

Adabas SMF Records

All Adabas SMF records have a common structure, with some sections appearing in all records and others generated according to specific events and parameters specified through ADARUN or operator commands. The ASMFREC macro provides mapping DSECTs for all parts of the SMF record.

This chapter covers the following topics:

- Record Structure
 - Record Size Limits
 - Record Subtypes
 - Statistical Recording
 - Record Sections
 - ASMFREC Mapping Macro
 - SMF User Exit
 - IBM Type 89 SMF Records
-

Record Structure

Adabas follows the modern convention for SMF record formats. A single record has:

- A standard IBM-type header
- A self-defining section that describes a variable number of detail sections
- A product ID detail section
- User-selected detail sections

Each detail section is described by an eight-byte entry in the self-defining section containing three fields. An entry is also called a *triplet*.

- A 4-byte offset from the beginning of the record to the detail section
- A 2-byte count of the number of instances of the detail section
- A 2-byte length of each detail section instance. If there are no detail section instances of a given type, the triplet is all zeros.

Field ASNumD in the product ID section specifies the number of triplets in the self-defining section.

Record Size Limits

SMF records are z/OS V-format records with a system-imposed maximum length of 32,756 bytes. Most Adabas SMF records fit within this limit for most reasonable types of ADARUN nucleus specifications. However, detail sections such as File Activity could potentially have several thousand detail section instances.

If the entire set of instances will not fit in the space remaining in the record, Adabas will include only as many as there are room for and write the record. The SMF record is reset by clearing the triplets for all detail sections except the product ID section and then adding as many of the remaining instances as will fit, repeating until all detail sections are processed. Field `ASSegNo` in the product ID section will start at 1 and be incremented with each additional record, and field `ASSegL` will be set to zero for the last (or only) record for an interval or event.

If any detail section is so large that even one instance would cause the record size limit to be exceeded after resetting the SMF record, that detail type is deactivated.

Record Subtypes

The header section field `ASSTy` identifies a record subtype.

- Subtype 1 (ASStI) - Adabas Nucleus Initialization
- Subtype 2 (ASStT) - Adabas Nucleus Termination
- Subtype 3 (ASStI) - Adabas Interval Statistics
- Subtype 4 (ASStP) - Adabas Parameter Change

Subtype 1 (ASStI) - Adabas Nucleus Initialization

A record of this subtype is generated during nucleus initialization. In addition to the header, self-defining and product ID sections, it contains ADARUN parameter and user sections if these have been selected by `SMFDETAIL ADARUN` parameter or operator commands.

Subtype 2 (ASStT) - Adabas Nucleus Termination

A record of this subtype is generated during nucleus termination. In addition to the header, self-defining and product ID sections, it contains all detail sections specified by the `SMFDETAIL ADARUN` parameter or operator commands except for the ADARUN parameter section. Statistics in Adabas detail sections reflect totals for the entire nucleus session.

Subtype 3 (ASStI) - Adabas Interval Statistics

If interval recording has been specified by the `SMFINTERVAL ADARUN` parameter or operator command, a record of this subtype is generated at the expiration of each interval. In addition to the header, self-defining and product ID sections, it contains all detail sections specified by the `SMFDETAIL ADARUN` parameter or operator commands except for the ADARUN parameter section. Statistics in Adabas detail sections reflect activity since the previous interval ended, except where noted. This is also called a delta value.

Subtype 4 (ASStP) - Adabas Parameter Change

If the Adabas ADARUN parameter detail section has been specified by the SMFDETAIL ADARUN parameter or operator commands, a record of this subtype is generated whenever an ADARUN parameter value is changed after nucleus initialization. In addition to the header, self-defining, product ID and ADARUN parameter sections, it may also contain a user section if that has been selected by SMFDETAIL ADARUN parameter or operator commands.

Statistical Recording

The nucleus accumulates usage statistics on the resources it uses to accomplish its tasks. These statistics may be recorded at user or system-defined intervals (see ADARUN parameter SMFINTERVAL) and at termination.

Interval recording (Adabas SMF record subtype 3) provides the usage since the last interval ended for each detail section. Adabas SMF record intervals may be synchronized with one of the system-level intervals specified by PARMLIB member SMFPRMxx entries. This allows straightforward analysis of the usage by allowing direct comparison with other record interval data. For example, you can compare the Adabas interval record with RMF data for the same interval to better understand system performance.

Statistics at termination (Adabas SMF record subtype 2) will have cumulative statistics that reflect activity for the entire nucleus session in each specified detail section.

Record Sections

Every Adabas SMF record contains header, self-defining and product ID sections. You can select additional detail sections through the SMFDETAIL ADARUN parameter or operator commands. Each section is mapped by a DSECT generated by the ASMFREC mapping macro.

