090 Model 500S partitioned 3 way	966.00	0.001035
3090 Model 500E single image	658.82	0.001518
3090 Model 500S single image	892.50	0.001120
3090 Model 600E partitioned	724.71	0.001380
3090 Model 600S partitioned 3 way	966.00	0.001035
3090 Model 600E single image	625.88	0.001598
3090 Model 600S single image	850.50	0.001176
4381 Model 90E	250.00	0.004000
4381 Model 91E	308.82	0.003238
4381 Model 92E	262.50	0.003810

For installations with no prior service data, the task time reported in type 4, 5, 30, 34, and 35 SMF records could be converted to service units using the above table.

CPU time may not be identical for different runs of the same job step. One or more of the following factors may cause small variations in CPU time: CPU architecture (such as storage buffering), cycle stealing with integrated channels, and the amount of queue searching (see the publication *SPL: System Management Facilities*). An installation may find that SRB time is approximately 1/10 of task time.

I/O Service: SRM calculates I/O service using either I/O block (EXCP) counts or device connect time (DCTI), as specified on the IOSRVC keyword in the IEAIPSxx parmlib member. The calculations are (1) I/O service units equal I/O block count, or (2) I/O service units equal DCTI (in seconds) divided by 8.3 milliseconds. If DCTI is used to calculate I/O service, operations to VIO data sets and to devices that the channel measurement facility does not time are not included in the I/O service total.

When an address space executes in cross memory mode (that is, during either secondary addressing mode or a cross memory call), EXCP counts or the DCTI are not included in the I/O service total. This I/O service is *not* counted for the address

Parameter Descriptions

Keyword	Meaning	Value Range	Default Value
CCCUTT	<p>Specifies the low (CCCUTLOT) and high (CCCUTHIT) processor utilization threshold values used in CPU load balancing to determine the degree that the processor is over or underutilized.</p> <p>The syntax is: CCCUTT = (a,b)</p>	<p>a = 0.0-128.0 b = 0.0-128.0 percent</p>	<p>85 100</p>
CCCSIGUR	<p>Specifies the minimum mean-time-to-wait threshold value for heavy CPU users. Users exceeding this threshold may be considered for CPU load balancing. This constant is also used to determine the range of mean-time-to-wait values which are assigned to each of the ten mean-time-to-wait dispatching priorities.</p> <p>The specified real time value is adjusted by relative processor speed to become SRM time in order to insure consistent SRM control across various processors.</p>	<p>0-32767 milliseconds</p>	<p>45</p>
CNTCLIST	<p>Specifies if the individual commands in a TSO CLIST are treated as separate commands for transaction control. To use this function, TSO Extensions must be installed. The syntax is: CNTCLIST = option</p> <p>where option is either YES or NO. CNTCLIST = NO specifies that the CLIST is treated as a single transaction. CNTCLIST = YES specifies that each command is to be treated as an individual transaction. By specifying CNTCLIST = YES, SRM control of a TSO command becomes the same whether the command is executed explicitly or as part of a CLIST. The RTO parameter, however, does not affect commands within a CLIST, even if they are treated as individual transactions.</p>	<p>NO YES</p>	<p>NO</p>
CNTNSW	<p>Specifies if non-swappable address spaces are to be included in the current MPL of their domain.</p> <p>The syntax is: CNTNSW = option</p> <p>where option is either YES or NO. CNTNSW = NO specifies that they are not to be counted. CNTNSW = YES specifies that they are to be counted in the current MPL.</p> <p>When CNTNSW = YES is specified, the minimum MPL for each domain should allow for the permanently non-swappable address spaces in the domain. The system component address spaces, such as PC/AUTH and ALLOCAS, are non-swappable and are usually in the default domain for started tasks; count these address spaces when setting the minimum for the domain.</p>	<p>NO YES</p>	<p>NO</p>
CPENABLE	<p>Specifies the low (ICCTPILO) and high (ICCTPIHI) threshold values for the percentage of I/O interruptions to be processed through the test pending interrupt (TPI) instruction path in IOS.</p> <p>SRM uses these thresholds to control the number of processors enabled for I/O interruptions. The syntax is: CPENABLE = (a,b)</p>	<p>a = 0-100 b = 0-100 percentage</p>	<p>10 30</p>