The following table summarizes the Adabas SMF record sections:

Detail Section Description	ASMFREC Macro or ADARUN Parameter Specification	Self-Defining Section Triplet Label Base	ASMFREC DSECT Name Produced by the ASMFREC Macro
Header and self-defining section	---	---	ASBase
Adabas command activity	CMD	ASTCmd	ASCmd
Adabas global cache activity by block type ¹	CSHB	ASTChB	ASChB
Adabas global cache activity by Adabas file number ¹	CSHF	ASTChF	ASChF
Adabas global cache activity ¹	CSHG	ASTChG	ASChG
Adabas Parallel Services cache activity ²	CSHP	ASTChP	ASChP
Adabas file activity	FILE	ASTFile	ASFile
Adabas global lock activity ¹	LOCK	ASTLok	ASLok
Adabas internucleus messaging control block activity	MSGB	ASTMsgB	ASMsgB
Adabas internucleus messaging counts	MSGC	ASTMsgC	ASMsgC
Adabas internucleus messaging service time histogram	MSGH	ASTMsgH	ASMsgH
ADARUN parameter values	PARM	ASTParm	ASParm
I/O by DD name	IODD	ASTIODD	ASIODD
Product ID ³	ID	ASTPID	ASPID
Storage pool	STG	ASTStg	ASStg
Thread activity	THRD	ASTThrd	ASThrd
User-defined	USER	ASTUsr	user-defined

1. The detail section is available only in cluster environments when either Adabas Cluster Services or Adabas Parallel Services are installed.
2. The detail section is available only in cluster environments when Adabas Parallel Services is installed.

- The product ID section is always included in every SMF record. It may not be specified in the SMFDETAIL ADARUN parameter or in operator commands.

This section describes the different detail record sections:

Note:

The DSECTs provided in the following sections may not be the most current. To see the most current versions of the DSECTs, generate them using the ASMFREC macro.

- Header Section
- Self-Defining Section
- Product ID Section: ID
- Adabas Command Activity Section: CMD
- Adabas File Activity Section: FILE
- Adabas Global Cache Activity by Block Type Section: CSHB
- Adabas Global Cache Activity by Adabas File Number Section: CSHF
- Adabas Global Cache Activity Section: CSHG
- Adabas Global Lock Activity Section: LOCK
- Adabas Internucleus Messaging Control Block Activity Section: MSGB
- Adabas Internucleus Messaging Counts Section: MSGC
- Adabas Internucleus Messaging Service Time Histogram Section: MSGH
- Adabas Parallel Services Cache Activity Section: CSHP
- ADARUN Parameter Value Section: PARM
- I/O by DD Name Section: IODD
- Storage Pool Section: STG
- Thread Activity Section: THRD

Header Section

IBM has defined a standard format for the initial part of all SMF records in *z/OS MVS System Management Facilities (SMF)*, IBM document SA22-7630. This section begins every Adabas SMF record.

ASBase	D	Sect	,	Base segment
*				
*				Standard SMF Header
*				
ASRDW	DS	0B14		Record descriptor word
ASLen	DS	B12		Record length
ASSeg	DS	B12		Segment descriptor
ASFlg	DS	B1.8		System indicator flags

ASFStV	Equ	x'40'	Subtypes are valid
ASFV4	Equ	x'10'	MVS/SP V4 and above
ASFV3	Equ	x'08'	MVS/SP V3 and above
ASFV2	Equ	x'04'	MVS/SP V2 and above
ASFVS2	Equ	x'02'	VS2
ASRTy	DS	B11	Adabas record type
ASTme	DS	B14	Time since midnight when record was + moved into SMF buffer in 1/100 sec
ASDte	DS	P14	Date when record was moved into SMF + buffer as 0cyydddF
ASSID	DS	C14	System identifier (SMFPRMxx SID)
ASSSI	DS	C14	Subsystem identifier
ASSty	DS	B12	Subtype
ASStI	Equ	1	Adabas initialization
ASStT	Equ	2	Adabas termination
ASStS	Equ	3	Interval statistics
ASStP	Equ	4	Parameter change
ASStA	Equ	9	Ad hoc record
*			
ASBaseL	Equ	*-ASBase	Length of standard header

Self-Defining Section

The self-defining section follows immediately after the header section. It is part of the header section DSECT.

Each detail section triplet is identified by a base label as shown in the table at the beginning of this section. The base label begins with the prefix specified in the ASMFREC invocation followed by the letter T (for triplet), and then followed by a mnemonic detail section identifier. The base label with suffix O is the offset, with suffix L is the length, and with suffix N is the number of instances.

Here is an example of some triplets.

```

*
*                               Self-Defining Section
*
ASSDS    DS    0B                Self-defining section
*                               Map of typical section triplet
ASSDSO   DS    B14               Offset to section from start of +
                               record
ASSDSL   DS    B12               Length of section
ASSDSN   DS    B12               Number of section(s)
                               Org    ASSDS
*
ASTID    DS    0B18              ID Section (always present)
ASTIDO   DS    B14               Offset to ID section from start +
                               of record
ASTIDL   DS    B12               Length of ID section
ASTIDN   DS    B12               Number of ID section(s)
*
ASTUser  DS    0B18              User-Defined Section
ASTUserO DS    B14               Offset to User-Defined section from+
                               start of record
ASTUserL DS    B12               Length of User-Defined section
ASTUserN DS    B12               Number of User-Defined section(s)
*
ASTParm  DS    0B18              ADARUN Parameter Section
ASTParmO DS    B14               Offset to detail section from start+
                               of record
ASTParmL DS    B12               Length of each detail section

```

```

ASTParmN DS      B12          Number of detail section(s)
*
*
*
ASSDSLn  Equ     *-ASSDS      Length of self-defining section
ASSDSNT  Equ     ASSDSLn/8    Number of triplets

```

Product ID Section: ID

The product ID section is always present in every Adabas SMF record with one instance. It describes the nucleus generating the SMF record and provides information about the record's contents.

The 2-byte version code consists of a major version and a minor version. A change, such as adding a new triplet or extending a detail section, will increment the minor version. All existing programs should continue to operate as no existing displacements have changed. A more disruptive change will increment the major version and require existing programs to (at least) be reassembled.

```

ASPID    DSect ,          Product ID Detail Section          +
          (always present in SMF record)
ASSMFV   DS      0B12    SMF record version
ASSMFVM  DS      B11    SMF record major version
ASSMFVN  DS      B11    SMF record minor version
ASSMFVC  Equ     ASSMFV11 Current version: 1.1
ASSMFV11 Equ     x'0101'  Version 1.1 - Initial release
ASSegNo  DS      B11    Record segment number
ASSegL   DS      B11    Last segment when = 0
ASNumD   DS      B12    Number of detail type triplets
ASPNm    DS      C18    Product name (ADABAS)
ASVRSC   DS      C18    Product ver/rlse/SM/cum: vvrsscc
ASSysN   DS      C18    System name
ASSypN   DS      C18    Sysplex name
ASVMN    DS      C18    Virtual machine name
ASJbN    DS      C18    Job name
ASStN    DS      C116   ProcStep/Step name
ASJNm    DS      C18    JES job identifier
ASPgm    DS      C18    Program name
ASGrp    DS      C18    Cluster messaging group name
ASST     DS      B18    Nucleus start time in STCK format
ASIST    DS      B18    Interval start time in STCK format
ASiET    DS      B18    Interval end time in STCK format
ASDBID   DS      B14    Database ID
ASNucX   DS      B12    External nucleus ID
ASNucI   DS      B11    Internal nucleus ID
ASSVC    DS      B11    Adabas SVC number
ASASID   DS      B12    Address space ID
ASASIDI  DS      B14    Reusable address space ID instance
ASComp   DS      B12    Completion code          +
          x'0ccc' System ABEND code ccc          +
          x'8ccc' User ABEND code ccc
ASARC    DS      B14    ABEND reason code
*
ASPIDL   Equ     *-ASPID

```

Adabas Command Activity Section: CMD

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Adabas command activity data is derived from data presented at nucleus shutdown.

There is one instance for each command group: A1/4, BT, CL, ET, E1/4, L1/4, L2/5, L3/6, L9, LF, N1/2, OP, UC, RC, RE, REST, S1/4, S2, S5, S8, S9, YA, YB, YF, YP, YCAL, V1, V2, V3, V4, U0, U1, U2 and U3. There are 34 possible instances but this is subject to change in future releases.

```

ASCmd      DSect ,           Adabas Command Activity
ASCmdNm    DS      C14      Command name
ASCmdCt    DS      B18      Number of times this command type +
                             was executed
ASCmdTm    DS      B18      Sum of this command type durations +
                             in microseconds
*
ASCmdL     Equ      *-ASCmd

```

Adabas File Activity Section: FILE

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Adabas file activity data is derived from data presented at nucleus shutdown or in response to a DFILUSE operator command. There is one instance for each file possible in the database as specified by ADADEF MAXFILES up to the highest file number with a non-zero use count. The file number is implied by the sequence number of the instance, starting with zero, which reflects commands such as OP that are not associated with a specific file.

```

ASFile     DSect ,           Adabas File Activity
ASFileCt   DS      B18      Number of commands executed against +
                             this file
*
ASFileL    Equ      *-ASFile

```

Adabas Global Cache Activity by Block Type Section: CSHB

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Global cache statistics are available only for Adabas Cluster Services and Adabas Parallel Services nuclei. They are derived from the ones presented at nucleus shutdown or in response to a DXCACHE operator command. There is one detail section instance for each type of block. Users should examine the block type and not rely on any observed order of the instances. The following block types are reported:

- AC: Address Converter
- DS: Data Storage
- DSST: Data Storage Space Table
- FCB: File Control Block
- NI: Normal Index
- UI: Upper Index
- OTHR: Any other block type

```

ASChB      DSect ,           Global Cache Activity by Block
ASBCN      DS      B12      Cache Number
ASCBRSv1   DS      B12      Unused
ASCBRT     DS      C14      Block type
ASCBRT     DS      B18      Reads - Total
ASCBRCs    DS      B18      Reads - Completed synchronous