Keyword	Meaning	Value Range	Default Value
CPU	Specifies the weighting factor by which the recommendation of the CPU load balancing factor is to be multiplied when an address space is evaluated for swapping. Example: CPU = 3.0	0.0-9.9	0.1
DVIO	Specifies whether directed VIO is to be active in the system or not. The syntax is: DVIO = option where option is either YES or NO. DVIO = YES, the default, specifies that directed VIO is to be active in the system; that is, the NONVIO keyword of the IEASYSxx parmlib member is honored. DVIO = NO specifies that directed VIO is not to be active in the system; the NONVIO parameter of the IEASYSxx parmlib member is ignored.	YES NO	YES
ERV	Specifies the number of CPU service units that the address space is allowed to absorb before being considered for a swap out based on a recommendation value analysis. An address space should not be swapped out during this period if one of the follow conditions is met: <ul style="list-style-type: none"> • The address space is enqueued on a system resource needed by another address space. • An authorized program in the address space obtains control of the resource (even if another address space does not need that resource) as a result of issuing a reserve for a DASD device which is SHARED. Example: ERV = 2 Note: SRM determines the execution time equivalent to the specified ERV by multiplying the ERV by the model-dependent time needed to accumulate 1 CPU service unit. In the example given, if an address space can consume 1 service unit in 10 milliseconds, it will be allowed to execute for 20 milliseconds, when it is enqueued on a resource requested by other address spaces, before it will be eligible for swap-out.	0-999999	500
ESCTBDS	Specifies the criteria age (in seconds) at which a hiperspace page will be send to extended storage. Example: ESCTBDS = 500 The hiperspace that exceeds this time is added to the processors reserve capacity. The formula is: $ma > ca$	0-32767	900
ESCTPOC(n)	Specifies the criteria age (in seconds) at which a <i>changed</i> page that is to be <i>paged out</i> will be sent to extended storage. Specify a value for n to indicate the type of user owning the page, where n can be 0 - nonswappable, common, or privileged users, 1 - all others not 0 or 2, or 2 - TSO users. Example: ESCTPOC(0) = 100 Changed paged-out pages are sent to extended storage when the migration age (ma) is greater than the criteria age (ca). The formula is: $ma > ca$	0-32767	0 = 100 1 = 100 2 = 100

Keyword	Meaning	Value Range	Default Value
ESCTPOU(n)	<p>Specifies the criteria age (in seconds) at which an <i>unchanged</i> page that is to be <i>paged out</i> will be sent to extended storage. Specify a value for n to indicate the type of user owning the page, where n can be</p> <p>0 - nonswappable, common, or privileged users, 1 - all others not 0 or 2, or 2 - TSO users.</p> <p>Example: ESCTPOU(1)=100</p> <p>Unchanged paged-out pages are sent to extended storage when the migration age (ma) is greater than the criteria age (ca). The formula is:</p> <p>ma > ca</p>	0-32767	0=100 1=100 2=100
ESCTSTC(n)	<p>Specifies the criteria age (in seconds) at which a <i>changed</i> page that is to be <i>stolen</i> will be sent to extended storage. Specify a value for n to indicate the type of user owning the page, where n can be</p> <p>0 - nonswappable, common, or privileged users, 1 - all others not 0 or 2, or 2 - TSO users.</p> <p>Example: ESCTSTC(2)=15</p> <p>Changed stolen pages are sent to extended storage when the migration age (ma) is greater than the criteria age (ca). The formula is:</p> <p>ma > ca</p>	0-32767	0=0 1=20 2=15
ESCTSTU(n)	<p>Specifies the criteria age (in seconds) at which an <i>unchanged</i> page that is to be <i>stolen</i> will be sent to extended storage. Specify a value for n to indicate the type of user owning the page, where n can be</p> <p>0 - nonswappable, common, or privileged users, 1 - all others not 0 or 2, or 2 - TSO users.</p> <p>Example: ESCTSTU(0)=0</p> <p>Unchanged stolen pages are sent to extended storage when the migration age (ma) is greater than the criteria age (ca). The formula is:</p> <p>ma > ca</p>	0-32767	0=0 1=20 2=15
ESCTSWTC(n)	<p>Specifies the criteria age (in seconds) at which a <i>changed</i> page that has been trimmed for a <i>swap out</i> will be sent to extended storage. Specify a value for n to indicate the type of user owning the page, where n can be</p> <p>0 - privileged users, 1 - all others not 0 or 2, or 2 - terminal wait swap users.</p> <p>Example: ESCTSWTC(1)=100</p> <p>Changed swap-out pages are sent to extended storage when the sum of the system-high unreferenced interval count (uic) and the migration age (ma) is greater than the criteria age (ca). The formula is:</p> <p>mc + ma > ca</p>	0-32767	0=100 1=100 2=60