```


ASCBRCA	DS	B18	Reads - Completed asynchronous
ASCBRIC	DS	B18	Reads - Data in cache
ASCBRNI	DS	B18	Reads - Data not in cache
ASCBRFS	DS	B18	Reads - Failed - Structure
ASCBRO	DS	B18	Reads - For cast-out
ASCBROS	DS	B18	Reads - For cast-out synchronous
ASCBROA	DS	B18	Reads - For cast-out asynchronous
ASCBWT	DS	B18	Writes - Total
ASCBWCS	DS	B18	Writes - Completed synchronous
ASCBWCA	DS	B18	Writes - Completed asynchronous
ASCBWDR	DS	B18	Writes - Data written
ASCBWNR	DS	B18	Writes - Data not written
ASCBWSF	DS	B18	Writes - Structure full
ASCBVI	DS	B18	Validates issued
ASCBVF	DS	B18	Validates failed
ASCBBD	DS	B18	Block deletes issued
ASCBDR	DS	B18	Deletes reissued due to timeout
ASCBUR	DS	B18	Number of times updates redone
*			
ASChBL	Equ	*-ASChB	

Adabas Global Cache Activity by Adabas File Number Section: CSHF

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Global cache statistics are available only for Adabas Cluster Services and Adabas Parallel Services nuclei. They are derived from the ones presented at nucleus shutdown or in response to a DXCACHE operator command. There is potentially one instance for each file possible in the database as specified by ADADEF MAXFILES. The size of this detail section precludes the ability to generate it for every possible file, so there is one detail section instance for each file that has non-zero usage. Users should examine the file number and not rely on any observed order of the instances.

ASChF	DSect	,	Global Cache Activity by File
ASFCn	DS	B12	Cache Number
ASCFRsv1	DS	B12	Unused
ASCFNum	DS	B14	File number
ASCFRT	DS	B18	Reads - Total
ASCFRCS	DS	B18	Reads - Completed synchronous
ASCFRCA	DS	B18	Reads - Completed a synchronous
ASCFRIC	DS	B18	Reads - Data in cache
ASCFRNI	DS	B18	Reads - Data not in cache
ASCFRFS	DS	B18	Reads - Failed - Structure
ASCFRO	DS	B18	Reads - For cast-out
ASCFROS	DS	B18	Reads - For cast-out synchronous
ASCFROA	DS	B18	Reads - For cast-out asynchronous
ASCFWT	DS	B18	Writes - Total
ASCFWCS	DS	B18	Writes - Completed synchronous
ASCFWCA	DS	B18	Writes - Completed a synchronous
ASCFWDR	DS	B18	Writes - Data written
ASCFWNR	DS	B18	Writes - Data not written
ASCFWSF	DS	B18	Writes - Structure full
ASCFVI	DS	B18	Validates issued
ASCFVF	DS	B18	Validates failed
ASCFBD	DS	B18	Block deletes issued
ASCFDR	DS	B18	Deletes reissued due to timeout
ASCFUR	DS	B18	Number of times updates redone
*			
ASChFL	Equ	*-ASChF	

Adabas Global Cache Activity Section: CSHG

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Global cache statistics are available only for Adabas Cluster Services and Adabas Parallel Services nuclei. They are derived from the ones presented at nucleus shutdown or in response to a DXCACHE operator command. This detail section appears with one instance.

ASChG	Dsect	,	Global Cache Activity Section
ASCGCN	DS	B12	Cache Number
ASCGRsv1	DS	B12	Unused
ASCGCOD	DS	B18	Cast-out directory reads issued
ASCGCODA	DS	B18	Cast-out directory - async
ASCGCODS	DS	B18	Cast-out directory - sync
ASCGCOU	DS	B18	Unlock cast-out locks issued
ASCGCOUA	DS	B18	Unlock cast-out locks - async
ASCGCOUS	DS	B18	Unlock cast-out locks - sync
ASCGDR	DS	B18	Directory reads issued
ASCGDRA	DS	B18	Directory reads issued - sync
ASCGDRS	DS	B18	Directory reads issued - async
ASCGPub	DS	(0*9)B18	Publishing requests
ASCGSync	DS	B18	Update sync
ASCGXEnd	DS	B18	BT/CL/ET transaction end
ASCGRedo	DS	B18	Redo threshold
ASCGFull	DS	B18	Full buffer pool
ASCGAll	DS	B18	All blocks
ASCGRABN	DS	B18	Specific RABN
ASCGDS	DS	B18	File DS blocks
ASCGDSST	DS	B18	DSST blocks
ASCGNI	DS	B18	File NI blocks
*			
ASChGL	Equ	*-ASChG	

Adabas Global Lock Activity Section: LOCK

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Global lock statistics are available only for Adabas Cluster Services and Adabas Parallel Services nuclei. They are derived from the ones presented at nucleus shutdown or in response to a DXLOCK operator command. There is one detail section instance for each lock type. The lock type is implied by the sequence number of the instance, starting with one.