Keyword	Meaning	Value Range	Default Value
ESCTSWTU(n)	<p>Specifies the criteria age (in seconds) at which an <i>unchanged</i> page that has been trimmed for a <i>swap out</i> will be sent to extended storage. Specify a value for n to indicate the type of user owning the page, where n can be</p> <p>0 - privileged users, 1 - all others not 0 or 2, or 2 - terminal wait swap users.</p> <p>Example: ESCTSWTU(2)=60</p> <p>Unchanged swap-out pages are sent to extended storage when the sum of the system-high unreferenced interval count (uic) and the migration age (ma) is greater than the criteria age (ca). The formula is:</p> $\text{uic} + \text{ma} > \text{ca}$	0-32767	0 = 100 1 = 100 2 = 60
ESCTSWWS(n)	<p>Specifies the criteria age (in seconds) at which a <i>working set</i> page that is ready for a <i>swap out</i> will be sent to extended storage. Specify a value for n to indicate the type of user owning the page, where n can be</p> <p>0 - privileged users, 1 - all others not 0 or 2, or 2 - terminal wait swap users.</p> <p>Example: ESCTSWWS(0)=100</p> <p>Working set swap-out pages are sent to extended storage when the sum of the system-high unreferenced interval count (uic) and the migration age (ma) is greater than the criteria age (ca). The formula is:</p> $\text{uic} + \text{ma} > \text{ca}$ <p>Terminal wait pages are sent to extended storage when the sum of the system-high unreferenced interval count (uic) and the migration age (ma) minus the think time (tt) is greater than the criteria age (ca). The formula is:</p> $(\text{uic} + \text{ma}) - \text{tt} > \text{ca}$	0-32767	0 = 100 1 = 100 2 = 50
ESCTVF	<p>Specifies the criteria age (in seconds) at which a <i>virtual fetch</i> page will be sent to extended storage.</p> <p>Example: ESCTVF = 15</p> <p>Virtual fetch pages are sent to extended storage when the migration age (ma) is greater than the criteria age (ca). The formula is:</p> $\text{ma} > \text{ca}$	0-32767	15
ESCTVIO	<p>Specifies the criteria age (in seconds) at which a <i>virtual I/O</i> page will be sent to extended storage.</p> <p>Example: ESCTVIO = 900</p> <p>Virtual I/O pages are sent to extended storage when the migration age (ma) is greater than the criteria age (ca). The formula is:</p> $\text{ma} > \text{ca}$	0-32767	900
ICCLPB (TAPE)	<p>Specifies the low (ICCLPBLO(1)) and high (ICCLPBHI(1)) logical path utilization threshold values for paths used by tape devices. The syntax is:</p> $\text{ICCLPB}(\text{TAPE}) = (\text{a}, \text{b})$	a = 0-100 b = 0-100 percent	50 80



.

4



.

.





.

2




.

.









1



2





1



2





GN28-1311-00