ASLok	Dsect	,	Global Lock Section
*			Lock Types
ASLokGC	Equ	1	GCB
ASLokSE	Equ	2	Security
ASLokFS	Equ	3	FST
ASLokUF	Equ	4	UFT
ASLokSO	Equ	5	Save Online
ASLokFL	Equ	6	Flush
ASLokES	Equ	7	Global ET Synchronization
ASLokRC	Equ	8	Recovery
ASLokUT	Equ	9	UFT-File
ASLokIU	Equ	10	Index Update
ASLokHI	Equ	11	Hold ISN
ASLokUD	Equ	12	Unique DE
ASLokET	Equ	13	ETID
ASLokLT	Equ	14	LOB Tracker
ASLokCM	Equ	15	Command Manager User
ASLokDI	Equ	16	Data Increment
ASLokCP	Equ	17	Checkpoint

ASLokDT	Equ	18	Net-Work DBID Target Assignment
ASLokGU	Equ	19	Global Update Commmand Sync
ASLokPM	Equ	20	Parameter
ASLokDS	Equ	21	DSF
ASLokRG	Equ	22	RLOG
ASLokSP	Equ	23	SPATS
ASLokCA	Equ	24	Cancel
ASLokWR	Equ	25	TBWK4A/E Table
ASLokWU	Equ	26	PUTUA/E Table
ASLokXI	Equ	27	XIDE
ASLokRH	Equ	28	Replication Handshake
ASLokRI	Equ	29	Read file/ISN
ASLokFA	Equ	30	Format AC/AC1
ASLokOC	DS	B18	Obtains - Conditional
ASLokOG	DS	B18	Obtains - Granted
ASLokOR	DS	B18	Obtains - Rejected
ASLokOU	DS	B18	Obtains - Unconditional
ASLokOS	DS	B18	Obtains - Synchronous
ASLokOA	DS	B18	Obtains - Asynchronous
ASLokAC	DS	B18	Alters - Conditional
ASLokAG	DS	B18	Alters - Granted
ASLokAR	DS	B18	Alters - Rejected
ASLokAU	DS	B18	Alters - Unconditional
ASLokAD	DS	B18	Alters - Deadlock/Rejected
ASLokAS	DS	B18	Alters - Synchronous
ASLokAA	DS	B18	Alters - Asynchronous
ASLokRL	DS	B18	Releases
ASLokRS	DS	B18	Releases - Synchronous
ASLokRA	DS	B18	Releases - Asynchronous
*			
ASLokL	Equ	*-ASLok	

Adabas Internucleus Messaging Control Block Activity Section: MSGB

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Internucleus messaging statistics are available only for Adabas Cluster Services and Adabas Parallel Services nuclei. They are derived from the ones presented at nucleus shutdown or in response to a DXMSG operator command. This detail section appears with one instance. The number of blocks allocated (ASMsgBBA) and the high water mark (ASMsgBBH) reflect the entire nucleus session in interval records.

ASMsgB	Dsect	,	Inter-Nucleus Messaging Counts
ASMsgBBA	DS	B18	Message control blocks allocated
ASMsgBBH	DS	B18	Message control blocks used + (high water mark)
ASMsgBBR	DS	B18	Message control block requests
*			
ASMsgBL	Equ	*-ASMsgB	

Adabas Internucleus Messaging Counts Section: MSGC

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Internucleus messaging statistics are available only for Adabas Cluster Services and Adabas Parallel Services nuclei. They are derived from the ones presented at nucleus shutdown or in response to a DXMSG operator command. This detail section appears with one instance. Adabas Parallel Services nuclei report only the count of messages sent.

```

ASMsgC      DSect ,           Inter-Nucleus Messaging Counts
ASMsgCMT    DS      C14       Message type
ASMsgCMS    DS      B18       Messages sent
ASMsgCMI    DS      B18       Messages incoming (arrived)
ASMsgCMA    DS      B18       Messages accepted
ASMsgCRS    DS      B18       Replies sent
*
ASMsgCL     Equ    *-ASMsgC

```

Adabas Internucleus Messaging Service Time Histogram Section: MSGH

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Internucleus messaging statistics are available only for Adabas Cluster Services and Adabas Parallel Services nuclei. They are derived from the ones presented at nucleus shutdown or in response to a DXMSG operator command. This detail section appears with two instances:

1. The first represents messages subject to the MXMSG timeout parameter.
2. The second represents certain control messages not subject to MXMSG.

The two instances may be summed for a single representation of all messages. All message times are in microseconds. The minimum and maximum durations (ASMsgHMn and ASMsgHMx) reflect the entire nucleus session in interval records. Field ASMsgMD2 is an extended (16-byte) floating point sum of the squares of all message durations. It may be used to compute a standard deviation.

```

ASMsgH      DSect ,           Inter-Nucleus Messaging Histogram
ASMsgHXP    DS      C14       Transport service
ASMsgHMM    DS      B14       MXMSG or zero for messages not      +
                               subject to MXMSG
ASMsgHMC    DS      B18       Message count
ASMsgHMD    DS      B18       Sum of all message durations
ASMsgHMS    DS      B116      Sum of squares, all msg durations    +
                               (extended hex floating point)
ASMsgHMn    DS      B14       Minimum duration (us)
ASMsgHMx    DS      B14       Maximum duration (us)
ASMsgHct    Equ    9         Number of histogram buckets
ASMsgHG     DS      (0*ASMsgHct)B18 Histogram buckets
ASMsgH10    DS      B18       > 1000 s
ASMsgH09    DS      B18       > 100 s, <= 1000 s
ASMsgH08    DS      B18       > 10 s, <= 100 s
ASMsgH07    DS      B18       > 1 s, <= 10 s
ASMsgH06    DS      B18       > 100 ms, <= 1 s
ASMsgH05    DS      B18       > 10 ms, <= 100 ms
ASMsgH04    DS      B18       > 1 ms, <= 10 ms
ASMsgH03    DS      B18       > 100 us, <= 1 ms
ASMsgH02    DS      B18       <= 100 us
*
ASMsgHL     Equ    *-ASMsgH

```

Adabas Parallel Services Cache Activity Section: CSHP

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Parallel services cache statistics are available only for Adabas Parallel Services nuclei. They are derived from the ones presented at nucleus shutdown or in response to a DXCACHE operator command. This detail section appears with one instance. The directory high water mark ASCPDHiN and in-use count ASCPDirI reflect the entire nucleus session in interval records.

```

ASChP      DSect ,           Parallel Services Cache Activity
ASPCPN     DS      B12      Cache Number
ASCPRSv1   DS      B12      Unused
ASCPNDir   DS      B18      Number of directory elements
ASCPNDiI   DS      B18      Number of directory index elements
*
*
ASCPDHiN   DS      B18      High-water mark, this nucleus
ASCPDirI   DS      B18      In-use, this nucleus
*
ASCPDRA    DS      B18      Read
ASCPDRF    DS      B18      Located active
ASCPDRC    DS      (0*4)B18  Obtained from free pool
ASCPDNN    DS      B18      Reclaim criteria categories
ASCPDND    DS      B18      First choice criteria
ASCPDIN    DS      B18      Second choice criteria
ASCPDID    DS      B18      Third choice criteria
ASCPDCF    DS      B18      Fourth choice criteria
ASCPDRT    DS      B18      Unable to obtain (cache full)
ASCPDWF    DS      B18      Tested for reclaim
*
ASCPDWF    DS      B18      Write
ASCPDWF    DS      B18      Obtained from free pool
*
*
ASCPDRF    DS      B18      Space Management Statistics
ASCPDRF    DS      B18      Request Statistics
ASCPDRF    DS      B18      Sufficient preallocated space
ASCPDRF    DS      B18      Free space allocated
ASCPDRF    DS      B18      Reclaim space, first choice
ASCPDRF    DS      B18      Reclaim space, second choice
ASCPDRF    DS      B18      Space unavailable (cache full)
ASCPDRF    DS      B18      Searched part of space chain
ASCPDRF    DS      B18      Searched entire space chain
ASCPDRF    DS      B18      Number of space seqs tested
*
ASCPDRF    DS      B18      Element Reclaim Statistics
ASCPDRF    DS      B18      First choice criteria
ASCPDRF    DS      B18      Second choice criteria
*
*
ASCPDRF    DS      B18      Latch management statistics
ASCPDRF    DS      B18      Cache Space Chain
ASCPDRF    DS      B18      Get Exclusive
ASCPDRF    DS      B18      WaitFor Exclusive
ASCPDRF    DS      B18      Release Exclusive
*
ASCPDRF    DS      B18      Cache Directory Index
ASCPDRF    DS      B18      Get Exclusive
ASCPDRF    DS      B18      Get Shared
ASCPDRF    DS      B18      Upgrade Exclusive
ASCPDRF    DS      B18      WaitFor Exclusive
ASCPDRF    DS      B18      WaitFor Shared
ASCPDRF    DS      B18      WaitFor Upgrade
ASCPDRF    DS      B18      Release Exclusive
ASCPDRF    DS      B18      Release Shared
*
ASCPDRF    DS      B18      Cache Directory
ASCPDRF    DS      B18      Get Exclusive
ASCPDRF    DS      B18      Get Shared
ASCPDRF    DS      B18      Upgrade Exclusive
ASCPDRF    DS      B18      WaitFor Exclusive
ASCPDRF    DS      B18      WaitFor Shared
ASCPDRF    DS      B18      Release Exclusive
ASCPDRF    DS      B18      Release Shared
*
ASCPDRF    DS      B18      Cast-Out Class
ASCPDRF    DS      B18      Get Exclusive
ASCPDRF    DS      B18      Get Shared
ASCPDRF    DS      B18      WaitFor Exclusive
ASCPDRF    DS      B18      WaitFor Shared

```

```

ASCPCORE DS B18 Release Exclusive
ASCPCORS DS B18 Release Shared
*
ASChPL Equ *-ASChP

```

ADARUN Parameter Value Section: PARM

This selectable detail section may appear in initialization records or whenever an ADARUN parameter is changed while the nucleus is running (subtypes 1 and 4 in the header section). It will not be generated for interval or termination records (subtypes 2 and 3). This section has a fixed-length portion containing most parameters, followed by variable-length areas for parameters capable of multiple values or lists of values.

Where possible, the individual field names are formed by prefixing the shortest allowable form of the parameter with ASP. In general, the SMF record will report character parameters in EBCDIC and numeric parameters in binary.

Parameters with limited enumerated values (YES or NO, for example) are reported in 1-byte fields if the possible values are unambiguous. Otherwise, the field length is that used by the nucleus, usually 4 bytes.

Here are some sample entries:

```

ASParm DSect ,
ASPAO DS C11 AOslog
ASPARE DS B14 ARExclude Offset to file table
ASPARMN DS C116 ARMname
ASPASS DS C11 ASSocache
ASPASY DS C11 ASYtvs
. . .
ASPVI DS C11 Vista
ASPV64B DS C11 V64Bit
ASPWO DS C11 Workcache
ASParmV DS 0B Begin variable part
ASParmL Equ *-ASParm Length of ID section

ASParm DSect , ADARUN Parameters
*
* Adabas Nucleus Parameters
*
ASPAO DS C11 AOslog
ASPARE DS B14 ARExclude Offset to file table
ASPARMN DS C116 ARMname
ASPASS DS C11 ASSocache
ASPASY DS C11 ASYtvs
. . .
ASPVI DS C11 VIsta

```

The ADARUN AREXCLUDE parameter is a variable length list of values. The base parameter entry will be an offset from the beginning of the detail section to the table of values. The table is a 4-byte inclusive field followed by 4-byte file numbers. A separate DSECT maps the AREXCLUDE file exclusion table:

```

ASPAFE DSect , ARM File Exclusion Table
ASPAFEN DS B14 Inclusive length of table
ASPAFEF DS 0B14 First file number entry

```

I/O by DD Name Section: IODD

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). I/O by DD data is derived from data presented at nucleus shutdown. There is one instance for each DD statement administered by the nucleus. You should examine the DD name and not rely on any observed order of the instances. An Adabas nucleus may open and close the same DD name multiple times. Multiple uses of a DD name are summed.

You might see this many DD statements in a single nucleus:

Statement Type	Number of DD Statements
ASSO	99
CLOG	8
DATA	99
ECS	1
PLOG	8
RLOG	1
WORK	2

```

ASIODD  DSect ,           I/O Activity by DD
ASIODDNm DS   C18         DD Name
ASIODDRd DS   B18         Reads
ASIODDWt DS   B18         Writes
*
ASIODDL  Equ   *-ASIODD

```

Storage Pool Section: STG

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Storage pool statistics are derived from statistics presented at nucleus shutdown or in response to a DRES operator command. There is one instance for each storage pool with a non-zero size. Be sure to examine the pool name and not rely on any observed order of the instances.

Storage pool statistics are reported two ways: in bytes and also in units such as a user might specify as an ADARUN parameter, for example, NC. When the units are bytes, the two sets of statistics are the same.

Normally an interval record would show the change from the previous interval, but that isn't meaningful for storage pools. Thus the interval and termination record subtypes all reflect total usage for the nucleus session.

```

ASStg  DSect ,           Storage Pool Usage
ASStgNm DS   C14         Storage pool name
ASStgBSz DS   B18         Size in bytes
ASStgBHW DS   B18         High water mark in bytes
ASStgUSz DS   B18         Size in units from ADARUN parameter
ASStgUHW DS   B18         High water mark in ADARUN units
*
ASStgL  Equ   *-ASStg   Length of Storage Pool section

```

These are the possible storage pools:

Storage Pools		
Pool Name	ADARUN Parameter	Description
AB	NAB	Attached buffers
CQ	NC	Command queue
DUQ	LDEUQP	Unique descriptor
FI	LFP	Internal format buffers
HQ	NH	Hold queue
PLIO	NPLOGBUFFERS	PLOG I/O buffers
REDO	LRDP	Deferred publishing
RPL	LRPL	Replication pool
SC	LCP	Security information
TBI	LI	ISN table
TBS	NQ	Sequential command table
UQ	NU	User queue element
UQF	NU	User queue file elements
WKIO	NWORK1BUFFERS	Work I/O buffers
WORK	LWP	Work
XID	NU	Transaction ID

Thread Activity Section: THRD

This selectable detail section may appear in interval or termination records (subtypes 2 and 3). Thread activity data is derived from data presented at nucleus shutdown or in response to a DTH operator command. The ADARUN parameter NTHREAD defines the number of user threads for the nucleus session. There is one instance for each defined user thread. The thread number is implied by the sequence number of the instance.

```

ASthrd  DSect ,          Thread Activity
ASthrdCt DS      B18      Number of commands executed in this +
                          thread
*
ASthrdL  Equ      *-ASthrd
    
```

ASMFREC Mapping Macro

Use the ASMFREC macro to generate the latest SMF record DSECTs. The ASMFREC macro will always generate the header and self-defining section DSECT. Detail section DSECTs will be generated as specified. The header and self-defining sections are mapped by a single DSECT. Each detail section is mapped by its own DSECT. The syntax of the ASMFREC macro is:


```

label    ASMFREC    Prefix={AS | prefix},
                               Detail={All | (type [,type]...)},
                               Title = {'Adabas SMF Records' | 'string'}
    
```

Prefix

Specify a character string to be used as the initial characters for all DSECT and field names. The default is Prefix=AS.

Detail

Identify which detail section DSECTs are to be included in the expansion. "All" is the default and will include all detail sections. Alternatively, a comma-delimited list of types (enclosed in parentheses) can be specified; only the types specified will be included. The valid types are shown in the following table. A null value (Detail=) will inhibit all detail section DSECTs.

ASMFREC Macro Specification	Detail Section Description	ASMFREC DSECT Name Produced by the ASMFREC Macro ¹
CMD	Adabas command activity	xxCmd
CSHB	Adabas global cache activity by block type	xxChB
CSHF	Adabas global cache activity by Adabas file number	xxChF
CSHG	Adabas global cache activity	xxChG
CSHP	Adabas Parallel Services cache activity	xxChP
FILE	Adabas file activity	xxFile
LOCK	Adabas global lock activity	xxLok
MSGB	Adabas internucleus messaging control block activity	xxMsgB
MSGC	Adabas internucleus messaging counts	xxMsgC
MSGH	Adabas internucleus messaging service time histogram	xxMsgH
PARAM	ADARUN parameter values	xxParm
IODD	I/O by DD name	xxIODD
STG	Storage pool	xxStg
THRD	Thread activity	xxThrd
USER	User-defined	user-defined

1. Where *xx* is the prefix specified in the ASMFREC macro.

Title

If the Title default "Adabas SMF Record " or another quoted string is specified, an assembler Title statement is generated before the header section DSECT. A null value (Title=) for this operand will inhibit a title in the DSECT.

SMF User Exit

You can provide a user exit if you want to add a detail section to the Adabas SMF record. The user exit is a separate load module whose name must be provided in the ADARUN UEXSMF parameter. For complete information about the user exit, read *SMF User Exit*, in the *Adabas User, Hyperdescriptor, Collation Descriptor, and SMF Exits Manual*.

IBM Type 89 SMF Records

An Adabas nucleus can register with z/OS to have CPU usage statistics included in IBM type 89 SMF records. These records are described in *z/OS MVS System Management Facilities (SMF)*, IBM document SA22-7630.

To activate type 89 recording for Adabas, specify ADARUN parameters SMF=YES and SMF89=YES. During initialization Adabas will register the nucleus address space with z/OS SMF and have its CPU statistics included in subtype 1 of the type 89 records. The address space is deregistered at nucleus termination. Each Adabas nucleus appears as a separate type 89 entry.

The type 89 entries include CPU usage and a number of descriptive registration parameters. Adabas nuclei use these descriptive fields in type 89 entries as follows:

SMF Type 89 Descriptive Fields				
Name	Length	Format	Description	Value
SMF89UPO	16	EBCDIC	Product owner or vendor name	SOFTWARE AG
SMF89UPN	16	EBCDIC	Product name	ADABAS
SMF89UPV	8	EBCDIC	Product version	The eight-byte product version has two-byte numeric values for the Adabas version, release, SM level, and cumulative level.
SMF89UPQ	8	Binary	Product qualifier	<p>The product qualifier is a seven-byte string that may be used to distinguish among several nucleus instances. It contains a series of binary fields:</p> <p>SVC (1 byte) DBID (4 bytes) NucID (2 bytes)</p> <p>Use both the SVC and DBID to identify instances of Adabas Cluster or Parallel Service nuclei for the same database on any one system.</p>
SMF89UPI	8	EBCDIC	Product ID	<p>The product ID is a string of up to eight single characters to show what add-on products are being used. The characters may appear in any order:</p> <p>C (Adabas Cluster Services) D (Adabas Delta Save) F (Adabas Fastpath) M (Adabas Review) P (Adabas Parallel Services) R (Event Replicator for Adabas) S (Adabas Cache Facility) T (Adabas Transaction Manager) U (Adabas Security) V (Adabas Vista)</p> <p>C, P, and R are mutually exclusive. D, F, M, T and V are exclusive with R.</p>